

Replacement of Faulty Components on Server UCS C240 M4 - CPAR

Contents

[Introduction](#)

[Background Information](#)

[Abbreviations](#)

[Workflow of MoP](#)

[Prerequisites](#)

[Backup](#)

[Component RMA - Compute Node](#)

[Identify VMs Hosted in Compute Node](#)

[1. CPAR Application Shutdown](#)

[2. VM Snapshot Task](#)

[VM Snapshot](#)

[Graceful Power Off](#)

[Replace Faulty Component from Compute Node](#)

[Restore VMs](#)

[Recover Instance with Snapshot](#)

[Create and Assign Floating IP Address](#)

[Enable SSH](#)

[Establish SSH Session](#)

[Component RMA - OSD Compute Node](#)

[Identify VMs Hosted in OSD-Compute Node](#)

[1. CPAR Application Shutdown](#)

[2. VM Snapshot Task](#)

[VM Snapshot](#)

[Put CEPH in Maintenance Mode](#)

[Graceful Power Off](#)

[Replace Faulty Component from OSD-Compute Node](#)

[Move CEPH Out of Maintenance Mode](#)

[Restore VMs](#)

[Recover Instance with Snapshot](#)

[Component RMA - Controller Node](#)

[Pre-Check](#)

[Move Controller Cluster to Maintenance Mode](#)

[Replace Faulty Component from Controller Node](#)

[Power On Server](#)

Introduction

This document describes the steps required to replace faulty components mentioned here in a

Unified Computing System (UCS) server in an Ultra-M setup.

This procedure applies for an Openstack environment with the use of NEWTON version where ESC does not manage CPAR and CPAR is installed directly on the VM deployed on Openstack.

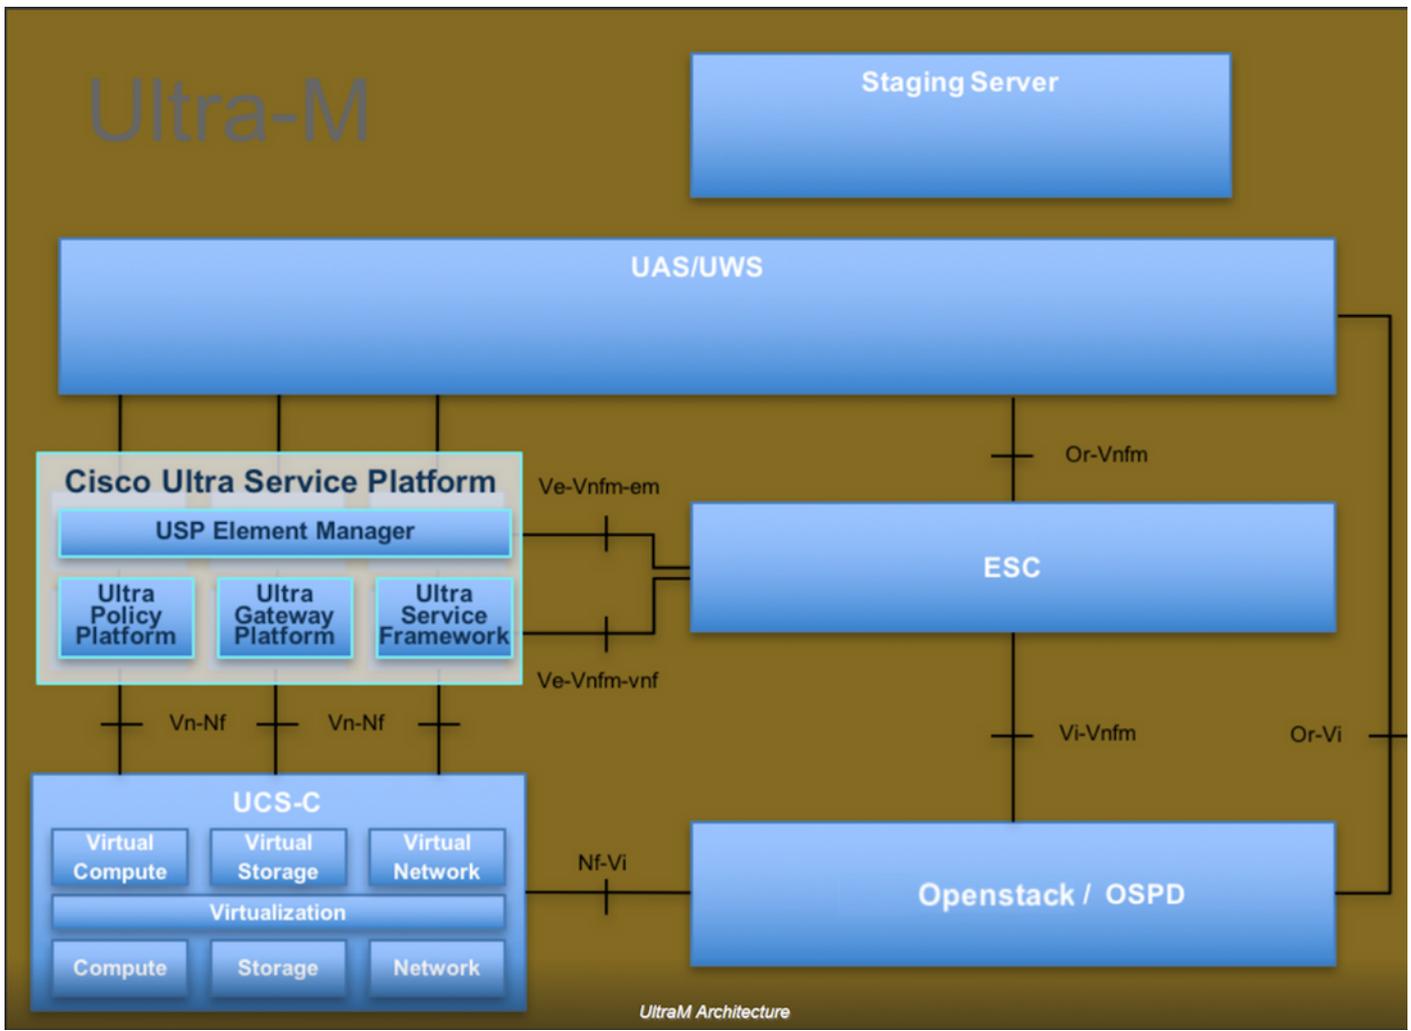
- Dual In-line Memory Module (DIMM) Replacement MOP
- FlexFlash Controller Failure
- Solid State Drive (SSD) Failure
- Trusted Platform Module (TPM) Failure
- Raid Cache Failure
- Raid Controller/ Hot Bus Adapter (HBA) Failure
- PCI Riser Failure
- PCIe adapter Intel X520 10G Failure
- Modular LAN-on Motherboard (MLOM) Failure
- Fan tray RMA
- CPU Failure

Background Information

Ultra-M is a pre-packaged and validated virtualized mobile packet core solution that is designed in order to simplify the deployment of VNFs. OpenStack is the Virtualized Infrastructure Manager (VIM) for Ultra-M and consists of these node types:

- Compute
- Object Storage Disk - Compute (OSD - Compute)
- Controller
- OpenStack Platform - Director (OSPD)

The high-level architecture of Ultra-M and the components involved are depicted in this image:



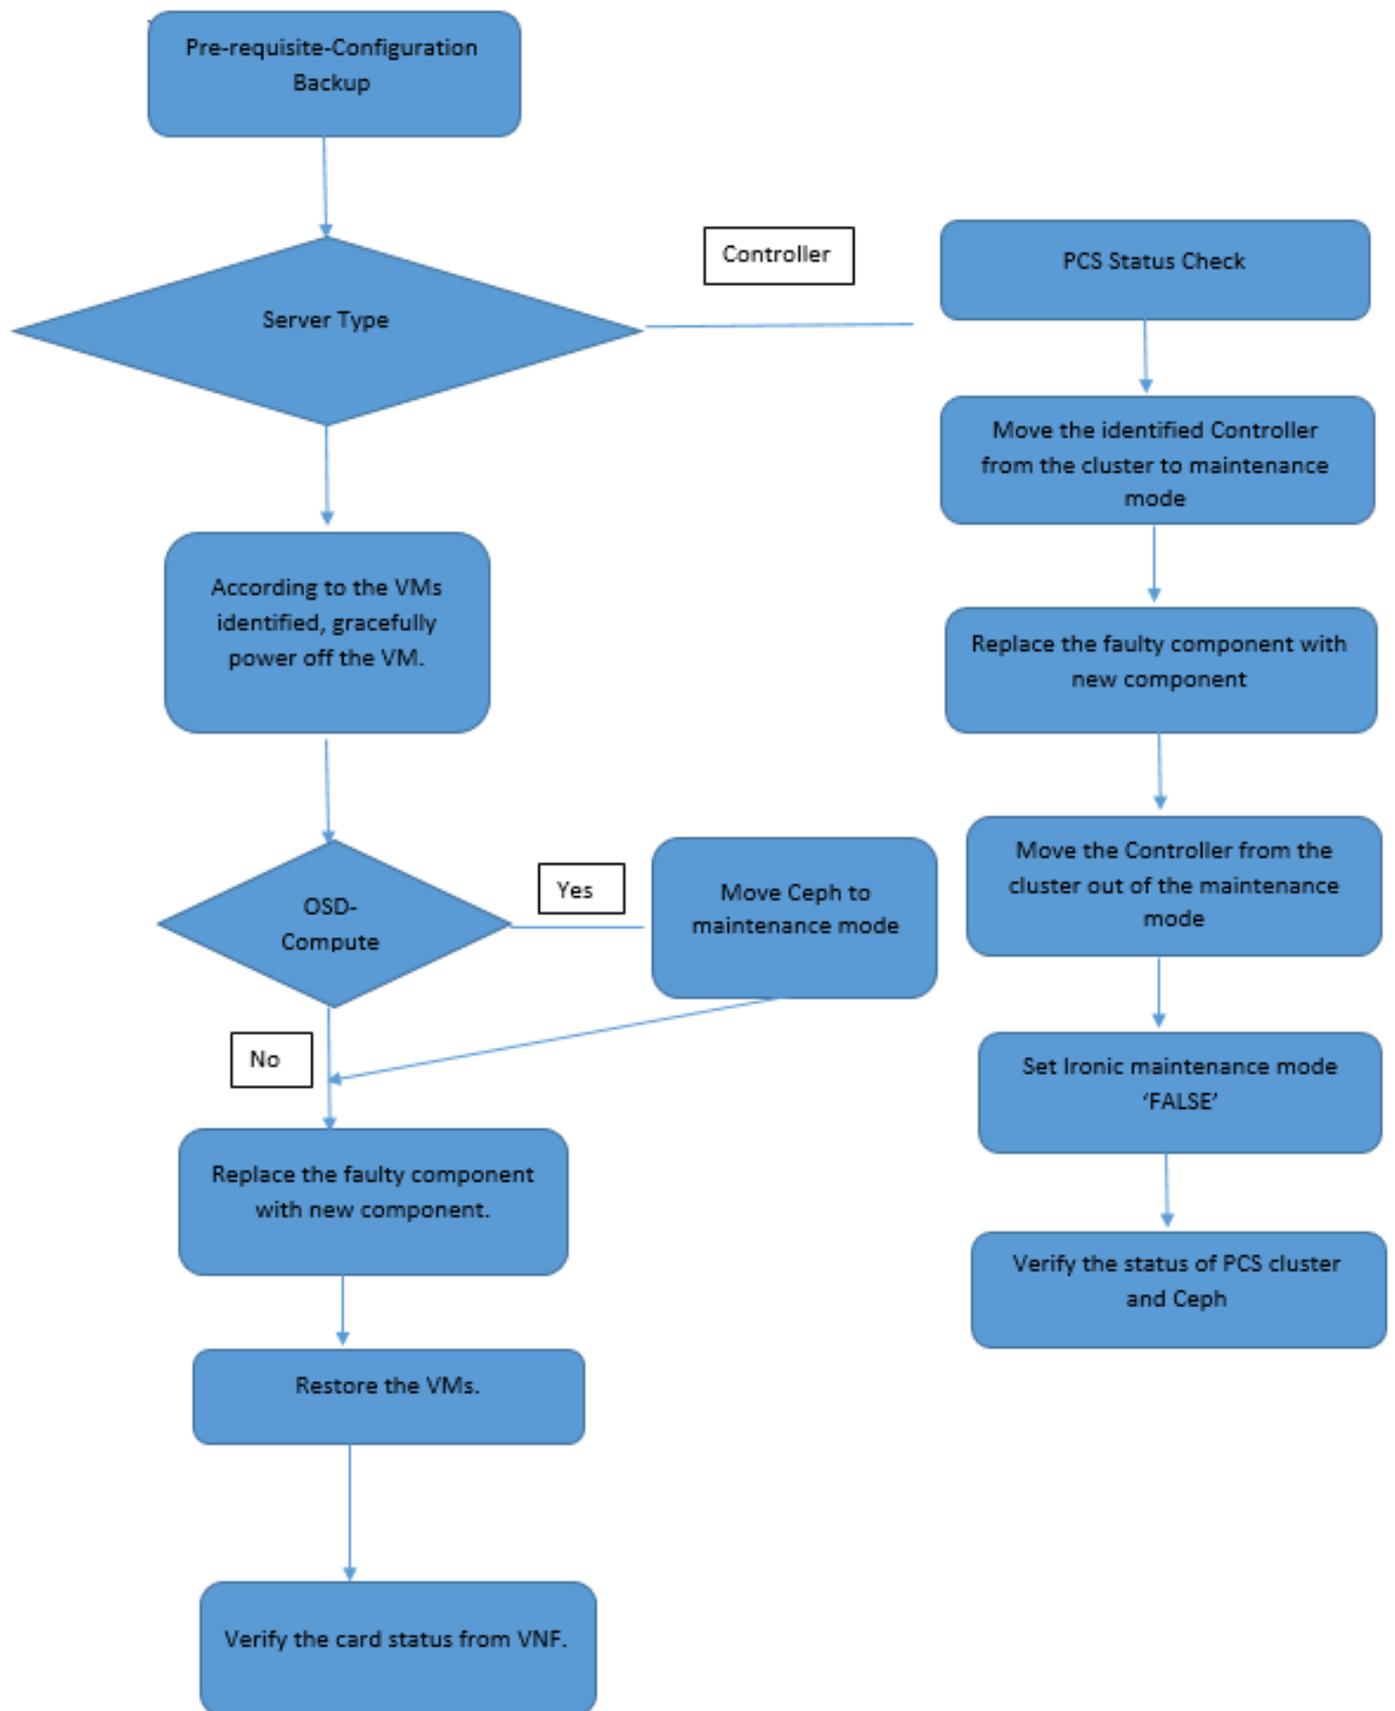
This document is intended for Cisco personnel who are familiar with Cisco Ultra-M platform and it details the steps required to be carried out at OpenStack and Redhat OS.

Note: Ultra M 5.1.x release is considered in order to define the procedures in this document.

Abbreviations

| | |
|------|--------------------------------|
| MoP | Method of Procedure |
| OSD | Object Storage Disks |
| OSPD | OpenStack Platform Director |
| HDD | Hard Disk Drive |
| SSD | Solid State Drive |
| VIM | Virtual Infrastructure Manager |
| VM | Virtual Machine |
| EM | Element Manager |
| UAS | Ultra Automation Services |
| UUID | Universally Unique Identifier |

Workflow of MoP



Prerequisites

Backup

Before you replace a faulty component, it is important to check the current state of your Red Hat

OpenStack Platform environment. It is recommended that you check the current state in order to avoid complications when the replacement process is on. It can be achieved by this flow of replacement.

In case of recovery, Cisco recommends to take a backup of the OSPD database with the use of these steps:

```
[root@director ~]# mysqldump --opt --all-databases > /root/undercloud-all-databases.sql
[root@director ~]# tar --xattrs -czf undercloud-backup-`date +%F`.tar.gz /root/undercloud-all-databases.sql
/etc/my.cnf.d/server.cnf /var/lib/glance/images /srv/node /home/stack
tar: Removing leading `/' from member names
```

This process ensures that a node can be replaced without affecting the availability of any instances. Also, it is recommended to back up the StarOS configuration especially if the compute/OSD-compute node to be replaced hosts the Control Function (CF) Virtual Machine (VM).

Note: If the Server is the Controller node, proceed to the section "", otherwise continue with the next section. Ensure that you have the snapshot of the instance so that you can restore the VM when needed. Follow the procedure on how to take snapshot of the VM.

Component RMA - Compute Node

Identify VMs Hosted in Compute Node

Identify the VMs that are hosted on the server.

```
[stack@al03-pod2-ospd ~]$ nova list --field name,host
```

| ID | Name | Host |
|--------------------------------------|---------------------------|----------------------------------|
| 46b4b9eb-a1a6-425d-b886-a0ba760e6114 | AAA-CPAR-testing-instance | pod2-stack-compute-4.localdomain |
| 3bc14173-876b-4d56-88e7-b890d67a4122 | aaa2-21 | pod2-stack-compute-3.localdomain |
| f404f6ad-34c8-4a5f-a757-14c8ed7fa30e | aaa21june | pod2-stack-compute-3.localdomain |

Note: In the output shown here, the first column corresponds to the UUID, the second column is the VM name and the third column is the hostname where the VM is present. The parameters from this output will be used in subsequent sections.

Backup: SNAPSHOT PROCESS

1. CPAR Application Shutdown

Step 1. Open any SSH client connected to the TMO Production network and connect to the CPAR instance.

It is important not to shutdown all 4 AAA instances within one site at the same time, do it in a one by one fashion.

Step 2. In order to shut down the CPAR application, run the command:

```
/opt/CSCOar/bin/arserver stop
```

A Message "Cisco Prime Access Registrar Server Agent shutdown complete." must show up.

Note: If a user left a CLI session open, the **arserver stop** command won't work and this message is displayed:

```
ERROR:      You cannot shut down Cisco Prime Access Registrar while the
            CLI is being used.  Current list of running
            CLI with process id is:
2903 /opt/CSCOar/bin/aregcmd -s
```

In this example, the highlighted process id 2903 needs to be terminated before CPAR can be stopped. If this is the case, terminate this process by running the command:

```
ERROR:      You cannot shut down Cisco Prime Access Registrar while the
            CLI is being used.  Current list of running
            CLI with process id is:
2903 /opt/CSCOar/bin/aregcmd -s
```

Then, repeat the step 1.

Step 3. In order to verify that the CPAR application was indeed shutdown, run the command:

```
ERROR:      You cannot shut down Cisco Prime Access Registrar while the
            CLI is being used.  Current list of running
            CLI with process id is:
2903 /opt/CSCOar/bin/aregcmd -s
```

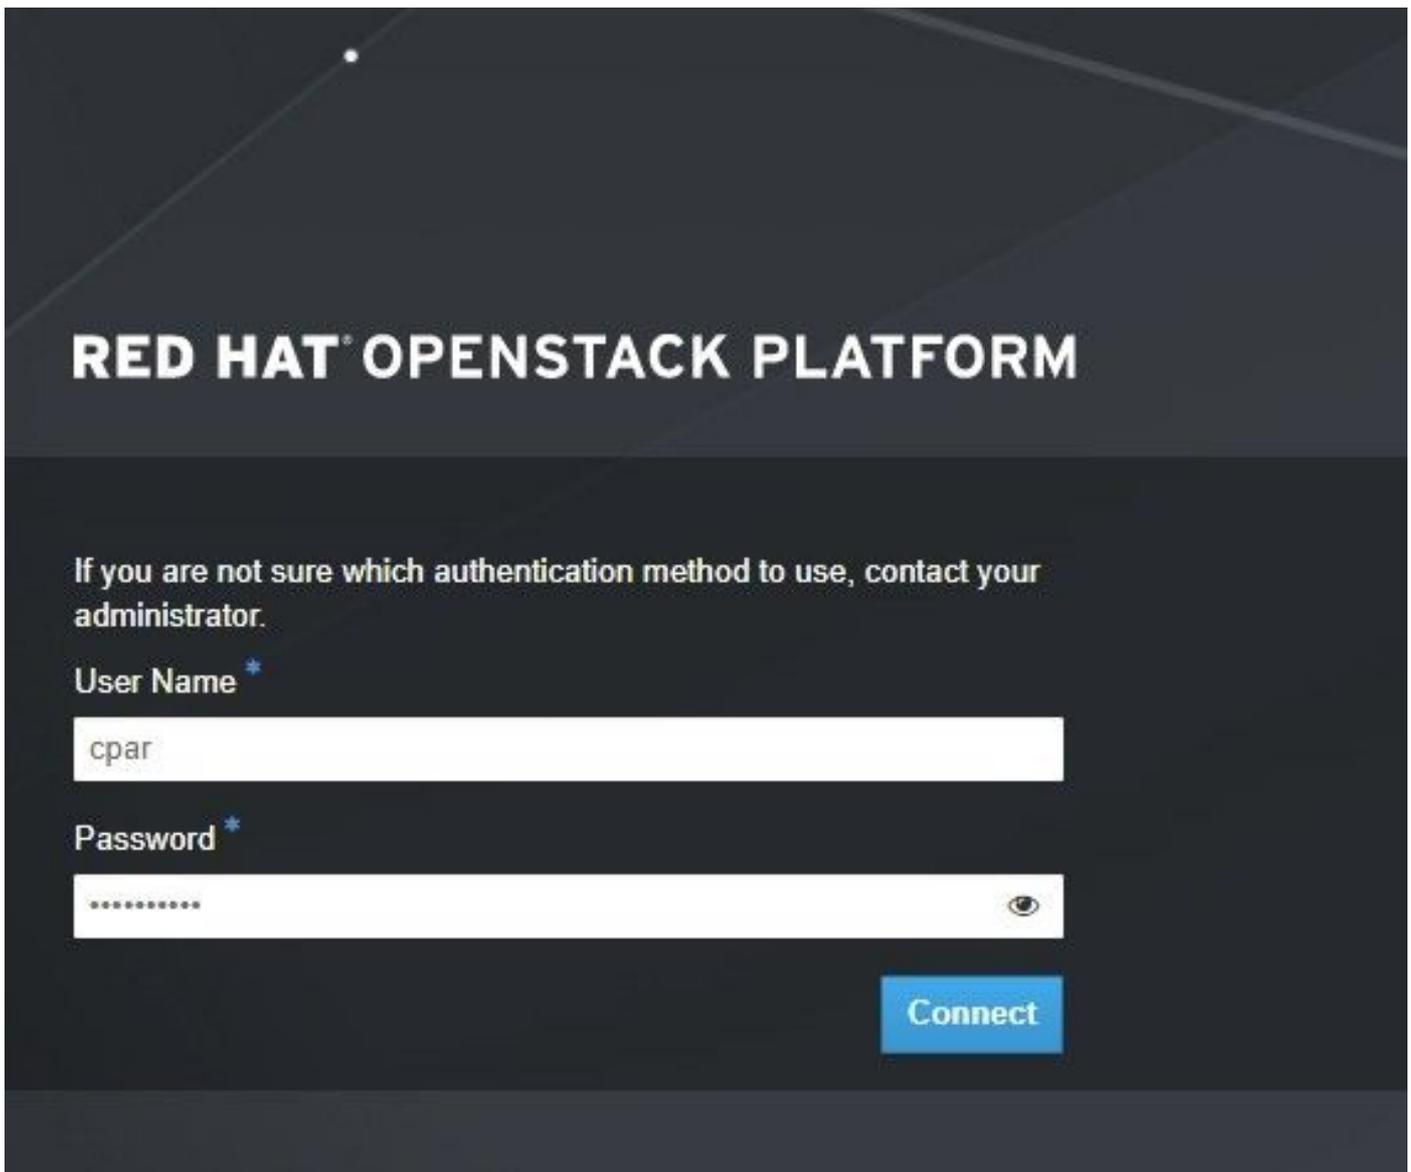
This messages must appear:

```
ERROR:      You cannot shut down Cisco Prime Access Registrar while the
            CLI is being used.  Current list of running
            CLI with process id is:
2903 /opt/CSCOar/bin/aregcmd -s
```

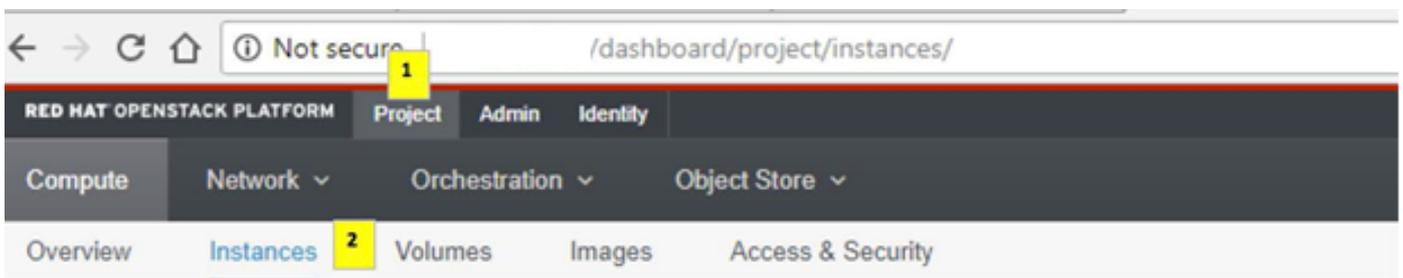
2. VM Snapshot Task

Step 1. Enter the Horizon GUI website that corresponds to the Site (City) currently being worked on.

When you access Horizon, this screen is observed.



Step 2. Navigate to **Project > Instances** as shown in this image.



If the user used was cpar, then only the 4 AAA instances appear in this menu.

Step 3. Shut down only one instance at a time, repeat the whole process in this document. In order to shutdown the VM, navigate to **Actions > Shut Off Instance** as shown in this image and confirm your selection..



Step 4. Validate that the instance was indeed shut down by checking the Status = **Shutoff** and Power State = **Shut Down** as shown in this image.

| Size | Key Pair | Status | Availability Zone | Task | Power State | Time since created | Actions |
|----------|----------|---------|-------------------|------|-------------|--------------------|----------------|
| AAA-CPAR | - | Shutoff | AZ-dalaaa09 | None | Shut Down | 3 months, 2 weeks | Start Instance |

This step ends the CPAR shutdown process.

VM Snapshot

Once the CPAR VMs are down, the snapshots can be taken in parallel, as they belong to independent computes.

The four QCOW2 files are created in parallel.

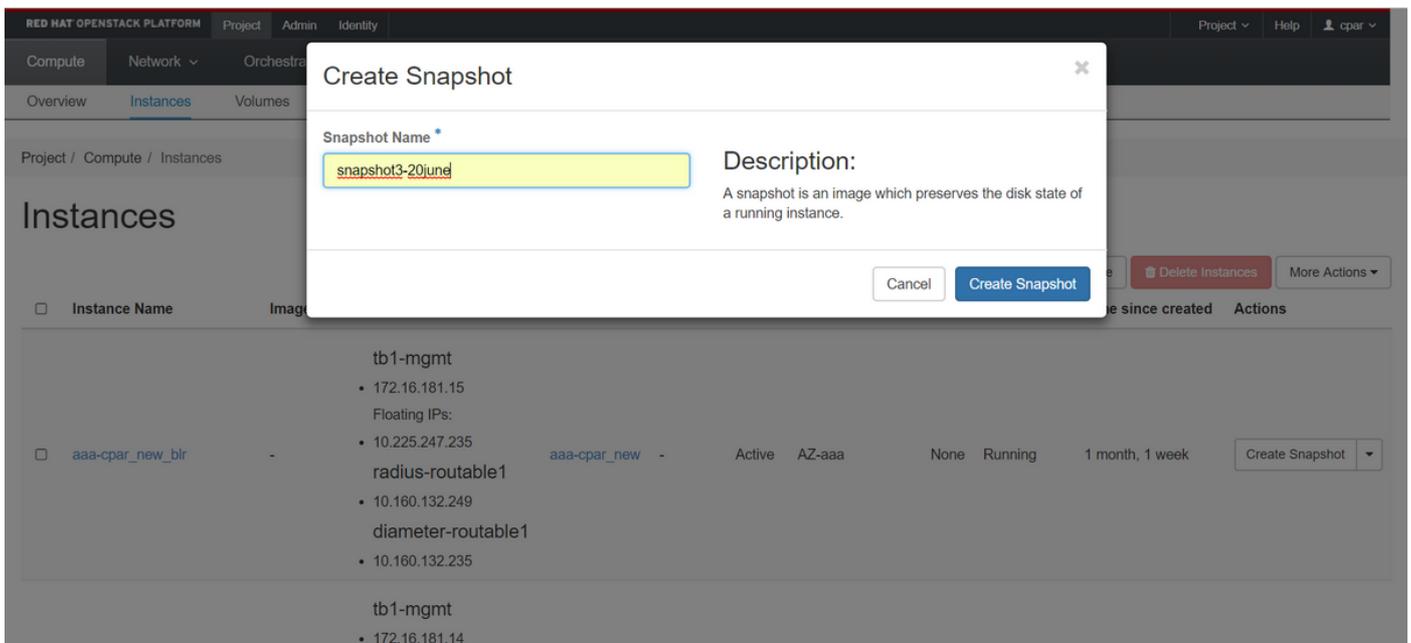
Take a snapshot of each AAA instance (25 minutes -1 hour) (25 minutes for instances that used a qcow image as a source and 1 hour for instances that user a raw image as a source)

1. Login to POD's Openstack's Horizon GUI.
2. Once logged in, navigate to **PROJECT > COMPUTE > INSTANCES** section on the top menu and look for the AAA instances as shown in this image.

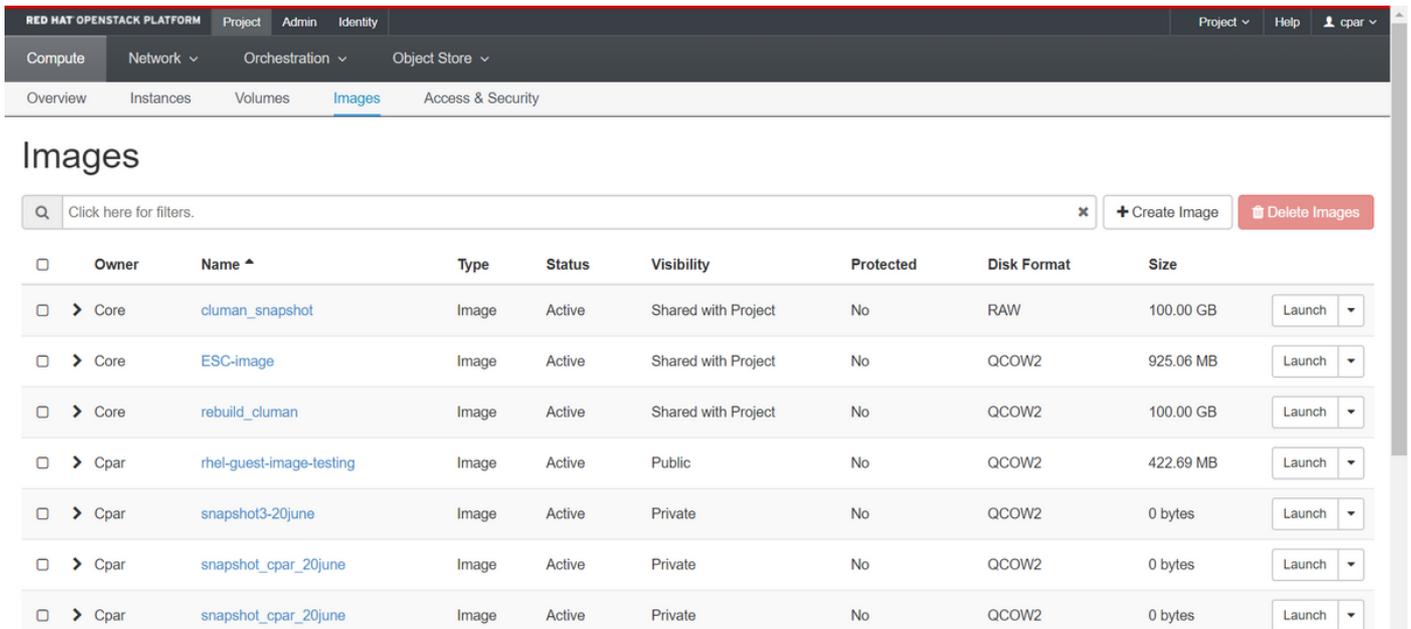
The screenshot shows the OpenStack Horizon 'Instances' page. The breadcrumb navigation is 'Project / Compute / Instances'. The main heading is 'Instances'. There are filters for 'Instance Name', 'Filter', 'Launch Instance', 'Delete Instances', and 'More Actions'. The table below shows the following instance:

| Instance Name | Image Name | IP Address | Size | Key Pair | Status | Availability Zone | Task | Power State | Time since created | Actions |
|------------------|------------|--|--------------|----------|--------|-------------------|------|-------------|--------------------|-----------------|
| aaa-cpar_new_blr | - | tb1-mgmt • 172.16.181.15 Floating IPs: • 10.225.247.235 radius-routable1 • 10.160.132.249 diameter-routable1 • 10.160.132.235 | aaa-cpar_new | - | Active | AZ-aaa | None | Running | 1 month, 1 week | Create Snapshot |

3. Click **Create Snapshot** in order to proceed with the snapshot creation (this needs to be executed on the corresponding AAA instance) as shown in this image.



4. Once the snapshot is executed, navigate to **Images** menu and verify that all finish and report no problems as shown in this image.



5. The next step is to download the snapshot on a QCOW2 format and transfer it to a remote entity, in case the OSPD is lost during this process. In order to achieve this, identify the snapshot by running the command **glance image-list** at OSPD level.

```
ERROR: You cannot shut down Cisco Prime Access Registrar while the
        CLI is being used. Current list of running
        CLI with process id is:
2903 /opt/CSCOar/bin/aregcmd -s
```

6. Once you identify the snapshot to download (the one marked in green), you can download it on a QCOW2 format with the command **glance image-download** as depicted here.

```
ERROR: You cannot shut down Cisco Prime Access Registrar while the
        CLI is being used. Current list of running
        CLI with process id is:
```

```
2903 /opt/CSCOar/bin/aregcmd -s
```

- The **&** sends the process to background. It can take some time to complete this action, once it is done, the image can be located at **/tmp** directory.
- On sending the process to background , if connectivity is lost, then the process is also stopped.
- Run the command **disown -h** so that in case SSH connection is lost, the process still runs and finishes on the OSPD.

7. Once the download process finishes, a compression process needs to be executed as that snapshot can be filled with ZEROES because of processes, tasks and temporary files handled by the Operating System (OS). The command to be used for file compression is **virt-sparsify**.

```
ERROR:      You cannot shut down Cisco Prime Access Registrar while the
            CLI is being used.   Current list of running
            CLI with process id is:
```

```
2903 /opt/CSCOar/bin/aregcmd -s
```

This process can take some time (around 10-15 minutes). Once finished, the resulting file is the one that needs to be transferred to an external entity as specified on next step.

Verification of the file integrity is required, in order to achieve this, run the next command and look for the “corrupt” attribute at the end of its output.

```
ERROR:      You cannot shut down Cisco Prime Access Registrar while the
            CLI is being used.   Current list of running
            CLI with process id is:
```

```
2903 /opt/CSCOar/bin/aregcmd -s
```

- In order to avoid a problem where the OSPD is lost, the recently created snapshot on QCOW2 format needs to be transferred to an external entity. Before you start the file transfer, you have to check if the destination have enough available disk space, use the command **df -kh** in order to verify the memory space. One advice is to transfer it to another site's OSPD temporarily with the use of SFTP **sftproot@x.x.x.x** where x.x.x.x is the IP of a remote OSPD. In order to speed up the transfer, the destination can be sent to multiple OSPDs. In the same way, you can run the command **scp *name_of_the_file*.qcow2 root@ x.x.x.x:/tmp** (where x.x.x.x is the IP of a remote OSPD) in order to transfer the file to another OSPD.

Graceful Power Off

- Power off Node

1. In order to power off the instance : **nova stop <INSTANCE_NAME>**
2. You can see the instance name with the status shutoff.

```
ERROR:      You cannot shut down Cisco Prime Access Registrar while the
            CLI is being used.   Current list of running
            CLI with process id is:
```

```
2903 /opt/CSCOar/bin/aregcmd -s
```

Replace Faulty Component from Compute Node

Power off the specified server. The steps in order to replace a faulty component on UCS C240 M4 server can be referred from:

[Replacing the Server Components](#)

Restore VMs

Recover Instance with Snapshot

Recovery process

It is possible to redeploy the previous instance with the snapshot taken in previous steps.

Step 1. [optional] If there is no previous VMsnapshot available, then connect to the OSPD node where the backup was sent and SFTP the backup back to its original OSPD node. With **sftproot@x.x.x.x** where x.x.x.x is the IP of a the original OSPD. Save the snapshot file in **/tmp** directory.

Step 2. Connect to the OSPD node where the instance can be re-deployed as shown in the image.

```
Last login: wed May 9 06:42:27 2018 from 10.169.119.213
[root@daucs01-ospd ~]#
```

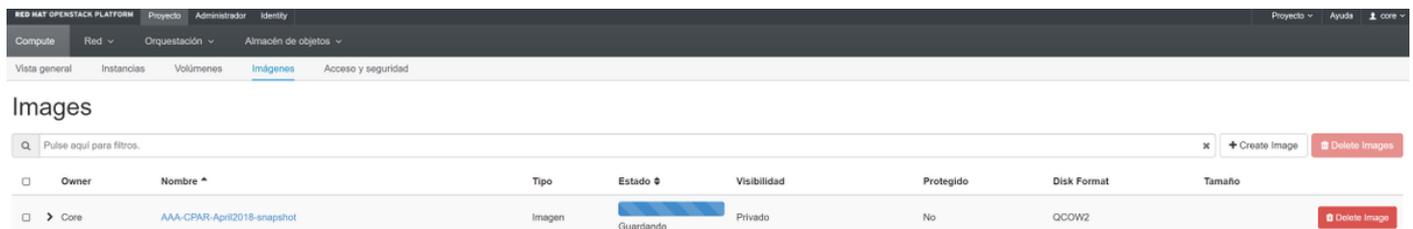
Source the environment variables with this command:

```
ERROR: You cannot shut down Cisco Prime Access Registrar while the
        CLI is being used. Current list of running
        CLI with process id is:
2903 /opt/CSCOar/bin/aregcmd -s
```

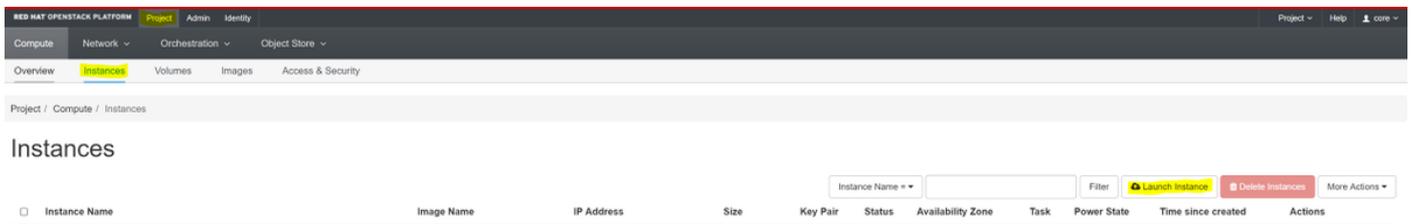
Step 3. In order to use the snapshot as an image it is necessary to upload it to the horizon as such. Run the next command to do so.

```
ERROR: You cannot shut down Cisco Prime Access Registrar while the
        CLI is being used. Current list of running
        CLI with process id is:
2903 /opt/CSCOar/bin/aregcmd -s
```

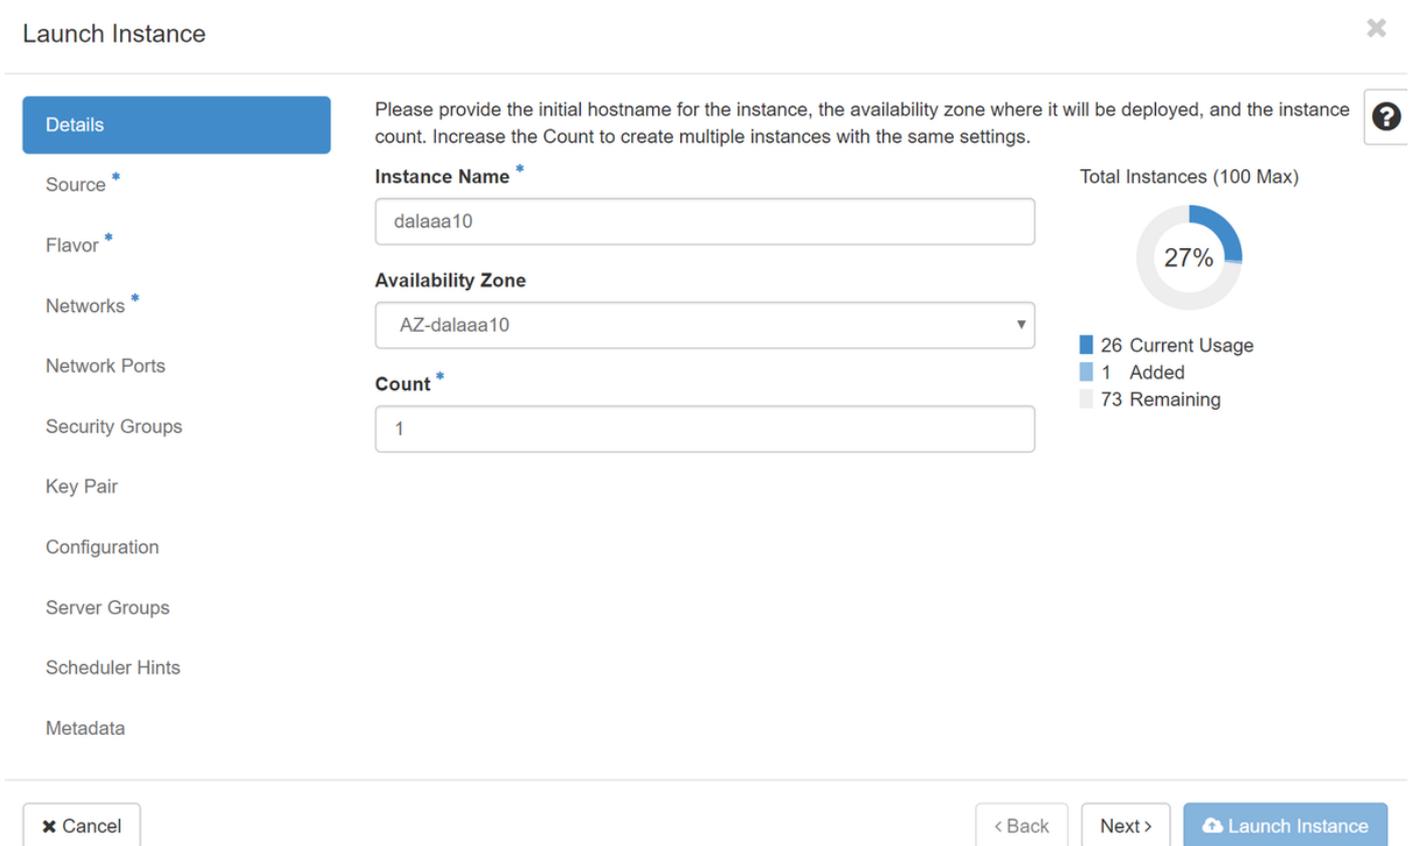
The process can be seen in horizon and as shown in this image.



Step 4. In Horizon, navigate to **Project > Instances** and click on **Launch Instance** as shown in this image.



Step 5. Enter the **Instance Name** and choose the **Availability Zone** as shown in this image.



Step 6. In the Source tab, choose the image in order to create the instance. In the **Select Boot Source** menu select **image**, a list of images are shown, choose the one that was previously uploaded by clicking on its + sign and as shown in this image.

Instance source is the template used to create an instance. You can use a snapshot of an existing instance, an image, or a volume (if enabled). You can also choose to use persistent storage by creating a new volume.

Source

Select Boot Source: Create New Volume:

Allocated

| Name | Updated | Size | Type | Visibility | |
|-------------------------------|-----------------|---------|-------|------------|---|
| > AAA-CPAR-April2018-snapshot | 5/10/18 9:56 AM | 5.43 GB | qcow2 | Private | - |

▼ Available 8 Select one

🔍 Click here for filters. ✕

| Name | Updated | Size | Type | Visibility | |
|---------------------------------|------------------|-----------|-------|------------|---|
| > redhat72-image | 4/10/18 1:00 PM | 469.87 MB | qcow2 | Private | + |
| > tmobile-pcrf-13.1.1.qcow2 | 9/9/17 1:01 PM | 2.46 GB | qcow2 | Public | + |
| > tmobile-pcrf-13.1.1.iso | 9/9/17 8:13 AM | 2.76 GB | iso | Private | + |
| > AAA-Temporary | 9/5/17 2:11 AM | 180.00 GB | qcow2 | Private | + |
| > CPAR_AAATEMPLATE_AUGUST222017 | 8/22/17 3:33 PM | 16.37 GB | qcow2 | Private | + |
| > tmobile-pcrf-13.1.0.iso | 7/11/17 7:51 AM | 2.82 GB | iso | Public | + |
| > tmobile-pcrf-13.1.0.qcow2 | 7/11/17 7:48 AM | 2.46 GB | qcow2 | Public | + |
| > ESC-image | 6/27/17 12:45 PM | 925.06 MB | qcow2 | Private | + |

✕ Cancel < Back Next > Launch Instance

Step 7. In the **Flavor** tab, choose the AAA flavor by clicking on the + sign as shown in this image.

Flavors manage the sizing for the compute, memory and storage capacity of the instance.

Allocated

| Name | VCPUS | RAM | Total Disk | Root Disk | Ephemeral Disk | Public | |
|------------|-------|-------|------------|-----------|----------------|--------|---|
| > AAA-CPAR | 36 | 32 GB | 180 GB | 180 GB | 0 GB | No | - |

Networks *
Select one

Network Ports
Click here for filters.

| Name | VCPUS | RAM | Total Disk | Root Disk | Ephemeral Disk | Public | |
|--------------|-------|--------|------------|-----------|----------------|--------|---|
| > pcrf-oam | 10 | 24 GB | 100 GB | 100 GB | 0 GB | Yes | + |
| > pcrf-pd | 12 | 16 GB | 100 GB | 100 GB | 0 GB | Yes | + |
| > pcrf-qns | 10 | 16 GB | 100 GB | 100 GB | 0 GB | Yes | + |
| > pcrf-arb | 4 | 16 GB | 100 GB | 100 GB | 0 GB | Yes | + |
| > esc-flavor | 4 | 4 GB | 0 GB | 0 GB | 0 GB | Yes | + |
| > pcrf-sm | 10 | 104 GB | 100 GB | 100 GB | 0 GB | Yes | + |
| > pcrf-cm | 6 | 16 GB | 100 GB | 100 GB | 0 GB | Yes | + |

✕ Cancel < Back Next > Launch Instance

Step 8. Finally, navigate to the **Network** tab and choose the networks that the instance will need by clicking on the + sign. For this case, select **diameter-soutable1**, **radius-routable1** and **tb1-mgmt** as shown in this image.

Details

Source

Flavor

Networks

Network Ports

Security Groups

Key Pair

Configuration

Server Groups

Scheduler Hints

Metadata

Networks provide the communication channels for instances in the cloud. ?

▼ Allocated 3 Select networks from those listed below.

| | Network | Subnets Associated | Shared | Admin State | Status | |
|-----|----------------------|------------------------|--------|-------------|--------|---|
| ↕ 1 | > radius-routable1 | radius-routable-subnet | Yes | Up | Active | − |
| ↕ 2 | > diameter-routable1 | sub-diameter-routable1 | Yes | Up | Active | − |
| ↕ 3 | > tb1-mgmt | tb1-subnet-mgmt | Yes | Up | Active | − |

▼ Available 16 Select at least one network

| | Network | Subnets Associated | Shared | Admin State | Status | |
|---|----------------|--------------------|--------|-------------|--------|---|
| > | Internal | Internal | Yes | Up | Active | + |
| > | pcrf_dap2_ldap | pcrf_dap2_ldap | Yes | Up | Active | + |
| > | pcrf_dap2_usd | pcrf_dap2_usd | Yes | Up | Active | + |
| > | tb1-orch | tb1-subnet-orch | Yes | Up | Active | + |
| > | pcrf_dap1_usd | pcrf_dap1_usd | Yes | Up | Active | + |
| > | pcrf_dap1_sy | pcrf_dap1_sy | Yes | Up | Active | + |
| > | pcrf_dap1_gx | pcrf_dap1_gx | Yes | Up | Active | + |
| > | pcrf_dap1_nap | pcrf_dap1_nap | Yes | Up | Active | + |
| > | pcrf_dap2_sy | pcrf_dap2_sy | Yes | Up | Active | + |
| > | pcrf_dap2_rx | pcrf_dap2_rx | Yes | Up | Active | + |

✕ Cancel
< Back
Next >
Launch Instance

Finally, click on **Launch Instance** in order to create it. The progress can be monitored in Horizon:

RED HAT OPENSTACK PLATFORM Proyecto Administrador Identity Proyecto Ayuda core

Sistema

Vista general Hipervisores Agregados de host Instancias Volúmenes Sabores Imágenes Redes Routers IPs flotantes Predeterminados Definiciones de los metadatos Información del Sistema

Administrador / Sistema / Instancias

Instancias

Proyección Filtrar Eliminar Instancias

| <input type="checkbox"/> | Proyecto | Host | Nombre | Nombre de la imagen | Dirección IP | Tamaño | Estado | Tarea | Estado de energía | Tiempo desde su creación | Acciones |
|--------------------------|----------|----------------------------------|----------|-----------------------------|---|----------|-----------|-----------|-------------------|--------------------------|------------------|
| <input type="checkbox"/> | Core | pod1-stack-compute-5.localdomain | dataaa10 | AAA-CPAR-April2018-snapshot | tb1-mgmt • 172.16.181.11 radius-routable1 • 10.178.6.56 diameter-routable1 • 10.178.6.40 | AAA-CPAR | Construir | Generando | Sin estado | 1 minuto | Editar instancia |

After a few minutes, the instance is completely deployed and ready for use as shown in this image.



Create and Assign Floating IP Address

A floating IP address is a routable address, which means that it's reachable from the outside of Ultra M/Openstack architecture, and it's able to communicate with other nodes from the network.

Step 1. In the Horizon top menu, navigate to **Admin > Floating IPs**.

Step 2. Click **Allocate IP to Project**.

Step 3. In the **Allocate Floating IP** window, select the **Pool** from which the new floating IP belongs, the **Project** where it is going to be assigned, and the new **Floating IP Address** itself.

For example:

Allocate Floating IP

Pool *
10.145.0.192/26 Management

Project *
Core

Floating IP Address (optional) ⓘ
10.145.0.249

Description:
From here you can allocate a floating IP to a specific project.

Cancel Allocate Floating IP

Step 4. Click **Allocate Floating IP** button.

Step 5. In the Horizon top menu, navigate to **Project > Instances**.

Step 6. In the **Action** column, click on the arrow that points down in the **Create Snapshot** button, a menu is displayed. Select **Associate Floating IP** option.

Step 7. Select the corresponding floating IP address intended to be used in the **IP Address** field, and choose the corresponding management interface (eth0) from the new instance where this floating IP is going to be assigned in the **Port to be associated**. Refer to the next image as an example of this procedure.

Manage Floating IP Associations



IP Address *

Select the IP address you wish to associate with the selected instance or port.

Port to be associated *

Cancel

Associate

Step 8. Finally, click **Associate**.

Enable SSH

Step 1. In the Horizon top menu, navigate to **Project > Instances**.

Step 2. Click on the name of the instance/VM that was created in section **Launch a new instance**.

Step 3. Click on **Console** tab. This will display the CLI of the VM.

Step 4. Once the CLI is displayed, enter the proper login credentials as shown in the image:

Username:**root**

Password:**cisco123**

```
Red Hat Enterprise Linux Server 7.0 (Maipo)
Kernel 3.10.0-514.el7.x86_64 on an x86_64

aaa-cpar-testing-instance login: root
Password:
Last login: Thu Jun 29 12:59:59 from 5.232.63.159
[root@aaa-cpar-testing-instance ~]#
```

Step 5. In the CLI, run the command **vi /etc/ssh/sshd_config** in order to edit SSH configuration.

Step 6. Once the SSH configuration file is open, press **I** to edit the file. Then look for the section and change the first line from **PasswordAuthentication no** to **PasswordAuthentication yes** as shown in this image.

```
# To disable tunneled clear text passwords, change to no here!  
PasswordAuthentication yes_  
#PermitEmptyPasswords no  
PasswordAuthentication no
```

Step 7. Press **ESC** and run **:wq!** in order to save **sshd_config** file changes.

Step 8. Run the command **service sshd restart** as shown in the image.

```
[root@aaa-cpar-testing-instance ssh]# service sshd restart  
Redirecting to /bin/systemctl restart sshd.service  
[root@aaa-cpar-testing-instance ssh]#
```

Step 9. In order to test SSH configuration changes have been correctly applied, open any SSH client and try to establish a remote secure connection **using the floating IP** assigned to the instance (i.e. 10.145.0.249) and the user **root** as shown in the image.

```
[2017-07-13 12:12.09] ~  
[dieaguil.DIEAGUIL-CWRQ7] > ssh root@10.145.0.249  
Warning: Permanently added '10.145.0.249' (RSA) to the list of known hosts  
.  
root@10.145.0.249's password:  
X11 forwarding request failed on channel 0  
Last login: Thu Jul 13 12:58:18 2017  
[root@aaa-cpar-testing-instance ~]#  
[root@aaa-cpar-testing-instance ~]#
```

Establish SSH Session

Step 1. Open a SSH session with the IP address of the corresponding VM/server where the application is installed as shown in the image.

```
[dieaguil.DIEAGUIL-CWRQ7] > ssh root@10.145.0.59  
X11 forwarding request failed on channel 0  
Last login: Wed Jun 14 17:12:22 2017 from 5.232.63.147  
[root@dalaaa07 ~]#
```

CPAR instance start

Follow these steps, once the activity has been completed and CPAR services can be re-established in the Site that was shut down.

Step 1. Login back to Horizon, navigate to **Project > Instance > Start Instance**

Step 2. Verify that the status of the instance is **Active** and the power state is **Running** as seen in this image.

Instances

| Instance Name | Image Name | IP Address | Size | Key Pair | Status | Availability Zone | Task | Power State | Time since created | Actions |
|---------------|---------------------|--|----------|----------|--------|-------------------|------|-------------|--------------------|-----------------|
| dl1aaa04 | dl1aaa01-sept092017 | diameter-routable1 • 10.160.132.231 radius-routable1 • 10.160.132.247 tb1-mgmt • 172.16.181.16 Floating IPs: • 10.250.122.114 | AAA-CPAR | - | Active | AZ-dl1aaa04 | None | Running | 3 months | Create Snapshot |

9. Post-activity Health Check

Step 1. Run the command **/opt/CSCOar/bin/arstatus** at OS level:

```
ERROR: You cannot shut down Cisco Prime Access Registrar while the
        CLI is being used. Current list of running
        CLI with process id is:
2903 /opt/CSCOar/bin/aregcmd -s
```

Step 2. Run the command **/opt/CSCOar/bin/aregcmd** at OS level and enter the admin credentials. Verify that CPAR Health is 10 out of 10 and the exit CPAR CLI.

```
ERROR: You cannot shut down Cisco Prime Access Registrar while the
        CLI is being used. Current list of running
        CLI with process id is:
2903 /opt/CSCOar/bin/aregcmd -s
```

Step 3. Run the command **netstat | grep diameter** and verify that all DRA connections are established.

The output mentioned here is for an environment where Diameter links are expected. If fewer links are displayed, this represents a disconnection from the DRA that needs to be analyzed.

```
ERROR: You cannot shut down Cisco Prime Access Registrar while the
        CLI is being used. Current list of running
        CLI with process id is:
2903 /opt/CSCOar/bin/aregcmd -s
```

Step 4. Check that the TPS log shows requests being processed by CPAR. The values highlighted represent the TPS and those are the ones you need to pay attention to.

The value of TPS must not exceed 1500.

```
[root@wscaaa04 ~]# tail -f /opt/CSCOar/logs/tps-11-21-2017.csv
11-21-2017,23:57:35,263,0
11-21-2017,23:57:50,237,0
11-21-2017,23:58:05,237,0
11-21-2017,23:58:20,257,0
11-21-2017,23:58:35,254,0
```

```
11-21-2017,23:58:50,248,0
11-21-2017,23:59:05,272,0
11-21-2017,23:59:20,243,0
11-21-2017,23:59:35,244,0
11-21-2017,23:59:50,233,0
```

Step 5. Look for any “error” or “alarm” messages in name_radius_1_log

```
[root@wscaaa04 ~]# tail -f /opt/CSC0ar/logs/tps-11-21-2017.csv
11-21-2017,23:57:35,263,0
11-21-2017,23:57:50,237,0
11-21-2017,23:58:05,237,0
11-21-2017,23:58:20,257,0
11-21-2017,23:58:35,254,0
11-21-2017,23:58:50,248,0
11-21-2017,23:59:05,272,0
11-21-2017,23:59:20,243,0
11-21-2017,23:59:35,244,0
11-21-2017,23:59:50,233,0
```

Step 6. Verify the amount of memory that the CPAR process uses by running command:

```
[root@wscaaa04 ~]# tail -f /opt/CSC0ar/logs/tps-11-21-2017.csv
11-21-2017,23:57:35,263,0
11-21-2017,23:57:50,237,0
11-21-2017,23:58:05,237,0
11-21-2017,23:58:20,257,0
11-21-2017,23:58:35,254,0
11-21-2017,23:58:50,248,0
11-21-2017,23:59:05,272,0
11-21-2017,23:59:20,243,0
11-21-2017,23:59:35,244,0
11-21-2017,23:59:50,233,0
```

```
[root@sfraaa02 ~]# top | grep radius
27008 root      20    0 20.228g 2.413g 11408 S 128.3  7.7  1165:41 radius
```

This highlighted value must be lower than 7Gb, which is the maximum allowed at application level.

Component RMA - OSD Compute Node

Identify VMs Hosted in OSD-Compute Node

Identify the VMs that are hosted on the OSD-Compute server.

```
[root@sfraaa02 ~]# top | grep radius
27008 root      20    0 20.228g 2.413g 11408 S 128.3  7.7  1165:41 radius
```

Note: In the output shown here, the first column corresponds to the UUID, the second column is the VM name and the third column is the hostname where the VM is present. The parameters from this output will be used in subsequent sections.

Backup: SNAPSHOT PROCESS

1. CPAR Application Shutdown

Step 1. Open any SSH client connected to the TMO Production network and connect to the CPAR instance.

It is important not to shut down all 4 AAA instances within one site at the same time, do it in a one by one fashion.

Step 2. In order to shut down CPAR application, run the command:

```
[root@sfraaa02 ~]# top | grep radius
27008 root      20   0 20.228g 2.413g 11408 S 128.3  7.7  1165:41 radius
```

A Message "Cisco Prime Access Registrar Server Agent shutdown complete." must show up.

Note: If a user left a CLI session open, the **arserver stop** command won't work and this message is displayed:

```
ERROR:      You cannot shut down Cisco Prime Access Registrar while the
            CLI is being used.  Current list of running
            CLI with process id is:
2903 /opt/CSCOar/bin/aregcmd -s
```

In this example, the highlighted process id 2903 needs to be terminated before CPAR can be stopped. If this is the case, terminate the process by running the command:

```
ERROR:      You cannot shut down Cisco Prime Access Registrar while the
            CLI is being used.  Current list of running
            CLI with process id is:
2903 /opt/CSCOar/bin/aregcmd -s
```

Then repeat the step 1.

Step 3. Verify that the CPAR application was indeed shutdown by running the command:

```
ERROR:      You cannot shut down Cisco Prime Access Registrar while the
            CLI is being used.  Current list of running
            CLI with process id is:
2903 /opt/CSCOar/bin/aregcmd -s
```

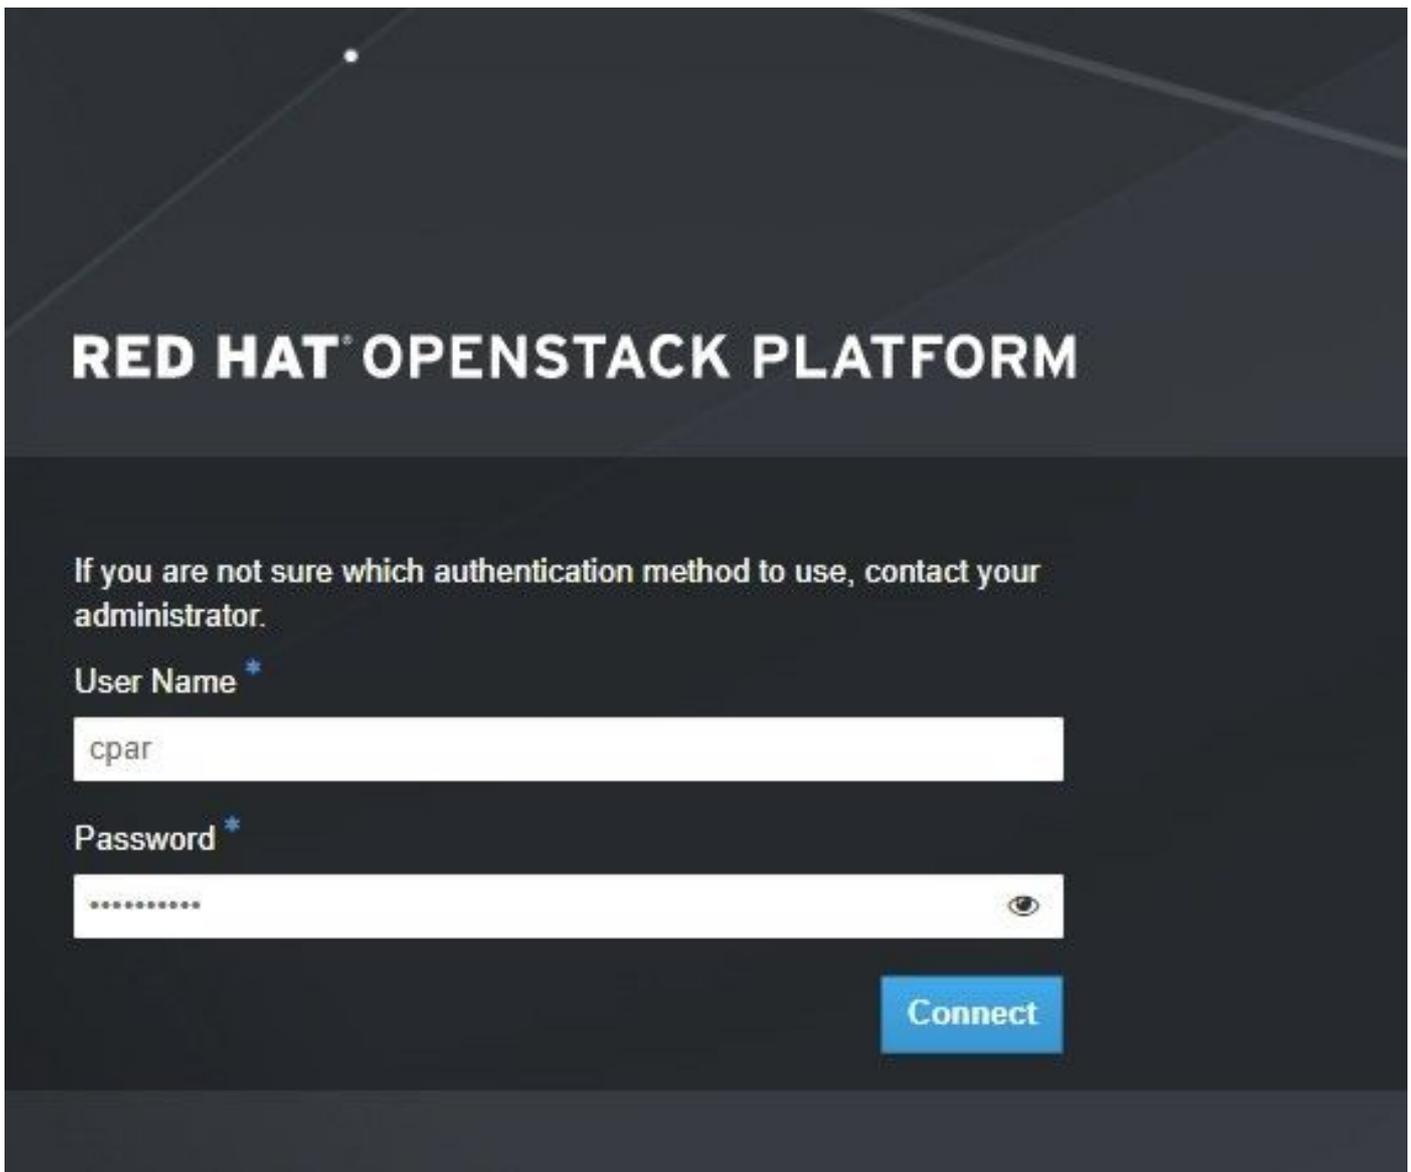
These messages must appear:

```
ERROR:      You cannot shut down Cisco Prime Access Registrar while the
            CLI is being used.  Current list of running
            CLI with process id is:
2903 /opt/CSCOar/bin/aregcmd -s
```

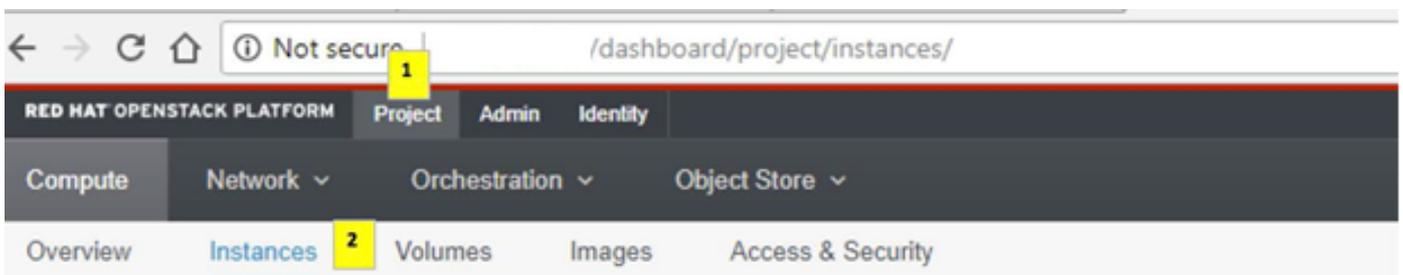
2. VM Snapshot Task

Step 1. Enter the Horizon GUI website that corresponds to the Site (City) currently being worked on.

When you access Horizon, this screen can be observed.



Step 2. Navigate to **Project > Instances** as shown in this image.

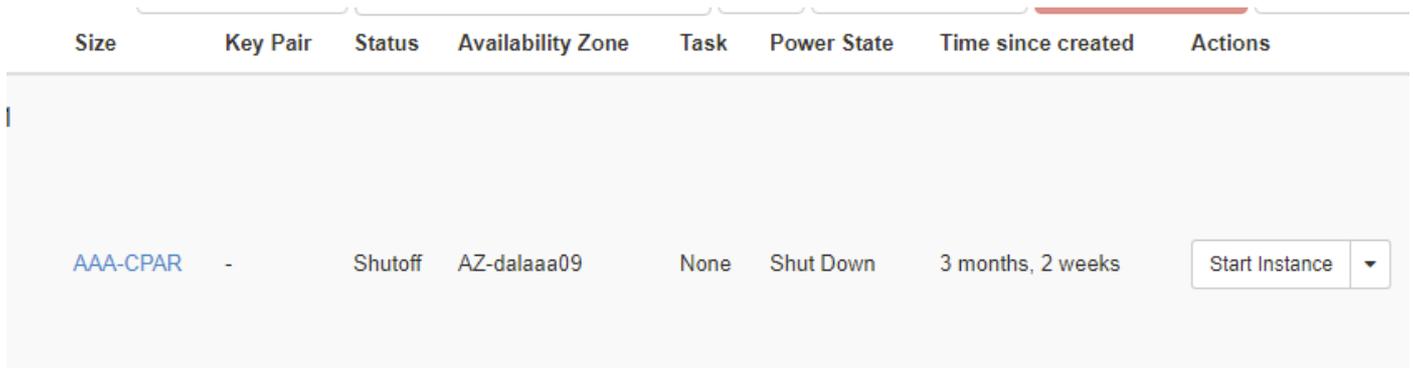


If the user used was CPAR, then only the 4 AAA instances can appear in this menu.

Step 3. Shut down only one instance at a time, repeat the whole process in this document. In order to shutdown the VM, navigate to **Actions > Shut Off Instance** as shown in the image and confirm your selection.



Step 4. Validate that the instance was indeed shut down by checking the Status = **Shutoff** and Power State = **Shut Down** as shown in the image.



| Size | Key Pair | Status | Availability Zone | Task | Power State | Time since created | Actions |
|----------|----------|---------|-------------------|------|-------------|--------------------|----------------|
| AAA-CPAR | - | Shutoff | AZ-dalaaa09 | None | Shut Down | 3 months, 2 weeks | Start Instance |

This step ends the CPAR shutdown process.

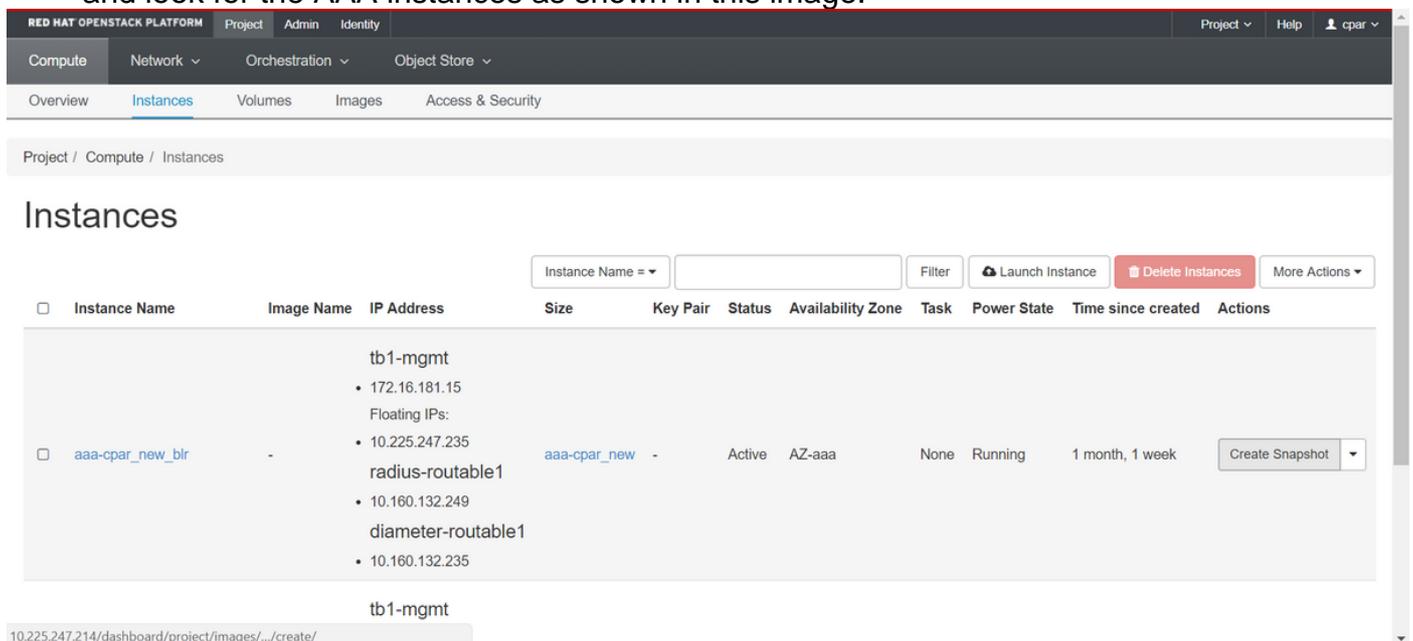
VM Snapshot

Once the CPAR VMs are down, the snapshots can be taken in parallel, as they belong to independent computes.

The four QCOW2 files are created in parallel.

Take a snapshot of each AAA instance. (25 minutes -1 hour) (25 minutes for instances that used a qcow image as a source and 1 hour for instances that user a raw image as a source)

1. Login to POD's Openstack's Horizon **GUI**
2. Once logged in, navigate to the **Project > Compute > Instances** section on the top menu and look for the AAA instances as shown in this image.



RED HAT OPENSTACK PLATFORM Project Admin Identity Project Help cpar

Compute Network Orchestration Object Store

Overview Instances Volumes Images Access & Security

Project / Compute / Instances

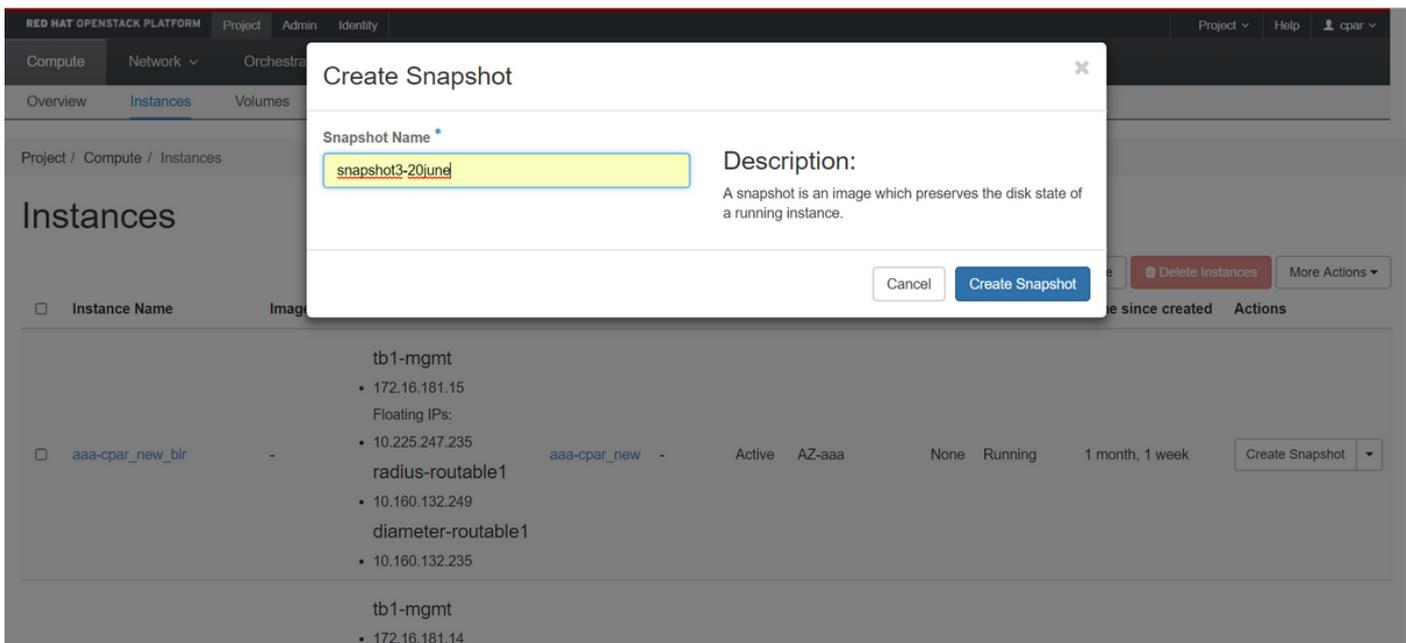
Instances

Instance Name = [] Filter [] Launch Instance [] Delete Instances [] More Actions []

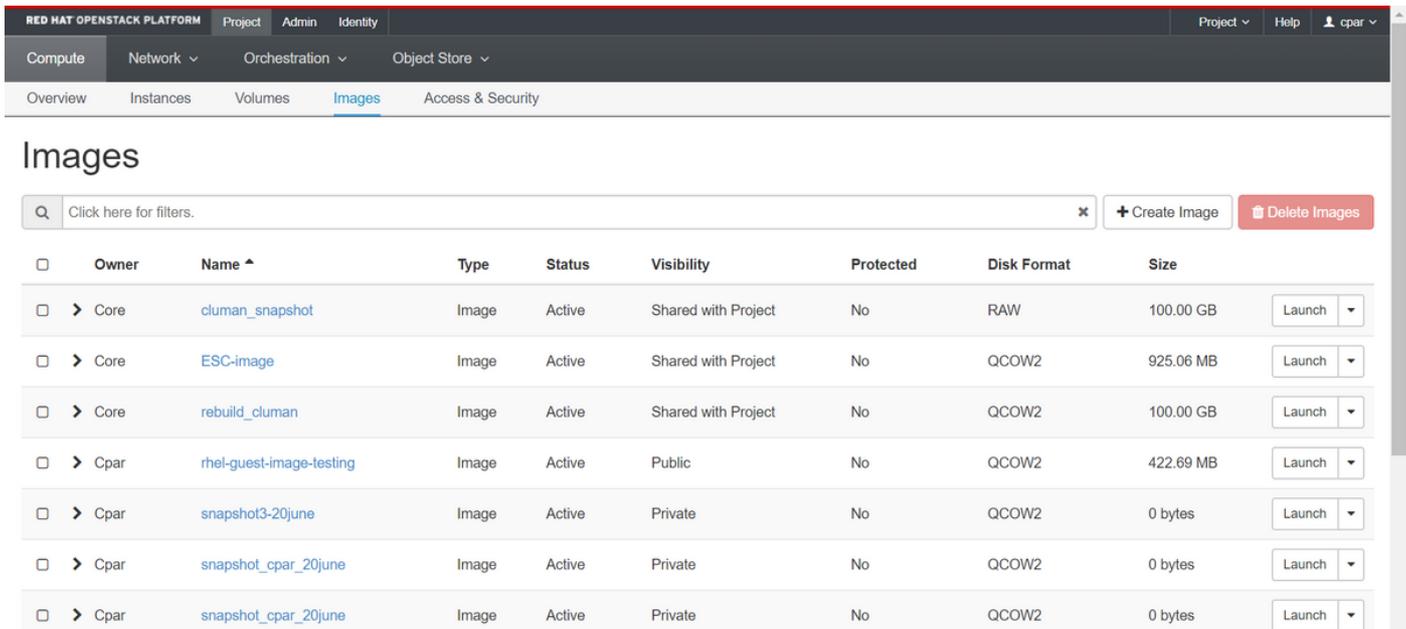
| Instance Name | Image Name | IP Address | Size | Key Pair | Status | Availability Zone | Task | Power State | Time since created | Actions |
|---|------------|--|--------------|----------|--------|-------------------|------|-------------|--------------------|-----------------|
| <input type="checkbox"/> aaa-cpar_new_blr | - | tb1-mgmt • 172.16.181.15 Floating IPs: • 10.225.247.235 radius-routable1 • 10.160.132.249 diameter-routable1 • 10.160.132.235 tb1-mgmt | aaa-cpar_new | - | Active | AZ-aaa | None | Running | 1 month, 1 week | Create Snapshot |

10.225.247.214/dashboard/project/images/.../create/

3. Click **Create Snapshot** in order to proceed with snapshot creation (this needs to be executed on the corresponding AAA instance) as shown in the image.



4. Once the snapshot is executed, navigate to the **Images** menu and verify that all finish and report no problems as seen in this image.



5. The next step is to download the snapshot on a QCOW2 format and transfer it to a remote entity, in case the OSPD is lost during this process. In order to achieve this, identify the snapshot by running the command **glance image-list** at OSPD level.

```
ERROR: You cannot shut down Cisco Prime Access Registrar while the
        CLI is being used. Current list of running
        CLI with process id is:
2903 /opt/CSC0ar/bin/aregcmd -s
```

6. Once you identify the snapshot to be downloaded (the one marked in green), you can download it on a QCOW2 format with the command **glance image-download** as depicted here.

```
ERROR: You cannot shut down Cisco Prime Access Registrar while the
        CLI is being used. Current list of running
        CLI with process id is:
```

```
2903 /opt/CSC0ar/bin/aregcmd -s
```

- The **&** sends the process to background. It can take some time to complete this action, once it is done, the image can be located at **/tmp** directory.
- On sending the process to background , if connectivity is lost, then the process is also stopped.
- Run the command **disown -h** so that in case SSH connection is lost, the process still runs and finishes on the OSPD.

7. Once the download process finishes, a compression process needs to be executed as that snapshot can be filled with ZEROES because of processes, tasks and temporary files handled by the OS. The command to be used for file compression is **virt-sparsify**.

```
ERROR:      You cannot shut down Cisco Prime Access Registrar while the
            CLI is being used.   Current list of running
            CLI with process id is:
```

```
2903 /opt/CSC0ar/bin/aregcmd -s
```

This process can take some time (around 10-15 minutes). Once finished, the resulting file is the one that needs to be transferred to an external entity as specified on next step.

Verification of the file integrity is required, in order to achieve this, run the next command and look for the “corrupt” attribute at the end of its output.

```
ERROR:      You cannot shut down Cisco Prime Access Registrar while the
            CLI is being used.   Current list of running
            CLI with process id is:
```

```
2903 /opt/CSC0ar/bin/aregcmd -s
```

- In order to avoid a problem where the OSPD is lost, the recently created snapshot on QCOW2 format needs to be transferred to an external entity. Before you start the file transfer, you have to check if the destination have enough available disk space, run the command **df -kh** in order to verify the memory space. One advice is to transfer it to another site's OSPD temporarily by using SFTP **sftproot@x.x.x.x** where x.x.x.x is the IP of a remote OSPD. In order to speed up the transfer, the destination can be sent to multiple OSPDs. In the same way, you can run the command **scp *name_of_the_file*.qcow2 root@ x.x.x.x:/tmp** (where x.x.x.x is the IP of a remote OSPD) in order to transfer the file to another OSPD.

Put CEPH in Maintenance Mode

Note: If the faulty component is to be replaced on OSD-Compute node, put the **Ceph** into Maintenance on the server before you proceed with the component replacement.

- Verify **ceph osd tree status** are up in the server.

```
ERROR:      You cannot shut down Cisco Prime Access Registrar while the
            CLI is being used.   Current list of running
            CLI with process id is:
```

```
2903 /opt/CSC0ar/bin/aregcmd -s
```

- Log into the OSD Compute node and put CEPH in the maintenance mode.

```
[root@pod2-stack-osd-compute-0 ~]# sudo ceph osd set norebalance
[root@pod2-stack-osd-compute-0 ~]# sudo ceph osd set noout
```

```
[root@pod2-stack-osd-compute-0 ~]# sudo ceph status
```

```
cluster eb2bb192-b1c9-11e6-9205-525400330666
health HEALTH_WARN
noout,norebalance,sortbitwise,require_jewel_osds flag(s) set
monmap e1: 3 mons at {pod2-stack-controller-0=11.118.0.10:6789/0,pod2-stack-controller-1=11.118.0.11:6789/0,pod2-stack-controller-2=11.118.0.12:6789/0}
election epoch 10, quorum 0,1,2 pod2-stack-controller-0,pod2-stack-controller-1,pod2-stack-controller-2
osdmap e79: 12 osds: 12 up, 12 in
flags noout,norebalance,sortbitwise,require_jewel_osds
pgmap v22844323: 704 pgs, 6 pools, 804 GB data, 423 kobjects
2404 GB used, 10989 GB / 13393 GB avail
704 active+clean
client io 3858 kB/s wr, 0 op/s rd, 546 op/s wr
```

Note: When CEPH is removed, VNF HD RAID goes into Degraded state but hd-disk must still be accessible.

Graceful Power Off

- Power off Node

1. In order to power off the instance : **nova stop <INSTANCE_NAME>**
2. You can see the instance name with the status shutoff.

```
[root@pod2-stack-osd-compute-0 ~]# sudo ceph osd set norebalance
[root@pod2-stack-osd-compute-0 ~]# sudo ceph osd set noout
```

```
[root@pod2-stack-osd-compute-0 ~]# sudo ceph status
```

```
cluster eb2bb192-b1c9-11e6-9205-525400330666
health HEALTH_WARN
noout,norebalance,sortbitwise,require_jewel_osds flag(s) set
monmap e1: 3 mons at {pod2-stack-controller-0=11.118.0.10:6789/0,pod2-stack-controller-1=11.118.0.11:6789/0,pod2-stack-controller-2=11.118.0.12:6789/0}
election epoch 10, quorum 0,1,2 pod2-stack-controller-0,pod2-stack-controller-1,pod2-stack-controller-2
osdmap e79: 12 osds: 12 up, 12 in
flags noout,norebalance,sortbitwise,require_jewel_osds
pgmap v22844323: 704 pgs, 6 pools, 804 GB data, 423 kobjects
2404 GB used, 10989 GB / 13393 GB avail
704 active+clean
client io 3858 kB/s wr, 0 op/s rd, 546 op/s wr
```

Replace Faulty Component from OSD-Compute Node

Power off the specified server. The steps in order to replace a faulty component on UCS C240 M4 server can be referred from:

[Replacing the Server Components](#)

Move CEPH Out of Maintenance Mode

- Log into the OSD Compute node and move CEPH out of the maintenance mode.

```
[root@pod2-stack-osd-compute-0 ~]# sudo ceph osd unset norebalance
[root@pod2-stack-osd-compute-0 ~]# sudo ceph osd unset noout
```

```
[root@pod2-stack-osd-compute-0 ~]# sudo ceph status
```

```
cluster eb2bb192-b1c9-11e6-9205-525400330666
health HEALTH_OK
monmap e1: 3 mons at {pod2-stack-controller-0=11.118.0.10:6789/0,pod2-stack-controller-1=11.118.0.11:6789/0,pod2-stack-controller-2=11.118.0.12:6789/0}
election epoch 10, quorum 0,1,2 pod2-stack-controller-0,pod2-stack-controller-1,pod2-stack-controller-2
osdmap e81: 12 osds: 12 up, 12 in
flags sortbitwise,require_jewel_osds
pgmap v22844355: 704 pgs, 6 pools, 804 GB data, 423 kobjects
2404 GB used, 10989 GB / 13393 GB avail
704 active+clean
client io 3658 kB/s wr, 0 op/s rd, 502 op/s wr
```

Restore VMs

Recover Instance with Snapshot

Recovery Process

It is possible to redeploy the previous instance with the snapshot taken in previous steps.

Step 1. [OPTIONAL] If there is no previous VMsnapshot available then connect to the OSPD node where the backup was sent and sftp the backup back to its original OSPD node. Using [sftproot@x.x.x.x](#) where x.x.x.x is the IP of a the original OSPD. Save the snapshot file in **/tmp** directory.

Step 2. Connect to the OSPD node where the instance will be re-deploy.

```
Last login: wed May 9 06:42:27 2018 from 10.169.119.213
[root@dauacs01-ospd ~]# █
```

Source the environment variables with this command:

```
[root@pod2-stack-osd-compute-0 ~]# sudo ceph osd unset norebalance
[root@pod2-stack-osd-compute-0 ~]# sudo ceph osd unset noout
```

```
[root@pod2-stack-osd-compute-0 ~]# sudo ceph status
```

```
cluster eb2bb192-b1c9-11e6-9205-525400330666
health HEALTH_OK
monmap e1: 3 mons at {pod2-stack-controller-0=11.118.0.10:6789/0,pod2-stack-controller-1=11.118.0.11:6789/0,pod2-stack-controller-2=11.118.0.12:6789/0}
election epoch 10, quorum 0,1,2 pod2-stack-controller-0,pod2-stack-controller-1,pod2-stack-controller-2
osdmap e81: 12 osds: 12 up, 12 in
flags sortbitwise,require_jewel_osds
pgmap v22844355: 704 pgs, 6 pools, 804 GB data, 423 kobjects
2404 GB used, 10989 GB / 13393 GB avail
704 active+clean
```

```
client io 3658 kB/s wr, 0 op/s rd, 502 op/s wr
```

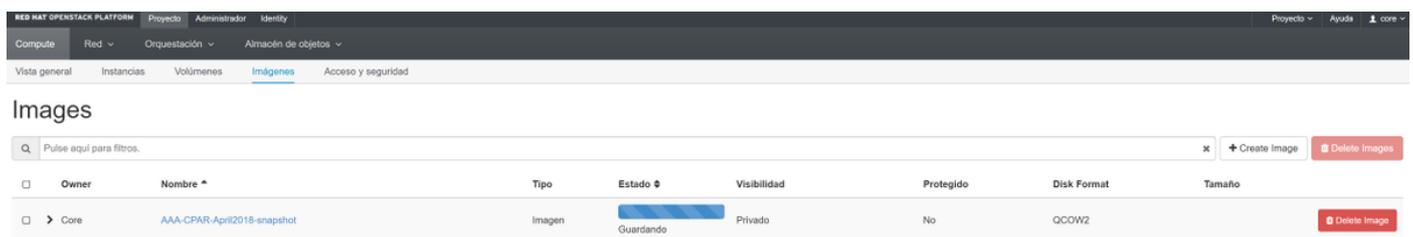
Step 3. In order to use the snapshot as an image it is necessary to upload it to horizon as such. Run the next command to do so.

```
[root@pod2-stack-osd-compute-0 ~]# sudo ceph osd unset norebalance  
[root@pod2-stack-osd-compute-0 ~]# sudo ceph osd unset noout
```

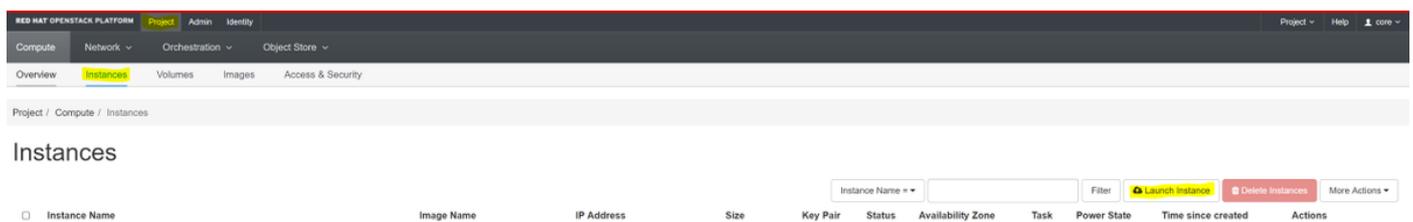
```
[root@pod2-stack-osd-compute-0 ~]# sudo ceph status
```

```
cluster eb2bb192-b1c9-11e6-9205-525400330666  
health HEALTH_OK  
monmap e1: 3 mons at {pod2-stack-controller-0=11.118.0.10:6789/0,pod2-stack-controller-1=11.118.0.11:6789/0,pod2-stack-controller-2=11.118.0.12:6789/0}  
election epoch 10, quorum 0,1,2 pod2-stack-controller-0,pod2-stack-controller-1,pod2-stack-controller-2  
osdmap e81: 12 osds: 12 up, 12 in  
flags sortbitwise,require_jewel_osds  
pgmap v22844355: 704 pgs, 6 pools, 804 GB data, 423 kobjects  
2404 GB used, 10989 GB / 13393 GB avail  
704 active+clean  
client io 3658 kB/s wr, 0 op/s rd, 502 op/s wr
```

The process can be seen in horizon.



Step 4. In Horizon, navigate to **Project > Instances** and click on **Launch Instance** as shown in this image.



Step 5. Enter **Instance Name** and choose the **Availability Zone** as shown in the image.

Details

Source *
Flavor *
Networks *
Network Ports
Security Groups
Key Pair
Configuration
Server Groups
Scheduler Hints
Metadata

Please provide the initial hostname for the instance, the availability zone where it will be deployed, and the instance count. Increase the Count to create multiple instances with the same settings.

Instance Name *
dalaaa10

Availability Zone
AZ-dalaaa10

Count *
1

Total Instances (100 Max)
27%

- 26 Current Usage
- 1 Added
- 73 Remaining

✕ Cancel < Back Next > Launch Instance

Step 6. In the **Source** tab choose the image to create the instance. In the **Select Boot Source** menu select **Image**, a list of images are shown, choose the one that was previously uploaded by clicking on its **+** sign.

Details

Source

Flavor

Networks

Network Ports

Security Groups

Key Pair

Configuration

Server Groups

Scheduler Hints

Metadata

Instance source is the template used to create an instance. You can use a snapshot of an existing instance, an image, or a volume (if enabled). You can also choose to use persistent storage by creating a new volume.



Select Boot Source

Image

Create New Volume

Yes No

Allocated

| Name | Updated | Size | Type | Visibility | |
|-------------------------------|-----------------|---------|-------|------------|---|
| > AAA-CPAR-April2018-snapshot | 5/10/18 9:56 AM | 5.43 GB | qcow2 | Private | - |

Available 8

Select one

Click here for filters.

| Name | Updated | Size | Type | Visibility | |
|---------------------------------|------------------|-----------|-------|------------|---|
| > redhat72-image | 4/10/18 1:00 PM | 469.87 MB | qcow2 | Private | + |
| > tmobile-pcrf-13.1.1.qcow2 | 9/9/17 1:01 PM | 2.46 GB | qcow2 | Public | + |
| > tmobile-pcrf-13.1.1.iso | 9/9/17 8:13 AM | 2.76 GB | iso | Private | + |
| > AAA-Temporary | 9/5/17 2:11 AM | 180.00 GB | qcow2 | Private | + |
| > CPAR_AAATEMPLATE_AUGUST222017 | 8/22/17 3:33 PM | 16.37 GB | qcow2 | Private | + |
| > tmobile-pcrf-13.1.0.iso | 7/11/17 7:51 AM | 2.82 GB | iso | Public | + |
| > tmobile-pcrf-13.1.0.qcow2 | 7/11/17 7:48 AM | 2.46 GB | qcow2 | Public | + |
| > ESC-image | 6/27/17 12:45 PM | 925.06 MB | qcow2 | Private | + |

✕ Cancel

< Back

Next >

Launch Instance

Step 7. In the **Flavor** tab, choose the AAA flavor by clicking on the + sign.

Flavors manage the sizing for the compute, memory and storage capacity of the instance.

Allocated

| Name | VCPUS | RAM | Total Disk | Root Disk | Ephemeral Disk | Public | |
|------------|-------|-------|------------|-----------|----------------|--------|---|
| > AAA-CPAR | 36 | 32 GB | 180 GB | 180 GB | 0 GB | No | - |

Networks *
Network Ports
Security Groups
Key Pair
Configuration
Server Groups
Scheduler Hints
Metadata

Available 7 Select one

Q Click here for filters. ✕

| Name | VCPUS | RAM | Total Disk | Root Disk | Ephemeral Disk | Public | |
|--------------|-------|--------|------------|-----------|----------------|--------|---|
| > pcrf-oam | 10 | 24 GB | 100 GB | 100 GB | 0 GB | Yes | + |
| > pcrf-pd | 12 | 16 GB | 100 GB | 100 GB | 0 GB | Yes | + |
| > pcrf-qns | 10 | 16 GB | 100 GB | 100 GB | 0 GB | Yes | + |
| > pcrf-arb | 4 | 16 GB | 100 GB | 100 GB | 0 GB | Yes | + |
| > esc-flavor | 4 | 4 GB | 0 GB | 0 GB | 0 GB | Yes | + |
| > pcrf-sm | 10 | 104 GB | 100 GB | 100 GB | 0 GB | Yes | + |
| > pcrf-cm | 6 | 16 GB | 100 GB | 100 GB | 0 GB | Yes | + |

✕ Cancel < Back Next > Launch Instance

Step 8. Finally, navigate to the **Networks** tab and choose the networks that the instance will need by clicking on the + sign. For this case, select **diameter-soutable1**, **radius-routable1** and **tb1-mgmt** as shown in this image.

Details

Source

Flavor

Networks

Network Ports

Security Groups

Key Pair

Configuration

Server Groups

Scheduler Hints

Metadata

Networks provide the communication channels for instances in the cloud. ?

▼ Allocated 3 Select networks from those listed below.

| | Network | Subnets Associated | Shared | Admin State | Status | |
|-----|----------------------|------------------------|--------|-------------|--------|---|
| ↕ 1 | > radius-routable1 | radius-routable-subnet | Yes | Up | Active | − |
| ↕ 2 | > diameter-routable1 | sub-diameter-routable1 | Yes | Up | Active | − |
| ↕ 3 | > tb1-mgmt | tb1-subnet-mgmt | Yes | Up | Active | − |

▼ Available 16 Select at least one network

| | Network | Subnets Associated | Shared | Admin State | Status | |
|---|----------------|--------------------|--------|-------------|--------|---|
| > | Internal | Internal | Yes | Up | Active | + |
| > | pcrf_dap2_ldap | pcrf_dap2_ldap | Yes | Up | Active | + |
| > | pcrf_dap2_usd | pcrf_dap2_usd | Yes | Up | Active | + |
| > | tb1-orch | tb1-subnet-orch | Yes | Up | Active | + |
| > | pcrf_dap1_usd | pcrf_dap1_usd | Yes | Up | Active | + |
| > | pcrf_dap1_sy | pcrf_dap1_sy | Yes | Up | Active | + |
| > | pcrf_dap1_gx | pcrf_dap1_gx | Yes | Up | Active | + |
| > | pcrf_dap1_nap | pcrf_dap1_nap | Yes | Up | Active | + |
| > | pcrf_dap2_sy | pcrf_dap2_sy | Yes | Up | Active | + |
| > | pcrf_dap2_rx | pcrf_dap2_rx | Yes | Up | Active | + |

✕ Cancel
< Back
Next >
Launch Instance

Finally, click on **Launch Instance** to create it. The progress can be monitored in Horizon:

RED HAT OPENSTACK PLATFORM Proyecto Administrador Identity Proyecto Ayuda core

Sistema

Vista general Hipervisores Agregados de host **Instancias** Volúmenes Sabores Imágenes Redes Routers IPs flotantes Predeterminados Definiciones de los metadatos Información del Sistema

Administrador / Sistema / Instancias

Instancias

Proyecto ▼

Filtrar
Eliminar Instancias

| <input type="checkbox"/> | Proyecto | Host | Nombre | Nombre de la imagen | Dirección IP | Tamaño | Estado | Tarea | Estado de energía | Tiempo desde su creación | Acciones |
|--------------------------|----------|----------------------------------|----------|-----------------------------|---|----------|-----------|-----------|-------------------|--------------------------|--------------------|
| <input type="checkbox"/> | Core | pod1-stack-compute-5.localdomain | dataaa10 | AAA-CPAR-April2018-snapshot | tb1-mgmt • 172.16.181.11 radius-routable1 • 10.178.6.56 diameter-routable1 • 10.178.6.40 | AAA-CPAR | Construir | Generando | Sin estado | 1 minuto | Editar instancia ▼ |

After a few minutes the instance will be completely deployed and ready for use.

| Core | pod1-stack-compute-5.localdomain | dalaaa10 | AAA-CPAR-April2018-snapshot | tb1-mgmt | AAA-CPAR | Activo | Ninguno | Ejecutando | 8 minutos | Editar instancia |
|------|----------------------------------|----------|-----------------------------|--|----------|--------|---------|------------|-----------|------------------|
| | | | | <ul style="list-style-type: none"> 172.16.181.16 IPs flotantes: 10.145.0.62 radius-routable1 10.178.6.56 diameter-routable1 10.178.6.40 | | | | | | |

Create and assign a floating IP address

A floating IP address is a routable address, which means that it's reachable from the outside of Ultra M/Openstack architecture, and it's able to communicate with other nodes from the network.

Step 1. In the Horizon top menu, navigate to **Admin > Floating IPs**.

Step 2. Click **Allocate IP to Project**.

Step 3. In the **Allocate Floating IP** window select the **Pool** from which the new floating IP belongs, the **Project** where it is going to be assigned, and the new **Floating IP Address** itself.

For example:

Allocate Floating IP ✕

Pool *

10.145.0.192/26 Management ▼

Project *

Core ▼

Floating IP Address (optional) ⓘ

10.145.0.249

Description:

From here you can allocate a floating IP to a specific project.

Cancel
Allocate Floating IP

Step 4. Click **Allocate Floating IP**.

Step 5. In the Horizon top menu, navigate to **Project > Instances**.

Step 6. In the **Action** column click on the arrow that points down in the **Create Snapshot** button, a menu must be displayed. Select **Associate Floating IP** option.

Step 7. Select the corresponding floating IP address intended to be used in the **IP Address** field, and choose the corresponding management interface (eth0) from the new instance where this floating IP is going to be assigned in the **Port to be associated**. Refer to the next image as an example of this procedure.

Manage Floating IP Associations



IP Address *

 ▼ +

Select the IP address you wish to associate with the selected instance or port.

Port to be associated *

 ▼

Cancel

Associate

Step 8. Finally, click **Associate**.

Enable SSH

Step 1. In the Horizon top menu, navigate to **Project > Instances**.

Step 2. Click on the name of the instance/VM that was created in section **Launch a new instance**.

Step 3. Click on **Console** tab. This will display the command line interface of the VM.

Step 4. Once the CLI is displayed, enter the proper login credentials as shown in the image:

Username:**root**

Password:**cisco123**

```
Red Hat Enterprise Linux Server 7.0 (Maipo)
Kernel 3.10.0-514.el7.x86_64 on an x86_64

aaa-cpar-testing-instance login: root
Password:
Last login: Thu Jun 29 12:59:59 from 5.232.63.159
[root@aaa-cpar-testing-instance ~]#
```

Step 5. In the CLI, run the command **vi /etc/ssh/sshd_config** in order to edit ssh configuration.

Step 6. Once the ssh configuration file is open, press **I** to edit the file. Then look for this section

and change the first line from **PasswordAuthentication no** to **PasswordAuthentication yes**.

```
# To disable tunneled clear text passwords, change to no here!  
PasswordAuthentication yes_  
#PermitEmptyPasswords no  
PasswordAuthentication no
```

Step 7. Press **ESC** and enter **:wq!** to save sshd_config file changes.

Step 8. Run the command **service sshd restart**.

```
root@aaa-cpar-testing-instance ssh]# service sshd restart  
Redirecting to /bin/systemctl restart sshd.service  
root@aaa-cpar-testing-instance ssh]#
```

Step 9. In order to test SSH configuration changes have been correctly applied, open any SSH client and try to establish a remote secure connection **using the floating IP** assigned to the instance (i.e. 10.145.0.249) and the user **root**.

```
[2017-07-13 12:12:09] ~  
[dieaguil.DIEAGUIL-CWRQ7] > ssh root@10.145.0.249  
Warning: Permanently added '10.145.0.249' (RSA) to the list of known hosts  
.  
root@10.145.0.249's password:  
X11 forwarding request failed on channel 0  
Last login: Thu Jul 13 12:58:18 2017  
[root@aaa-cpar-testing-instance ~]#  
[root@aaa-cpar-testing-instance ~]#
```

Establish SSH session

Step 1. Open a SSH session using the IP address of the corresponding VM/server where the application is installed.

```
[dieaguil.DIEAGUIL-CWRQ7] > ssh root@10.145.0.59  
X11 forwarding request failed on channel 0  
Last login: Wed Jun 14 17:12:22 2017 from 5.232.63.147  
[root@dalaaa07 ~]#
```

CPAR instance start

Follow these steps, once the activity has been completed and CPAR services can be re-established in the Site that was shut down.

Step 1. Login back to Horizon, navigate to **Project > Instance > Start Instance**.

Step 2. Verify that the status of the instance is **Active** and the power state is **Running** as shown in the image.

Instances

| Instance Name | Image Name | IP Address | Size | Key Pair | Status | Availability Zone | Task | Power State | Time since created | Actions |
|-----------------------------------|---------------------|--|----------|----------|--------|-------------------|------|-------------|--------------------|-----------------|
| <input type="checkbox"/> dilaaa04 | dilaaa01-sept092017 | diameter-routable1 • 10.160.132.231 radius-routable1 • 10.160.132.247 tb1-mgmt • 172.16.181.16 Floating IPs: • 10.250.122.114 | AAA-CPAR | - | Active | AZ-dilaaa04 | None | Running | 3 months | Create Snapshot |

9. Post-activity Health Check

Step 1. Run the command `/opt/CSCOar/bin/arstatus` at OS level

```
[root@pod2-stack-osd-compute-0 ~]# sudo ceph osd unset norebalance
[root@pod2-stack-osd-compute-0 ~]# sudo ceph osd unset noout
```

```
[root@pod2-stack-osd-compute-0 ~]# sudo ceph status
```

```
cluster eb2bb192-b1c9-11e6-9205-525400330666
health HEALTH_OK
monmap e1: 3 mons at {pod2-stack-controller-0=11.118.0.10:6789/0,pod2-stack-controller-1=11.118.0.11:6789/0,pod2-stack-controller-2=11.118.0.12:6789/0}
election epoch 10, quorum 0,1,2 pod2-stack-controller-0,pod2-stack-controller-1,pod2-stack-controller-2
osdmap e81: 12 osds: 12 up, 12 in
flags sortbitwise,require_jewel_osds
pgmap v22844355: 704 pgs, 6 pools, 804 GB data, 423 kobjects
2404 GB used, 10989 GB / 13393 GB avail
704 active+clean
client io 3658 kB/s wr, 0 op/s rd, 502 op/s wr
```

Step 2. Run the command `/opt/CSCOar/bin/aregcmd` at OS level and enter the admin credentials. Verify that CPAR Health is 10 out of 10 and the exit CPAR CLI.

```
[root@pod2-stack-osd-compute-0 ~]# sudo ceph osd unset norebalance
[root@pod2-stack-osd-compute-0 ~]# sudo ceph osd unset noout
```

```
[root@pod2-stack-osd-compute-0 ~]# sudo ceph status
```

```
cluster eb2bb192-b1c9-11e6-9205-525400330666
health HEALTH_OK
monmap e1: 3 mons at {pod2-stack-controller-0=11.118.0.10:6789/0,pod2-stack-controller-1=11.118.0.11:6789/0,pod2-stack-controller-2=11.118.0.12:6789/0}
election epoch 10, quorum 0,1,2 pod2-stack-controller-0,pod2-stack-controller-1,pod2-stack-controller-2
osdmap e81: 12 osds: 12 up, 12 in
flags sortbitwise,require_jewel_osds
pgmap v22844355: 704 pgs, 6 pools, 804 GB data, 423 kobjects
2404 GB used, 10989 GB / 13393 GB avail
704 active+clean
client io 3658 kB/s wr, 0 op/s rd, 502 op/s wr
```

Step 3. Run the command `netstat | grep diameter` and verify that all DRA connections are established.

The output mentioned here is for an environment where Diameter links are expected. If fewer links are displayed, this represents a disconnection from the DRA that needs to be analyzed.

```
[root@pod2-stack-osd-compute-0 ~]# sudo ceph osd unset norebalance
[root@pod2-stack-osd-compute-0 ~]# sudo ceph osd unset noout
```

```
[root@pod2-stack-osd-compute-0 ~]# sudo ceph status
```

```
cluster eb2bb192-b1c9-11e6-9205-525400330666
health HEALTH_OK
monmap e1: 3 mons at {pod2-stack-controller-0=11.118.0.10:6789/0,pod2-stack-controller-1=11.118.0.11:6789/0,pod2-stack-controller-2=11.118.0.12:6789/0}
election epoch 10, quorum 0,1,2 pod2-stack-controller-0,pod2-stack-controller-1,pod2-stack-controller-2
osdmap e81: 12 osds: 12 up, 12 in
flags sortbitwise,require_jewel_osds
pgmap v22844355: 704 pgs, 6 pools, 804 GB data, 423 kobjects
2404 GB used, 10989 GB / 13393 GB avail
704 active+clean
client io 3658 kB/s wr, 0 op/s rd, 502 op/s wr
```

Step 4. Check that the TPS log shows requests being processed by CPAR. The values highlighted represent the TPS and those are the ones you need to pay attention to.

The value of TPS must not exceed 1500.

```
[root@wscaaa04 ~]# tail -f /opt/CSC0ar/logs/tps-11-21-2017.csv
11-21-2017,23:57:35,263,0
11-21-2017,23:57:50,237,0
11-21-2017,23:58:05,237,0
11-21-2017,23:58:20,257,0
11-21-2017,23:58:35,254,0
11-21-2017,23:58:50,248,0
11-21-2017,23:59:05,272,0
11-21-2017,23:59:20,243,0
11-21-2017,23:59:35,244,0
11-21-2017,23:59:50,233,0
```

Step 5. Look for any “error” or “alarm” messages in name_radius_1_log

```
[root@wscaaa04 ~]# tail -f /opt/CSC0ar/logs/tps-11-21-2017.csv
11-21-2017,23:57:35,263,0
11-21-2017,23:57:50,237,0
11-21-2017,23:58:05,237,0
11-21-2017,23:58:20,257,0
11-21-2017,23:58:35,254,0
11-21-2017,23:58:50,248,0
11-21-2017,23:59:05,272,0
11-21-2017,23:59:20,243,0
11-21-2017,23:59:35,244,0
11-21-2017,23:59:50,233,0
```

Step 6. Verify the amount of memory that the CPAR process uses by running the command:

```
[root@wscaaa04 ~]# tail -f /opt/CSC0ar/logs/tps-11-21-2017.csv
11-21-2017,23:57:35,263,0
11-21-2017,23:57:50,237,0
11-21-2017,23:58:05,237,0
11-21-2017,23:58:20,257,0
```

```
11-21-2017,23:58:35,254,0
11-21-2017,23:58:50,248,0
11-21-2017,23:59:05,272,0
11-21-2017,23:59:20,243,0
11-21-2017,23:59:35,244,0
11-21-2017,23:59:50,233,0
```

```
[root@sfraaa02 ~]# top | grep radius
27008 root      20   0 20.228g 2.413g 11408 S 128.3  7.7  1165:41 radius
```

This highlighted value must be lower than 7Gb, which is the maximum allowed at application level.

Component RMA - Controller Node

Pre-Check

- From OSPD, login to the controller and verify pcs is in a good state – all three controllers Online and Galera show all three controllers as Master.

Note: A healthy cluster requires 2 active controllers so verify that the two controllers that remain are Online and active.

```
[heat-admin@pod2-stack-controller-0 ~]$ sudo pcs status
Cluster name: tripleo_cluster
Stack: corosync
Current DC: pod2-stack-controller-2 (version 1.1.15-11.el7_3.4-e174ec8) - partition with quorum
Last updated: Fri Jul 6 09:03:37 2018Last change: Fri Jul 6 09:03:35 2018 by root via
crm_attribute on pod2-stack-controller-0
```

3 nodes and 19 resources configured

Online: [pod2-stack-controller-0 pod2-stack-controller-1 pod2-stack-controller-2]

Full list of resources:

```
ip-11.120.0.49(ocf::heartbeat:IPaddr2):Started pod2-stack-controller-1
Clone Set: haproxy-clone [haproxy]
Started: [ pod2-stack-controller-0 pod2-stack-controller-1 pod2-stack-controller-2 ]
Master/Slave Set: galera-master [galera]
Masters: [ pod2-stack-controller-1 pod2-stack-controller-2 ]
Slaves: [ pod2-stack-controller-0 ]
ip-192.200.0.110(ocf::heartbeat:IPaddr2):Started pod2-stack-controller-1
ip-11.120.0.44(ocf::heartbeat:IPaddr2):Started pod2-stack-controller-2
ip-11.118.0.49(ocf::heartbeat:IPaddr2):Started pod2-stack-controller-2
Clone Set: rabbitmq-clone [rabbitmq]
Started: [ pod2-stack-controller-1 pod2-stack-controller-2 ]
Stopped: [ pod2-stack-controller-0 ]
ip-10.225.247.214(ocf::heartbeat:IPaddr2):Started pod2-stack-controller-1
Master/Slave Set: redis-master [redis]
Masters: [ pod2-stack-controller-2 ]
Slaves: [ pod2-stack-controller-0 pod2-stack-controller-1 ]
ip-11.119.0.49(ocf::heartbeat:IPaddr2):Started pod2-stack-controller-2
openstack-cinder-volume(systemd:openstack-cinder-volume):Started pod2-stack-controller-1
```

Daemon Status:

```
corosync: active/enabled
pacemaker: active/enabled
pcsd: active/enabled
```

Move Controller Cluster to Maintenance Mode

- Run the **pcs cluster** on the controller that is updated in standby:

```
[heat-admin@pod2-stack-controller-0 ~]$ sudo pcs cluster standby
```

- Check the **pcs status** again and ensure that the pcs cluster stopped on this node:

```
[heat-admin@pod2-stack-controller-0 ~]$ sudo pcs status
Cluster name: tripleo_cluster
Stack: corosync
Current DC: pod2-stack-controller-2 (version 1.1.15-11.el7_3.4-e174ec8) - partition with quorum
Last updated: Fri Jul 6 09:03:10 2018Last change: Fri Jul 6 09:03:06 2018 by root via
crm_attribute on pod2-stack-controller-0
```

3 nodes and 19 resources configured

```
Node pod2-stack-controller-0: standby
Online: [ pod2-stack-controller-1 pod2-stack-controller-2 ]
```

Full list of resources:

```
ip-11.120.0.49(ocf::heartbeat:IPAddr2):Started pod2-stack-controller-1
Clone Set: haproxy-clone [haproxy]
Started: [ pod2-stack-controller-1 pod2-stack-controller-2 ]
Stopped: [ pod2-stack-controller-0 ]
Master/Slave Set: galera-master [galera]
Masters: [ pod2-stack-controller-0 pod2-stack-controller-1 pod2-stack-controller-2 ]
ip-192.200.0.110(ocf::heartbeat:IPAddr2):Started pod2-stack-controller-1
ip-11.120.0.44(ocf::heartbeat:IPAddr2):Started pod2-stack-controller-2
ip-11.118.0.49(ocf::heartbeat:IPAddr2):Started pod2-stack-controller-2
Clone Set: rabbitmq-clone [rabbitmq]
Started: [ pod2-stack-controller-0 pod2-stack-controller-1 pod2-stack-controller-2 ]
ip-10.225.247.214(ocf::heartbeat:IPAddr2):Started pod2-stack-controller-1
Master/Slave Set: redis-master [redis]
Masters: [ pod2-stack-controller-2 ]
Slaves: [ pod2-stack-controller-1 ]
Stopped: [ pod2-stack-controller-0 ]
ip-11.119.0.49(ocf::heartbeat:IPAddr2):Started pod2-stack-controller-2
openstack-cinder-volume(systemd:openstack-cinder-volume):Started pod2-stack-controller-1
```

```
Daemon Status:
corosync: active/enabled
pacemaker: active/enabled
pcsd: active/enabled
```

Also, pcs status on the other 2 controllers must show the node as standby.

Replace Faulty Component from Controller Node

Power off the specified server. The steps in order to replace a faulty component on UCS C240 M4 server can be referred from:

[Replacing the Server Components](#)

Power On Server

- Power on the server and verify server comes up:

```
[heat-admin@pod2-stack-controller-0 ~]$ sudo pcs status
Cluster name: tripleo_cluster
Stack: corosync
Current DC: pod2-stack-controller-2 (version 1.1.15-11.e17_3.4-e174ec8) - partition with quorum
Last updated: Fri Jul 6 09:03:10 2018Last change: Fri Jul 6 09:03:06 2018 by root via
crm_attribute on pod2-stack-controller-0
```

3 nodes and 19 resources configured

```
Node pod2-stack-controller-0: standby
Online: [ pod2-stack-controller-1 pod2-stack-controller-2 ]
```

Full list of resources:

```
ip-11.120.0.49(ocf::heartbeat:IPaddr2):Started pod2-stack-controller-1
Clone Set: haproxy-clone [haproxy]
Started: [ pod2-stack-controller-1 pod2-stack-controller-2 ]
Stopped: [ pod2-stack-controller-0 ]
Master/Slave Set: galera-master [galera]
Masters: [ pod2-stack-controller-0 pod2-stack-controller-1 pod2-stack-controller-2 ]
ip-192.200.0.110(ocf::heartbeat:IPaddr2):Started pod2-stack-controller-1
ip-11.120.0.44(ocf::heartbeat:IPaddr2):Started pod2-stack-controller-2
ip-11.118.0.49(ocf::heartbeat:IPaddr2):Started pod2-stack-controller-2
Clone Set: rabbitmq-clone [rabbitmq]
Started: [ pod2-stack-controller-0 pod2-stack-controller-1 pod2-stack-controller-2 ]
ip-10.225.247.214(ocf::heartbeat:IPaddr2):Started pod2-stack-controller-1
Master/Slave Set: redis-master [redis]
Masters: [ pod2-stack-controller-2 ]
Slaves: [ pod2-stack-controller-1 ]
Stopped: [ pod2-stack-controller-0 ]
ip-11.119.0.49(ocf::heartbeat:IPaddr2):Started pod2-stack-controller-2
openstack-cinder-volume(systemd:openstack-cinder-volume):Started pod2-stack-controller-1
```

Daemon Status:

```
corosync: active/enabled
pacemaker: active/enabled
pcsd: active/enabled
```

- Login to the impacted controller, remove standby mode with the use of **unstandby**. Verify that the controller comes online with cluster and Galera shows all three controllers as Master. This might take a few minutes:

```
[heat-admin@pod2-stack-controller-0 ~]$ sudo pcs cluster unstandby
```

```
[heat-admin@pod2-stack-controller-0 ~]$ sudo pcs status
Cluster name: tripleo_cluster
Stack: corosync
Current DC: pod2-stack-controller-2 (version 1.1.15-11.e17_3.4-e174ec8) - partition with quorum
Last updated: Fri Jul 6 09:03:37 2018Last change: Fri Jul 6 09:03:35 2018 by root via
crm_attribute on pod2-stack-controller-0
```

3 nodes and 19 resources configured

```
Online: [ pod2-stack-controller-0 pod2-stack-controller-1 pod2-stack-controller-2 ]
```

Full list of resources:

```
ip-11.120.0.49(ocf::heartbeat:IPaddr2):Started pod2-stack-controller-1
Clone Set: haproxy-clone [haproxy]
Started: [ pod2-stack-controller-0 pod2-stack-controller-1 pod2-stack-controller-2 ]
Master/Slave Set: galera-master [galera]
```

```

Masters: [ pod2-stack-controller-1 pod2-stack-controller-2 ]
Slaves: [ pod2-stack-controller-0 ]
ip-192.200.0.110(ocf::heartbeat:IPaddr2):Started pod2-stack-controller-1
ip-11.120.0.44(ocf::heartbeat:IPaddr2):Started pod2-stack-controller-2
ip-11.118.0.49(ocf::heartbeat:IPaddr2):Started pod2-stack-controller-2
Clone Set: rabbitmq-clone [rabbitmq]
Started: [ pod2-stack-controller-1 pod2-stack-controller-2 ]
Stopped: [ pod2-stack-controller-0 ]
ip-10.225.247.214(ocf::heartbeat:IPaddr2):Started pod2-stack-controller-1
Master/Slave Set: redis-master [redis]
Masters: [ pod2-stack-controller-2 ]
Slaves: [ pod2-stack-controller-0 pod2-stack-controller-1 ]
ip-11.119.0.49(ocf::heartbeat:IPaddr2):Started pod2-stack-controller-2
openstack-cinder-volume(systemd:openstack-cinder-volume):Started pod2-stack-controller-1

Daemon Status:
corosync: active/enabled
pacemaker: active/enabled
pcsd: active/enabled

```

- You can check some of the monitor services such as **ceph** that they are in a healthy state:

```

[heat-admin@pod2-stack-controller-0 ~]$ sudo ceph -s
  cluster eb2bb192-b1c9-11e6-9205-525400330666
health HEALTH_OK
monmap e1: 3 mons at {pod2-stack-controller-0=11.118.0.10:6789/0,pod2-stack-controller-1=11.118.0.11:6789/0,pod2-stack-controller-2=11.118.0.12:6789/0}
election epoch 10, quorum 0,1,2 pod2-stack-controller-0,pod2-stack-controller-1,pod2-stack-controller-2
osdmap e81: 12 osds: 12 up, 12 in
flags sortbitwise,require_jewel_osds
pgmap v22844355: 704 pgs, 6 pools, 804 GB data, 423 kobjects
2404 GB used, 10989 GB / 13393 GB avail
704 active+clean
client io 3658 kB/s wr, 0 op/s rd, 502 op/s wr

```