



Red Hat OpenStack Platform 16.1

High Availability for Compute Instances

Configure High Availability for Compute Instances

Red Hat OpenStack Platform 16.1 High Availability for Compute Instances

Configure High Availability for Compute Instances

OpenStack Team
rhos-docs@redhat.com

Legal Notice

Copyright © 2021 Red Hat, Inc.

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

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

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

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

Abstract

This guide describes how to manage Instance High Availability (Instance HA). With Instance HA, Red Hat OpenStack Platform (RHOSP) can automatically evacuate and re-create instances on a different Compute node when a Compute node fails.

Table of Contents

MAKING OPEN SOURCE MORE INCLUSIVE	3
PROVIDING FEEDBACK ON RED HAT DOCUMENTATION	4
CHAPTER 1. HOW INSTANCE HA WORKS	5
CHAPTER 2. PLANNING YOUR INSTANCE HA DEPLOYMENT	6
CHAPTER 3. INSTALLING AND CONFIGURING INSTANCE HA	7
3.1. CONFIGURING THE INSTANCE HA ROLE, FLAVOR, AND PROFILE	7
3.2. ENABLING FENCING	8
3.3. DEPLOYING THE OVERCLOUD	9
3.4. TESTING INSTANCE HA EVACUATION	9
CHAPTER 4. DESIGNATING SPECIFIC INSTANCES TO EVACUATE	11
CHAPTER 5. PERFORMING MAINTENANCE ON COMPUTE NODES AND CONTROLLER NODES WITH INSTANCE HA	12
CHAPTER 6. SHUTTING DOWN AND STARTING UP THE UNDERCLOUD AND OVERCLOUD WITH INSTANCE HA	14
6.1. PREREQUISITES	14
6.2. UNDERCLOUD AND OVERCLOUD SHUTDOWN ORDER	14
6.3. SHUTTING DOWN INSTANCES ON OVERCLOUD COMPUTE NODES	14
6.4. STOPPING INSTANCE HA SERVICES ON OVERCLOUD COMPUTE NODES	15
6.5. SHUTTING DOWN COMPUTE NODES	15
6.6. STOPPING SERVICES ON CONTROLLER NODES	16
6.7. SHUTTING DOWN CEPH STORAGE NODES	17
6.8. SHUTTING DOWN CONTROLLER NODES	18
6.9. SHUTTING DOWN THE UNDERCLOUD	18
6.10. PERFORMING SYSTEM MAINTENANCE	18
6.11. UNDERCLOUD AND OVERCLOUD STARTUP ORDER	18
6.12. STARTING THE UNDERCLOUD	19
6.13. STARTING CONTROLLER NODES	19
6.14. STARTING CEPH STORAGE NODES	20
6.15. STARTING THE PACEMAKER CLUSTER ON CONTROLLER NODES	21
6.16. STARTING COMPUTE NODES	21
6.17. STARTING INSTANCE HA SERVICES ON OVERCLOUD COMPUTE NODES	22
6.18. STARTING INSTANCES ON OVERCLOUD COMPUTE NODES	23

MAKING OPEN SOURCE MORE INCLUSIVE

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

PROVIDING FEEDBACK ON RED HAT DOCUMENTATION

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

Using the Direct Documentation Feedback (DDF) function

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

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

CHAPTER 1. HOW INSTANCE HA WORKS

When a Compute node fails, the Instance High Availability (HA) tool evacuates and re-creates the instances on a different Compute node.

Instance HA uses the following resource agents:

Agent name	Name inside cluster	Role
fence_compute	fence-nova	Marks a Compute node for evacuation when the node becomes unavailable.
NovaEvacuate	nova-evacuate	Evacuates instances from failed nodes. This agent runs on one of the Controller nodes.
Dummy	compute-undefence-trigger	Releases a fenced node and enables the node to run instances again.

The following events occur when a Compute node fails and triggers Instance HA:

1. At the time of failure, the **IPMI** agent performs first-layer fencing, which includes physically resetting the node to ensure that it shuts down and preventing data corruption or multiple identical instances on the overcloud. When the node is offline, it is considered fenced.
2. After the physical IPMI fencing, the **fence-nova** agent automatically performs second-layer fencing and marks the fenced node with the **"evacuate=yes"** cluster per-node attribute by running the following command:

```
$ attrd_updater -n evacuate -A name="evacuate" host="FAILEDHOST" value="yes"
```

FAILEDHOST is the name of the failed Compute node.

3. The **nova-evacuate** agent continually runs in the background and periodically checks the cluster for nodes with the **"evacuate=yes"** attribute. When **nova-evacuate** detects that the fenced node contains this attribute, the agent starts evacuating the node. The evacuation process is similar to the manual instance evacuation process that you can perform at any time. For more information about instance evacuation, see [Evacuating an instance](#).
4. When the failed node restarts after the IPMI reset, the **nova-compute** process on that node also starts automatically. Because the node was previously fenced, it does not run any new instances until Pacemaker un-fences the node.
5. When Pacemaker detects that the Compute node is online, it starts the **compute-undefence-trigger** resource agent on the node, which releases the node and so that it can run instances again.

Instance HA works with shared storage or local storage environments, which means that evacuated instances maintain the same network configuration, such as static IP and floating IP. The re-created instances also maintain the same characteristics inside the new Compute node.

CHAPTER 2. PLANNING YOUR INSTANCE HA DEPLOYMENT

Before you deploy Instance HA, review the following considerations:

- Compute node host names and Pacemaker remote resource names must comply with the W3C naming conventions. For more information, see [Declaring Namespaces](#) and [Names and Tokens](#) in the W3C documentation.
- Typically, Instance HA requires that you configure shared storage for disk images of instances. Therefore, if you attempt to use the **no-shared-storage** option, you might receive an **InvalidSharedStorage** error during evacuation, and the instances will not start on another Compute node.
However, if all your instances are configured to boot from an OpenStack Block Storage (**cinder**) volume, you do not need to configure shared storage for the disk image of the instances, and you can evacuate all instances using the **no-shared-storage** option.

During evacuation, if your instances are configured to boot from a Block Storage volume, any evacuated instances boot from the same volume on another Compute node. Therefore, the evacuated instances immediately restart their jobs because the OS image and the application data are stored on the OpenStack Block Storage volume.

- If you deploy Instance HA in a Spine-Leaf environment, you must define a single **internal_api** network for the Controller and Compute nodes. You can then define a subnet for each leaf. For more information about configuring Spine-Leaf networks, see [Creating a roles data file](#) in the *Spine Leaf Networking* guide.
- From Red Hat OpenStack Platform 13 and later, you use director to upgrade Instance HA as a part of the overcloud upgrade. For more information about upgrading the overcloud, see [Keeping Red Hat OpenStack Platform Updated](#) guide.
- Disabling Instance HA with the director after installation is not supported. For a workaround to manually remove Instance HA components from your deployment, see the article [How can I remove Instance HA components from the controller nodes?](#) .



IMPORTANT

This workaround is not verified for production environments. You must verify the procedure in a test environment before you implement it in a production environment.

CHAPTER 3. INSTALLING AND CONFIGURING INSTANCE HA

Red Hat OpenStack Platform (RHOSP) director deploys Instance High Availability (HA). However, you must perform additional steps to configure a new Instance HA deployment on a new overcloud. After you complete the steps, Instance HA will run on a subset of Compute nodes with a custom role.



IMPORTANT

To enable instance HA in a different environment, such as an existing overcloud that uses standard or custom roles, perform only the procedures that are relevant to your deployment and adapt your templates accordingly.

For general information about deploying the overcloud, see [Director Installation and Usage](#) guide. For more information on custom roles, see [Composable Services and Custom Roles](#) in the *Advanced Overcloud Customization* guide.

3.1. CONFIGURING THE INSTANCE HA ROLE, FLAVOR, AND PROFILE

Procedure

You can modify the example file and role names in the following procedure according to your environment.

1. Add the **ComputeInstanceHA** role to your **roles-data.yaml** file and regenerate the file.

```
$ openstack overcloud roles generate -o ~/my_roles_data.yaml Controller Compute
ComputeInstanceHA
```

The **ComputeInstanceHA** role includes all the services in the default **Compute** role, the **ComputeInstanceHA** services, and the **PacemakerRemote** services. For general information about custom roles and about the **roles-data.yaml** file, see [Roles](#) in the *Advanced Overcloud Customization* guide.

2. Create the **compute-instance-ha** flavor to tag the Compute nodes to manage with Instance HA.

```
$ source ~/stackrc
$ openstack flavor create --id auto --ram 6144 --disk 40 --vcpus 4 compute-instance-ha
$ openstack flavor set --property "cpu_arch"="x86_64" --property
"capabilities:boot_option"="local" --property "capabilities:profile"="compute-instance-ha"
compute-instance-ha
$ openstack flavor set --property resources:VCPU=0 --property resources:MEMORY_MB=0 -
-property resources:DISK_GB=0 --property resources:CUSTOM_BAREMETAL=1 compute-
instance-ha
```

3. Tag each Compute node that you want to manage with Instance HA with the **compute-instance-ha** profile, and replace **<NODE UUID>** with the actual UUID:

```
$ openstack baremetal node set --property capabilities='profile:compute-instance-
ha,boot_option:local' <NODE UUID>
```

4. Map the **ComputeInstanceHA** role to the **compute-instance-ha** flavor by creating an environment file with the following parameter:

■

```
parameter_defaults:
  OvercloudComputeInstanceHAFlavor: compute-instance-ha
```

3.2. ENABLING FENCING

Enable fencing on all Controller and Compute nodes in the overcloud by creating an environment file with fencing information.

Procedure

1. Create the environment file in an accessible location, such as `~/templates`, and include the following content:

```
parameter_defaults:
  EnableFencing: true
  FencingConfig:
    devices:
      - agent: fence_ipmilan
        host_mac: 00:ec:ad:cb:3c:c7
        params:
          login: admin
          ipaddr: 192.168.24.1
          ipport: 6230
          passwd: password
          lanplus: 1
      - agent: fence_ipmilan
        host_mac: 00:ec:ad:cb:3c:cb
        params:
          login: admin
          ipaddr: 192.168.24.1
          ipport: 6231
          passwd: password
          lanplus: 1
      - agent: fence_ipmilan
        host_mac: 00:ec:ad:cb:3c:cf
        params:
          login: admin
          ipaddr: 192.168.24.1
          ipport: 6232
          passwd: password
          lanplus: 1
      - agent: fence_ipmilan
        host_mac: 00:ec:ad:cb:3c:d3
        params:
          login: admin
          ipaddr: 192.168.24.1
          ipport: 6233
          passwd: password
          lanplus: 1
      - agent: fence_ipmilan
        host_mac: 00:ec:ad:cb:3c:d7
        params:
          login: admin
          ipaddr: 192.168.24.1
```

```

ipport: 6234
passwd: password
lanplus: 1

```

For more information about fencing and STONITH configuration, see [Fencing Controller Nodes with STONITH](#) in the *High Availability Deployment and Usage* guide.

- Instance HA uses shared storage by default. If shared storage is not configured for your Compute instance, add the following parameter to the environment file that you created:

```

parameter_defaults:
  ExtraConfig:
    tripleo::instanceha::no_shared_storage: true

```

For information about how to boot from an OpenStack Block Storage (cinder) volume instead of configuring shared storage for the disk image of an instance, see [Chapter 2, Planning your Instance HA deployment](#).

3.3. DEPLOYING THE OVERCLOUD

You can configure Instance HA for your overcloud at any time after you create the undercloud. If you already deployed the overcloud, you must rerun the **openstack overcloud deploy** command with the additional Instance HA files you created.

Procedure

- Use the **openstack overcloud deploy** command with the **-e** option for each environment file that you created and with the *compute-instanceha.yaml* environment file. Replace **<FLAVOR_ENV_FILE>** and **<FENCING_ENV_FILE>** with the appropriate file names in your environment:

```

$ openstack overcloud deploy --templates \
  -e <FLAVOR_ENV_FILE> \
  -e <FENCING_ENV_FILE> \
  -e /usr/share/openstack-tripleo-heat-templates/environments/compute-instanceha.yaml

```



NOTE

- Do not modify the **compute-instanceha.yaml** environment file.
- Include the full path to each environment file that you want to include in the overcloud deployment.

Result

After the deployment is complete, each Compute node includes a **STONITH** device and a **GuestNode** service.

3.4. TESTING INSTANCE HA EVACUATION

**WARNING**

The following procedure involves deliberately crashing a Compute node, which triggers the automated evacuation of instances with Instance HA.

Procedure

1. Start one or more instances on the overcloud.

```
stack@director $ . overcloudrc
stack@director $ openstack server create --image cirros --flavor 2 test-failover
stack@director $ openstack server list -c Name -c Status
```

2. Log in to the Compute node that hosts the instances and change to the **root** user. Replace **compute-n** with the name of the Compute node:

```
stack@director $ . stackrc
stack@director $ ssh -l heat-admin compute-n
heat-admin@compute-n $ su -
```

3. Crash the Compute node.

```
root@compute-n $ echo c > /proc/sysrq-trigger
```

4. Wait a few minutes for the node to restart, and then verify that the instances from the Compute node that you crash are re-created on another Compute node:

```
stack@director $ openstack server list -c Name -c Status
stack@director $ openstack compute service list
```

CHAPTER 4. DESIGNATING SPECIFIC INSTANCES TO EVACUATE

By default, Instance HA evacuates all instances from a failed node. You can also configure Instance HA to only evacuate instances with specific images or flavors.

Prerequisites

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

```
$ source ~/overcloudrc
```

Tagging an image to evacuate

- Tag an image and replace **IMAGE_ID** with the ID of the image that you want to evacuate:

```
(overcloud) $ openstack image set --tag evacuable IMAGE_ID
```

Tagging a flavor to evacuate

- Tag a flavor and replace **FLAVOR_ID** with the ID of the flavor that you want to evacuate:

```
(overcloud) $ openstack flavor set --property evacuable=true FLAVOR_ID
```

CHAPTER 5. PERFORMING MAINTENANCE ON COMPUTE NODES AND CONTROLLER NODES WITH INSTANCE HA

If you must perform maintenance on a Compute node or a Controller node with Instance HA, you stop the node by setting it in **standby** mode and disabling the Pacemaker resources on the node. After you complete the maintenance work, you start the node and check that the Pacemaker resources are healthy.

Prerequisites

- A running overcloud with Instance HA enabled

Procedure

1. Log in to a Controller node and stop the Compute or Controller node:

```
# pcs node standby <node UUID>
```

You must log in to a different node from the node you want to stop.

2. Disable the Pacemaker resources on the node:

```
# pcs resource disable <ocf::pacemaker:remote on the node>
```

3. Perform any maintenance work on the node.
4. Restore the IPMI connection and start the node. Wait until the node is ready before proceeding.
5. Enable the Pacemaker resources on the node and start the node:

```
# pcs resource enable <ocf::pacemaker:remote on the node>  
# pcs node unstandby <node UUID>
```

6. If you set the node to maintenance mode, source the credential file for your overcloud and unset the node from maintenance mode:

```
# source stackrc  
# openstack baremetal node maintenance unset <baremetal node UUID>
```

Verification

1. Check that the Pacemaker resources are active and healthy:

```
# pcs status
```

2. If any Pacemaker resources fail to start during the startup process, run the **pcs resource cleanup** command to reset the status and the fail count of the resource.
3. If you evacuated instances from a Compute node before you stopped the node, check that the instances are migrated to a different node:

```
# openstack server list --long  
# nova migration-list
```


-

CHAPTER 6. SHUTTING DOWN AND STARTING UP THE UNDERCLOUD AND OVERCLOUD WITH INSTANCE HA

If you must perform maintenance on the undercloud and overcloud, you must shut down and start up the undercloud and overcloud nodes in a specific order to ensure minimal issues when you start your overcloud.

6.1. PREREQUISITES

- A running undercloud and overcloud

6.2. UNDERCLOUD AND OVERCLOUD SHUTDOWN ORDER

To shut down the Red Hat OpenStack Platform environment, you must shut down the overcloud and undercloud in the following order:

1. Shut down instances on overcloud Compute nodes
2. Stop all instance HA services on overcloud Compute nodes
3. Shut down Compute nodes
4. Stop all high availability and OpenStack Platform services on Controller nodes
5. Shut down Ceph Storage nodes
6. Shut down Controller nodes
7. Shut down the undercloud

6.3. SHUTTING DOWN INSTANCES ON OVERCLOUD COMPUTE NODES

As a part of shutting down the Red Hat OpenStack Platform environment, shut down all instances on Compute nodes before shutting down the Compute nodes.

Prerequisites

- An overcloud with active Compute services

Procedure

1. Log in to the undercloud as the **stack** user.
2. Source the credentials file for your overcloud:

```
█ $ source ~/overcloudrc
```

3. View running instances in the overcloud:

```
█ $ openstack server list --all-projects
```

4. Stop each instance in the overcloud:

```
$ openstack server stop <INSTANCE>
```

Repeat this step for each instance until you stop all instances in the overcloud.

6.4. STOPPING INSTANCE HA SERVICES ON OVERCLOUD COMPUTE NODES

As a part of shutting down the Red Hat OpenStack Platform environment, you must shut down all Instance HA services on Compute nodes before stopping the instances and shutting down the Compute nodes.

Prerequisites

- An overcloud with active Compute services
- Instance HA is enabled on Compute nodes

Procedure

1. Log in as the **root** user to an overcloud node that runs Pacemaker.
2. Disable the Pacemaker remote resource on each Compute node:
 - a. Identify the Pacemaker Remote resource on Compute nodes:

```
# pcs resource status
```

These resources use the **ocf::pacemaker:remote** agent and are usually named after the Compute node host format, such as **overcloud-novacomputeiha-0**.

- b. Disable each Pacemaker Remote resource. The following example shows how to disable the resource for **overcloud-novacomputeiha-0**:

```
# pcs resource disable overcloud-novacomputeiha-0
```

3. Disable the Compute node STONITH devices:
 - a. Identify the Compute node STONITH devices:

```
# pcs stonith status
```

- b. Disable each Compute node STONITH device:

```
# pcs stonith disable <STONITH_DEVICE>
```

6.5. SHUTTING DOWN COMPUTE NODES

As a part of shutting down the Red Hat OpenStack Platform environment, log in to and shut down each Compute node.

Prerequisites

- Shut down all instances on the Compute nodes
- Stop instance HA services on the Compute nodes

Procedure

1. Log in as the **root** user to a Compute node.
2. Shut down the node:

```
# shutdown -h now
```

3. Perform these steps for each Compute node until you shut down all Compute nodes.

6.6. STOPPING SERVICES ON CONTROLLER NODES

As a part of shutting down the Red Hat OpenStack Platform environment, stop services on the Controller nodes before shutting down the nodes. This includes Pacemaker and systemd services.

Prerequisites

- An overcloud with active Pacemaker services

Procedure

1. Log in as the **root** user to a Controller node.
2. Stop the Pacemaker cluster.

```
# pcs cluster stop --all
```

This command stops the cluster on all nodes.

3. Wait until the Pacemaker services stop and check that the services stopped.
 - a. Check the Pacemaker status:

```
# pcs status
```

- b. Check that no Pacemaker services are running in Podman:

```
# podman ps --filter "name=.*-bundle.*"
```

4. Stop the Red Hat OpenStack Platform services:

```
# systemctl stop 'tripleo_*
```

5. Wait until the services stop and check that services are no longer running in Podman:

```
# podman ps
```

6.7. SHUTTING DOWN CEPH STORAGE NODES

As a part of shutting down the Red Hat OpenStack Platform environment, disable Ceph Storage services then log in to and shut down each Ceph Storage node.

Prerequisites

- A healthy Ceph Storage cluster
- Ceph MON services are running on standalone Ceph MON nodes or on Controller nodes

Procedure

1. Log in as the **root** user to a node that runs Ceph MON services, such as a Controller node or a standalone Ceph MON node.
2. Check the health of the cluster. In the following example, the **podman** command runs a status check within a Ceph MON container on a Controller node:

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

Ensure that the status is **HEALTH_OK**.

3. Set the **noout**, **norecover**, **norebalance**, **nobackfill**, **nodown**, and **pause** flags for the cluster. In the following example, the **podman** commands set these flags through a Ceph MON container on a Controller node:

```
# sudo podman exec -it ceph-mon-controller-0 ceph osd set noout
# sudo podman exec -it ceph-mon-controller-0 ceph osd set norecover
# sudo podman exec -it ceph-mon-controller-0 ceph osd set norebalance
# sudo podman exec -it ceph-mon-controller-0 ceph osd set nobackfill
# sudo podman exec -it ceph-mon-controller-0 ceph osd set nodown
# sudo podman exec -it ceph-mon-controller-0 ceph osd set pause
```

4. Shut down each Ceph Storage node:
 - a. Log in as the **root** user to a Ceph Storage node.
 - b. Shut down the node:

```
# shutdown -h now
```

- c. Perform these steps for each Ceph Storage node until you shut down all Ceph Storage nodes.

5. Shut down any standalone Ceph MON nodes:
 - a. Log in as the **root** user to a standalone Ceph MON node.
 - b. Shut down the node:

```
# shutdown -h now
```

- c. Perform these steps for each standalone Ceph MON node until you shut down all standalone Ceph MON nodes.

Additional resources

- ["What is the procedure to shutdown and bring up the entire ceph cluster?"](#)

6.8. SHUTTING DOWN CONTROLLER NODES

As a part of shutting down the Red Hat OpenStack Platform environment, log in to and shut down each Controller node.

Prerequisites

- Stop the Pacemaker cluster
- Stop all Red Hat OpenStack Platform services on the Controller nodes

Procedure

1. Log in as the **root** user to a Controller node.
2. Shut down the node:

```
█ # shutdown -h now
```

3. Perform these steps for each Controller node until you shut down all Controller nodes.

6.9. SHUTTING DOWN THE UNDERCLOUD

As a part of shutting down the Red Hat OpenStack Platform environment, log in to the undercloud node and shut down the undercloud.

Prerequisites

- A running undercloud

Procedure

1. Log in to the undercloud as the **stack** user.
2. Shut down the undercloud:

```
█ $ sudo shutdown -h now
```

6.10. PERFORMING SYSTEM MAINTENANCE

After you completely shut down the undercloud and overcloud, perform any maintenance to the systems in your environment and then start up the undercloud and overcloud.

6.11. UNDERCLOUD AND OVERCLOUD STARTUP ORDER

To start the Red Hat OpenStack Platform environment, you must start the undercloud and overcloud in the following order:

1. Start the undercloud

2. Start Controller nodes
3. Start Ceph Storage nodes
4. Start high availability services on the Controller nodes
5. Start Compute nodes
6. Start all instance HA services on overcloud Compute nodes
7. Start instances on overcloud Compute nodes

6.12. STARTING THE UNDERCLOUD

As a part of starting the Red Hat OpenStack Platform environment, power on the undercloud node, log in to the undercloud, and check the undercloud services.

Prerequisites

- A powered down undercloud

Procedure

1. Power on the undercloud and wait until the undercloud boots.

Verification

1. Log in to the undercloud as the **stack** user.
2. Check the services on the undercloud:

```
$ systemctl list-units 'tripleo_*
```

3. Source the credential file for your undercloud and run the validation command to check that all services and containers are active and healthy.

```
$ source stackrc  
$ openstack tripleo validator run --validation service-status --limit undercloud
```

Additional resources

- [Using the validation framework](#)

6.13. STARTING CONTROLLER NODES

As a part of starting the Red Hat OpenStack Platform environment, power on each Controller node and check the non-Pacemaker services on the node.

Prerequisites

- Powered down Controller nodes

Procedure

1. Power on each Controller node.

Verification

1. Log in to each Controller as the **root** user.
2. Check the services on the Controller node:

```
$ systemctl -t service
```

Only non-Pacemaker based services are running.

6.14. STARTING CEPH STORAGE NODES

As a part of starting the Red Hat OpenStack Platform environment, power on the Ceph MON and Ceph Storage nodes and enable Ceph Storage services.

Prerequisites

- A powered down Ceph Storage cluster
- Ceph MON services are enabled on powered down standalone Ceph MON nodes or on powered on Controller nodes

Procedure

1. If your environment has standalone Ceph MON nodes, power on each Ceph MON node.
2. Power on each Ceph Storage node.
3. Log in as the **root** user to a node that runs Ceph MON services, such as a Controller node or a standalone Ceph MON node.
4. Check the status of the cluster nodes. In the following example, the **podman** command runs a status check within a Ceph MON container on a Controller node:

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

Ensure that each node is powered on and connected.

5. Unset the **noout**, **norecover**, **norebalance**, **nobackfill**, **nodown** and **pause** flags for the cluster. In the following example, the **podman** commands unset these flags through a Ceph MON container on a Controller node:

```
# sudo podman exec -it ceph-mon-controller-0 ceph osd unset noout
# sudo podman exec -it ceph-mon-controller-0 ceph osd unset norecover
# sudo podman exec -it ceph-mon-controller-0 ceph osd unset norebalance
# sudo podman exec -it ceph-mon-controller-0 ceph osd unset nobackfill
# sudo podman exec -it ceph-mon-controller-0 ceph osd unset nodown
# sudo podman exec -it ceph-mon-controller-0 ceph osd unset pause
```

Verification

1. Check the health of the cluster. In the following example, the **podman** command runs a status check within a Ceph MON container on a Controller node:

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

Ensure the status is **HEALTH_OK**.

Additional resources

- ["What is the procedure to shutdown and bring up the entire ceph cluster?"](#)

6.15. STARTING THE PACEMAKER CLUSTER ON CONTROLLER NODES

As a part of starting the Red Hat OpenStack Platform environment, start the Pacemaker cluster on the Controller nodes. This process enables high availability services, such as the database, messaging, load balancing, and VIP assignment.

Prerequisites

- An overcloud with running Controller nodes

Procedure

1. Log in as the **root** user to a Controller node.
2. Start the Pacemaker cluster.

```
# pcs cluster start --all
```

This command starts the cluster on all nodes.

3. Wait until the Pacemaker services start and check that the services started.
 - a. Check the Pacemaker status:

```
# pcs status
```

- b. Check that the Pacemaker services are running in Podman:

```
# podman ps --filter "name=.*-bundle.*"
```

6.16. STARTING COMPUTE NODES

As a part of starting the Red Hat OpenStack Platform environment, power on each Compute node and check the services on the node.

Prerequisites

- Powered down Compute nodes

Procedure

1. Power on each Compute node.

Verification

1. Log in to each Compute as the **root** user.
2. Check the services on the Compute node:

```
$ systemctl -t service
```

6.17. STARTING INSTANCE HA SERVICES ON OVERCLOUD COMPUTE NODES

As a part of starting the Red Hat OpenStack Platform environment, start all Instance HA services on the Compute nodes.

Prerequisites

- An overcloud with running Compute nodes
- Instance HA is enabled on Compute nodes

Procedure

1. Log in as the **root** user to an overcloud node that runs Pacemaker.
2. Enable the STONITH device for a Compute node:

- a. Identify the Compute node STONITH device:

```
# pcs stonith status
```

- b. Clear any STONITH errors for the Compute node:

```
# pcs stonith confirm <COMPUTE_NODE>
```

This command returns the node to a clean STONITH state.

- c. Enable the Compute node STONITH device:

```
# pcs stonith enable <STONITH_DEVICE>
```

- d. Perform these steps for each Compute node with STONITH.

3. Enable the Pacemaker remote resource on each Compute node:

- a. Identify the Pacemaker remote resources on Compute nodes:

```
# pcs resource status
```

These resources use the **ocf::pacemaker:remote** agent and are usually named after the Compute node host format, such as **overcloud-novacomputeiha-0**.

- b. Enable each Pacemaker Remote resource. The following example shows how to enable the resource for **overcloud-novacomputeiha-0**:

```
# pcs resource enable overcloud-novacomputeiha-0
```

- c. Perform these steps for each Compute node with Pacemaker remote management.

6.18. STARTING INSTANCES ON OVERCLOUD COMPUTE NODES

As a part of starting the Red Hat OpenStack Platform environment, start the instances on on Compute nodes.

Prerequisites

- An active overcloud with active nodes

Procedure

1. Log in to the undercloud as the **stack** user.
2. Source the credentials file for your overcloud:

```
$ source ~/overcloudrc
```

3. View running instances in the overcloud:

```
$ openstack server list --all-projects
```

4. Start an instance in the overcloud:

```
$ openstack server start <INSTANCE>
```