Troubleshoot PCRF Cluster Manager VM Recovery- Openstack

Contents

Introduction

This document describes the steps to recover Cisco Virtual Policy and Charging Rules Function (vPCRF) instances deployed on Ultra-M/Openstack deployment.

Troubleshoot

Power on Cluster Manager from SHUTOFF State

If any instance is in SHUTOFF state due to a planned shutdown or some other reason, please use this procedure to start the instance and enable it's monitoring in Elastic Services Controller (ESC).

Step 1. Check the state of instance via OpenStack.

```
source /home/stack/destackovsrc-Pcrf
nova list --fields name,host,status | grep cm_0
| c5e4ebd4-803d-45c1-bd96-fd6e459b7ed6 | SVS1-tmo_cm_0_e3ac7841-7f21-45c8-9f86-3524541d6634 |
destackovs-compute-2  | SHUTOFF|
```

Step 2. Check if the compute is available and ensure that the state is up.

```
source /home/stack/destackovsrc
nova hypervisor-show destackovs-compute-2 | egrep 'status|state'
| state                    | up                                    |
| status                   | enabled                               |
```

Step 3. Login to ESC Master as admin user and check the state of instance in opdata.

```
/opt/cisco/esc/esc-confd/esc-cli/esc_nc_cli get esc_datamodel/opdata | grep cm_0
SVS1-tmo_cm_0_e3ac7841-7f21-45c8-9f86-3524541d6634 VM_ERROR_STATE
```

Step 4. Power on the instance from openstack.

```
source /home/stack/destackovsrc-Pcrf
nova start SVS1-tmo_cm_0_e3ac7841-7f21-45c8-9f86-3524541d6634
```

Step 5. Wait five minutes for the instance to boot up and come to active state.

```
source /home/stack/destackovsrc-Pcrf
nova list -fields name,status | grep cm_0
| c5e4ebd4-803d-45c1-bd96-fd6e459b7ed6 | SVS1-tmo_cm_0_e3ac7841-7f21-45c8-9f86-3524541d6634 |
ACTIVE
```

Step 6. Enable the VM Monitor in ESC after the instance is in active state.

```
/opt/cisco/esc/esc-confd/esc-cli/esc_nc_cli vm-action ENABLE_MONITOR SVS1-tmo_cm_0_e3ac7841-
7f21-45c8-9f86-3524541d6634
```

For Further recovery of instance configurations, refer instance type specific procedures provided here.

Recover any Instance from ERROR State

This procedure can be used if state of CPS instance in openstack is ERROR:

Step 1. Check the state of instance in OpenStack.

```
source /home/stack/destackovsrc-Pcrf
nova list --fields name,host,status | grep cm_0
| c5e4ebd4-803d-45c1-bd96-fd6e459b7ed6 | SVS1-tmo_cm_0_e3ac7841-7f21-45c8-9f86-3524541d6634 |
destackovs-compute-2  | ERROR|
```

Step 2. Check if the compute is available and runs fine.

```
source /home/stack/destackovsrc
nova hypervisor-show destackovs-compute-2 | egrep 'status|state'
| state                  | up                                          |
| status                 | enabled                                     |
```

Step 3. Login to ESC Master as admin user and check the state of instance in opdata.

```
/opt/cisco/esc/esc-confd/esc-cli/esc_nc_cli get esc_datamodel/opdata | grep cm_0
SVS1-tmo_cm_0_e3ac7841-7f21-45c8-9f86-3524541d6634 VM_ERROR_STATE
```

Step 4. Reset the state of instance to force the instance back to an active state instead of an error state, once done, reboot your instance.

```
source /home/stack/destackovsrc-Pcrf
nova reset-state -active SVS1-tmo_cm_0_e3ac7841-7f21-45c8-9f86-3524541d6634
nova reboot --hard  SVS1-tmo_cm_0_e3ac7841-7f21-45c8-9f86-3524541d6634
```

Step 5. Wait five minutes for the instance to boot up and come to active state.

```
source /home/stack/destackovsrc-Pcrf
```

```
nova list –fields name,status | grep cm_0
| c5e4ebd4-803d-45c1-bd96-fd6e459b7ed6 | SVS1-tmo_cm_0_e3ac7841-7f21-45c8-9f86-3524541d6634 |
ACTIVE |
```

Step 6. If, Cluster Manager changes state to ACTIVE after reboot, Enable VM Monitor in ESC after Cluster Manager instance is in active state.

```
/opt/cisco/esc/esc-confd/esc-cli/esc_nc_cli vm-action ENABLE_MONITOR SVS1-tmo_cm_0_e3ac7841-
7f21-45c8-9f86-3524541d6634
```

Post recovery to running/active state, refer instance type specific procedure to recover config/data from backup.

Rebuild CPS Cluster Manager through Snapshot

If Cisco Policy Suite (CPS) is stuck in ERROR state and unable to power on through procedures already described and the instance is available in openstack. It is suggested that you rebuild the instance through snapshot image.

Step 1. Ensure that the snapshot of last know good configuration is present as a QCOW file, use this previously generated file during backup, scp/sftp it back to the OpenStack Platform- Director (OSPD) compute. Use this procedure to convert it into a glance image:

```
source /home/stack/destackovsrc-Pcrf
glance image-create --name CPS_Cluman_13.1.1 --disk-format "qcow2" --container "bare" --file
/var/Pcrf/cluman_snapshot.raw

Alternatively,
glance image-create --name rebuild_cluman --file /home/stack/cluman_snapshot.raw --disk-format
qcow2 --container-format bare
```

Step 2. Use a nova rebuild command on OSPD to rebuild the Cluman VM instance with the uploaded snapshot as shown.

```
nova rebuild <instance_name> <snapshot_image_name>
```

Step 3. Wait five minutes for the instance to boot up and come to active state.

```
source /home/stack/destackovsrc-Pcrf
nova list –fields name,status | grep cm
| c5e4ebd4-803d-45c1-bd96-fd6e459b7ed6 |cm_0_170d9c14-0221-4609-87e3-d752e636f57f| ACTIVE |
```

Step 4. If, Cluster Manager changes state to ACTIVE after rebuild, check the state of instance in ESC and Enable VM Monitor in ESC if required.

```
echo "show esc_datamodel opdata tenants tenant Pcrf deployments * state_machine | tab" |
/opt/cisco/esc/confd/bin/confd_cli -u admin –C | grep cm
cm_0_170d9c14-0221-4609-87e3-d752e636f57f VM_ERROR_STATE
/opt/cisco/esc/esc-confd/esc-cli/esc_nc_cli vm-action ENABLE_MONITOR cm_0_170d9c14-0221-4609-
87e3-d752e636f57f
```

Step 5. Verify the Cinder volume associated with Cluster Manager Original ISO image is updated with the current time after the redeploy:

```
cinder list | grep tmobile-pcrf-13.1.1-1.iso
| 2f6d7deb-60d6-40fa-926f-a88536cf98a3 | in-use    | tmobile-pcrf-13.1.1-1.iso | 3    | -
        | true     | a3f3bc62-0195-483a-bbc0-692bccd37307 |
cinder show 2f6d7deb-60d6-40fa-926f-a88536cf98a3 | grep updated_at
| updated_at                | 2018-06-18T08:54:59.000000

updated_at                      | 2018-06-18T08:54:59.000000
```

Step 6. Attach backup disks or any other Cinder volume previously attached to Cluster Manager Instance if not auto attached in previous steps.

```
source /home/stack/destackovsrc-Pcrf

cinder list
```

```
+-------------------------------------+-----------+--------------------------+------+---------
----+----------+------------------------------------+
| ID                                  | Status    | Name                     | Size | Volume
Type | Bootable | Attached to                        |
+-------------------------------------+-----------+--------------------------+------+---------
----+----------+------------------------------------+
| 0e7ec662-b59e-4e3a-91a9-35c4ed3f51d7 | available | pcrf-atp1-mongo02        | 3    | -
        | false    |                                    |
| 2f6d7deb-60d6-40fa-926f-a88536cf98a3 | in-use    | tmobile-pcrf-13.1.1-1.iso | 3    | -
        | true     | a3f3bc62-0195-483a-bbc0-692bccd37307 |
| 4c553948-df75-4f0b-bf7b-0e64127dfda3 | available | pcrf-atp1-svn01          | 3    | -
        | false    |                                    |
| 594c052e-aaa3-4c82-867d-3b36162244b3 | available | tmobile-pcrf-13.1.1-2.iso | 3    | -
        | true     |                                    |
| 64953713-de86-40d5-a0e5-07db22d692f2 | in-use    | tmobile-pcrf-13.1.1.iso  | 3    | -
        | true     | 80a93e90-59e2-43bd-b67e-5d766d0a2f11 |
```

```
openstack server add volume <volume-ID> <Server-ID> --device <location of dev in Instance
example /dev/vdc>
```

Step 7. If the cluman snapshot is old and **config_br.py** backup is available of a date post snapshot was taken. Import the config from backup and if not then skip this step.

```
ssh <cluman-ip>
config_br.py -a import --svn --etc --grafanadb --auth-htpasswd --haproxy /mnt/backup/<file-
name.tgz>
```

Step 8. Rebuild all VM images from backup through **config_br.py** on cluster manager:

```
/var/qps/install/current/scripts/build/build_all.sh
```

Redeploy CPS Cluster Manager through Snapshot

If CPS Cluster Manager VM is lost (unable to recover) and rebuild process (as described it in 2.3) has also failed, you need to redeploy the instance through ESC. This procedure describes the process for the same:

Step 1. Ensure that snapshot of last know good configuration is present as a QCOW file, use this previously generated file during backup, scp/sftp it back to the OSPD compute.

```
ls -ltr /var/Pcrf/cluman_snapshot.qcow
-rw-r--r--. 1 root root 328514100 May 18 16:59 cluman_snapshot.qcow
```

Step 2. Use this procedure to convert it into a glance image.

```
source /home/stack/destackovsrc-Pcrf
glance image-create --name CPS_Cluman_13.1.1 --disk-format "qcow2" --container "bare" --file
/var/Pcrf/cluman_snapshot.qcow
```

Step 3. Once the image is available login to ESC and verify the state of Cluster Manager Instance in ESC opdata.

```
echo "show esc_datamodel opdata tenants tenant Pcrf deployments * state_machine | tab" |
/opt/cisco/esc/confd/bin/confd_cli -u admin -C | grep cm
cm_0_170d9c14-0221-4609-87e3-d752e636f57f VM_ERROR_STATE
```

Step 4. Ensure that the **/home/admin/PCRF_config.xml** file is present as backed up in 2.1.1

Step 5. Get the name of the deployment, tenant and vm_group for cluster manager to be recovered.

Sample Snippet:

```
<esc_datamodel xmlns="http://www.cisco.com/esc/esc">
     <tenants>
       <tenant>
         <name>Pcrf</name>                        --------------- Name of the tenant
         <managed_resource>false</managed_resource>
         <deployments>
           <deployment>
             <name>DEP1</name>        --------------- Name of the Deployment
-----
-----
-----
             <vm_group>
               <name>cm</name>                       --------------- Name of the vm_group
               <image>pcrf-13.1.1.qcow2</image> ------------- Name of the Image used
               <flavor>pcrf-cm</flavor>
               <bootup_time>600</bootup_time>
               <recovery_wait_time>30</recovery_wait_time>
```

Step 6. Trigger a delete of Cluster Manager vm from ESC:

> **Warning**: The command to remove the instance from opdata should be complete, incomplete command can delete the whole deployment. Please be cautious. The command should always contain all of the parameter i.e. tenant name, deployment name and vm_group name.

```
/opt/cisco/esc/confd/bin/confd_cli -u admin -C
esc-ha-01# config
esc-ha-01(config)# no esc_datamodel tenants tenant Pcrf deployments deployment DEP1 vm_group cm
esc-ha-01(config)# commit
esc-ha-01(config)# exit
```

Above step should remove the instance from openstack as well as ESC opdata. In other words, the Cluster Manager is now not a part of deployment.

Step 7. Verify that the Cluster Manager Instance is removed from deployment from **yangesc.log**, **escmanager.log** in ESC and nova list in OSPD node.

Step 8. Modify the **PCRF_config.xml** file backed up in step 2.1.1 and modify the name of the cluster manager image to the newly created image from snapshot in above steps:

Before Change	After Change
<vm_group>	**<vm_group>**
<name>cm</name>	**<name>cm</name>**
<image>pcrf-13.1.1.qcow2</image>	**<image>CPS_Cluman_13.1.1</image>**

Step 9. Modify the **PCRF_config.xml** and remove the cloud user-data file for Cluster Manager vm group. Sample xml snippet that is to be removed is shown here:

```
    <config_data>
                <configuration>
                  <dst>--user-data</dst>
                  <file>file:///opt/cisco/esc/cisco-cps/config/pcrf-cm_cloud.cfg</file>
                  <variable>
                    <name>CLUSTER_ID</name>
                    <val>P1</val>
                  </variable>
                  <variable>
                    <name>CM_IP_ADDR_PVT</name>
                    <val>192.168.1.107</val>
                  </variable>
                  <variable>
                    <name>PREFIX</name>
                    <val>vpc</val>
                  </variable>
                  <variable>
                    <name>SEQ</name>
                    <val>01</val>
                  </variable>
                  <variable>
                    <name>SITE_ID</name>
                    <val>DE</val>
                  </variable>
                </configuration>
              </config_data>
```

Step 10. Copy the file **PCRF_config.xml** to **/opt/cisco/esc/cisco-cps/config/** folder where all other configuration files are present.

Step 11. Load Merge the new configuration file to ESC opdata.

```
/opt/cisco/esc/confd/bin/confd_cli -u admin –C
esc-ha-01# config
esc-ha-01(config)# load merge /opt/cisco/esc/cisco-cps/config/PCRF_config.xml
esc-ha-01(config)# commit
esc-ha-01(config)# exit
```

Step 12. Monitor the **yangesc.log**, **escmanager.log** on ESC and nova list on OSPD to verify deployment of Cluster Manager.

```
source /home/stack/destackovsrc-Pcrf
nova list --fields name,status| grep cm
| 96a5647e-9970-4e61-ab5c-5e7285543a09 | cm_0_a11a9068-df37-4974-9bd8-566f825d5e39    | ACTIVE
```

Step 13. If, Cluster Manager changes state to ACTIVE after rebuild, check the state of instance in ESC and Enable VM Monitor in ESC if required.

```
echo "show esc_datamodel opdata tenants tenant Pcrf deployments * state_machine | tab" |
/opt/cisco/esc/confd/bin/confd_cli -u admin -C | grep cm
cm_0_170d9c14-0221-4609-87e3-d752e636f57f VM_ERROR_STATE
/opt/cisco/esc/esc-confd/esc-cli/esc_nc_cli vm-action ENABLE_MONITOR cm_0_170d9c14-0221-4609-
87e3-d752e636f57f
```

Step 14. Attach backup disks or any other Cinder volume previously attached to Cluster Manager Instance and not auto attached by esc in previous step.

```
source /home/stack/destackovsrc-Pcrf
cinder list
+--------------------------------------+--------+------------------------+------+-------------+--
-------+--------------------------------------+
| ID                                   | Status | Name                   | Size | Volume Type|
Bootable| Attached to                         |
+--------------------------------------+--------+------------------------+------+-------------+--
-------+--------------------------------------+
| 4c478cce-c746-455a-93f1-3f360acb87ce | in-use | CPS_14.0.0.release.iso | 3    | -           |
true    | 96a5647e-9970-4e61-ab5c-5e7285543a09 |
| 7e5573d9-29bc-4ea0-b046-c666bb1f7e06 | in-use | PCRF_backup            | 1024 | -           |
false   |                                     |
| d5ab1991-3e09-41f2-89f5-dd1cf8a9e172 | in-use | svn01                  | 2    | -           |
false   | 09f4bafa-dfb6-457f-9af5-69196eb31b13 |
| d74988a7-1f59-4241-9777-fc4f2d4f3e78 | in-use | svn02                  | 2    | -           |
false   | 86ea448d-09bc-4d2f-81a3-de05884f1e05 |
+--------------------------------------+--------+------------------------+------+-------------+--
-------+--------------------------------------+
openstack server add volume <volume-ID> <Server-ID> --device <location of dev in Instance
example /dev/vdc>
```

Step 15. If the cluman snapshot is old and **config_br.py** backup is available of a date post snapshot was taken. Import the config from backup, if not then skip this step.

```
ssh <cluman-ip>
config_br.py -a import --svn --etc --grafanadb --users --auth-htpasswd --haproxy
/mnt/backup/<file-name.tgz>
```

Step 16. Rebuild all VM images from backup through **config_br.py** on cluster manager:

```
/var/qps/install/current/scripts/build/build_all.sh
```

Verify

- Ping the cluster manager IP to ensure the connectivity is up.
- SSH the cluster manager to check the accessibility.
- Verify the diagnostics from Cluster Manager to ensure the health status of other VMs of CPS is not affected.