Backup and Recovery - Full Cluster Recovery - CPS

Contents

Introduction

This document describes the steps required to Recover an entire CPS Cluster in an Ultra-M setup that hosts CPS Virtual Network Functions (VNFs).

Background Information

Ultra-M is a pre-packaged and validated virtualized mobile packet core solution designed to simplify the deployment of VNFs. Ultra-M solution consists of the following Virtual Machine (VM) types:

- Elastic Services Controller (ESC)
- Cisco Policy Suite (CPS)

The high-level architecture of Ultra-M and the components involved are depicted in this image:
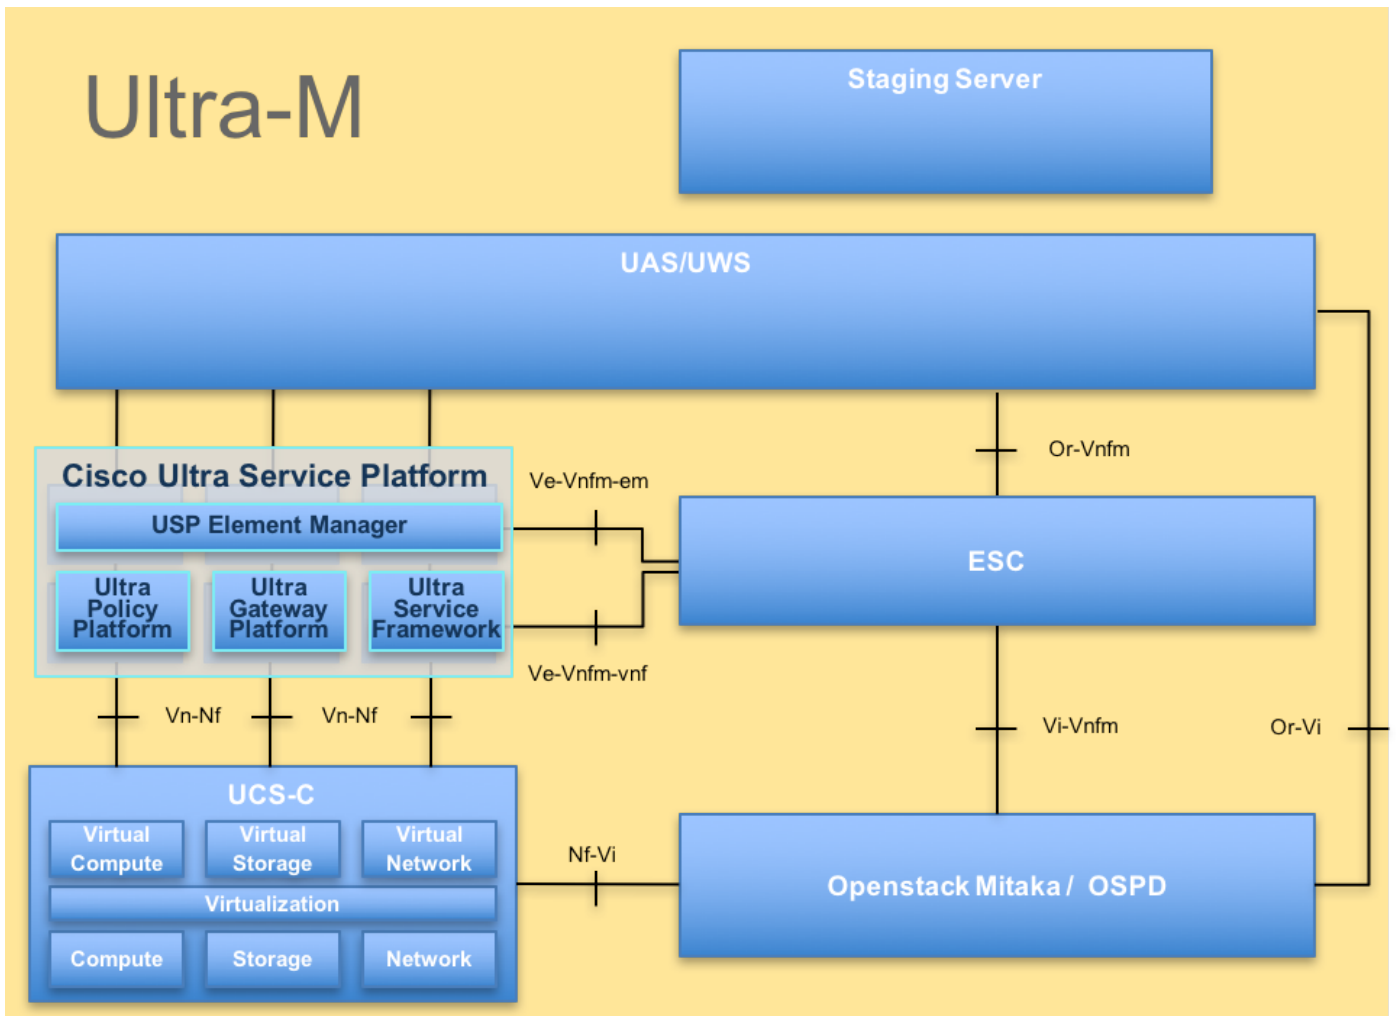
This document is intended for the Cisco personnel familiar with Cisco Ultra-M platform.

Note: Ultra M 5.1.x release is considered for defining the procedures in this document.

Abbreviations

VNF	Virtual Network Function
ESC	Elastic Service Controller
MOP	Method of Procedure
OSD	Object Storage Disks
HDD	Hard Disk Drive
SSD	Solid State Drive
VIM	Virtual Infrastructure Manager
VM	Virtual Machine
UUID	Universally Unique IDentifier

Assumption

For this procedure, the assumption is that only the CPS cluster is to be recovered and all components at the Openstack level are operational including the ESC

Restore Procedure

When ESC Fails to Start VM:

- In some cases, ESC fails to start the VM due to an unexpected state. A workaround is to perform an ESC switchover by rebooting the Master ESC. The ESC switchover takes about a minute. Run **health.sh** on the new Master ESC to verify it is up. When the ESC becomes Master, ESC may fix the VM state and start the VM. Since this operation is scheduled, you must wait 5-7 minutes for it to complete.

- You can monitor **/var/log/esc/yangesc.log** and **/var/log/esc/escmanager.log**. If you do NOT see VM getting recovered after 5-7 minutes, the user would need to go and do the manual recovery of the impacted VM(s).

- Once the VM is successfully restored and running; ensure all the syslog specific configuration is restored from the previous successful known backup. Ensure it is restored in all the ESC VMs

```
root@abautotestvnfm1em-0:/etc/rsyslog.d# pwd
/etc/rsyslog.d

root@abautotestvnfm1em-0:/etc/rsyslog.d# ll

total 28
drwxr-xr-x  2 root root 4096 Jun  7 18:38 ./
drwxr-xr-x 86 root root 4096 Jun  6 20:33 ../]
-rw-r--r--  1 root root  319 Jun  7 18:36 00-vnmf-proxy.conf
-rw-r--r--  1 root root  317 Jun  7 18:38 01-ncs-java.conf
-rw-r--r--  1 root root  311 Mar 17  2012 20-ufw.conf
-rw-r--r--  1 root root  252 Nov 23  2015 21-cloudinit.conf
-rw-r--r--  1 root root 1655 Apr 18  2013 50-default.conf

root@abautotestvnfm1em-0:/etc/rsyslog.d# ls /etc/rsyslog.conf
rsyslog.conf
```

CPS Recovery

1. Create a Backup of CPS Cluster-Manager

Step 1. Use the following command to view the nova instances and note the name of the cluster manager VM instance:

```
nova list
```
Stop the Cluman from ESC.

```
/opt/cisco/esc/esc-confd/esc-cli/esc_nc_cli vm-action STOP <vm-name>
```
Step 2. Verify Cluster Manager in SHUTOFF state.

```
admin@esc1 ~]$ /opt/cisco/esc/confd/bin/confd_cli
```

```
admin@esc1> show esc_datamodel opdata tenants tenant Core deployments * state_machine
```
Step 3. Create a nova snapshot image as shown in this command:

```
nova image-create --poll <cluman-vm-name> <snapshot-name>
```

Note:Ensure that you have enough disk space for the snapshot.

Important - In case if VM becomes unreachable after snapshot creation, check the status of VM using nova list command. If it is in "SHUTOFF" state, you need to start the VM manually.

Step 4. View the image list with the following command: nova image-list Figure 1: Example Output

```
+-------------------------------------+--------------------------------------+--------+-------------------------------------+|
ID                                    | Name                                 |Status |Server
+-------------------------------------+--------------------------------------+--------+-------------------------------------+
|146719e8-d8a0-4d5a-9b15-2a669cfab81f |CPS_10.9.9_20160803_100301_112.iso|ACTIVE|
|1955d56e-4ecf-4269-b53d-b30e73ad57f0 |base_vm                           |ACTIVE|
|2bbfb51c-cd05-4b7c-ad77-8362d76578db |cluman_snapshot                   |ACTIVE|4842ae5a-83a3-48fd-915b-6ca6361adb2c
+-------------------------------------+--------------------------------------+--------+-------------------------------------+
```

Step 5. When a snapshot is created, the snapshot image is stored in OpenStack Glance. To store the snapshot in a remote data store, download the snapshot and transfer the file in OSPD to (/home/stack/CPS_BACKUP)

To download the image, use the following command in OpenStack:

```
glance image-download --file For example: glance image-download --file snapshot.raw 2bbfb51c-
cd05-4b7c-ad77-8362d76578db
```
Step 6. List the downloaded images as shown in the following command:

```
ls -ltr *snapshot*

Example output: -rw-r--r--. 1 root root 10429595648 Aug 16 02:39 snapshot.raw
```
Step 7. Store the snapshot of the Cluster Manager VM to restore in the future.

2. Backup the configuration and database.

```
1. config_br.py -a export --all /var/tmp/backup/ATP1_backup_all_$(date +\%Y-\%m-\%d).tar.gz OR
2. config_br.py -a export --mongo-all /var/tmp/backup/ATP1_backup_mongoall$(date +\%Y-\%m-
\%d).tar.gz
3. config_br.py -a export --svn --etc --grafanadb --auth-htpasswd --haproxy
/var/tmp/backup/ATP1_backup_svn_etc_grafanadb_haproxy_$(date +\%Y-\%m-\%d).tar.gz
4. mongodump - /var/qps/bin/support/env/env_export.sh --mongo /var/tmp/env_export_$date.tgz
5. patches - cat /etc/broadhop/repositories, check which patches are installed and copy those
patches to the backup directory /home/stack/CPS_BACKUP on OSPD
6. backup the cronjobs by taking backup of the cron directory: /var/spool/cron/ from the
Pcrfclient01/Cluman. Then move the file to CPS_BACKUP on the OSPD.
```

Verify from the crontab -l if any other backup is needed

Transfer all the backups to the OSPD /home/stack/CPS_BACKUP

3. Backup yaml file from ESC Master.

```
/opt/cisco/esc/confd/bin/netconf-console --host 127.0.0.1 --port 830 -u <admin-user> -p <admin-
password> --get-config > /home/admin/ESC_config.xml
```

Transfer the file in OSPD /home/stack/CPS_BACKUP

4. Back up crontab -l entries

Create a txt file with crontab -l and ftp it to remote location (in OSPD /home/stack/CPS_BACKUP
)

5. Take a backup of the route files from LB and PCRF client.

```
Collect and scp the below conifgurations from both LBs and Pcrfclients
route -n /etc/sysconfig/network-script/route-*
```

Restore Cluster Manager VM in OpenStack

Step 1. Copy the cluster manager VM snapshot to the controller blade as shown in this command:

```
ls —ltr *snapshot*
```
Example output: -rw-r--r--. 1 root root 10429595648 Aug 16 02:39 snapshot.raw

Step 2. Upload the snapshot image to OpenStack from Datastore:

```
glance image-create --name --file --disk-format qcow2 --container-format bare
```

Step 3. Verify whether the snapshot is uploaded with a Nova command as shown in this example:

```
nova image-list
```
Figure 2: Example Output

```
+--------------------------------------+----------------------------------+--------+--------------------------------------+
| ID                                   | Name                             | Status | Server                               |
+--------------------------------------+----------------------------------+--------+--------------------------------------+
| 146719e8-d8a0-4d5a-9b15-2a669cfab81f | CPS_10.9.9_20160803_100301_112.iso | ACTIVE |                                      |
| 1955d56e-4ecf-4269-b53d-b30e73ad57f0 | base_vm                          | ACTIVE |                                      |
| 2bbfb51c-cd05-4b7c-ad77-8362d76578db | cluman_snapshot                  | ACTIVE | 4842ae5a-83a3-48fd-915b-6ca6361adb2c |
| 5eebff44-658a-49a5-a170-1978f6276d18 | imported_image                   | ACTIVE |                                      |
+--------------------------------------+----------------------------------+--------+--------------------------------------+
```

Step 4. Depending on whether the cluster manager VM exists or not, you can choose to create the
cluman or rebuild the cluman:

If the Cluster Manager VM instance does not exist, create the Cluman VM with a Heat or Nova
command as shown in the following example:

Create the Cluman VM with ESC

```
/opt/cisco/esc/esc-confd/esc-cli/esc_nc_cli edit-config /opt/cisco/esc/cisco-
cps/config/gr/tmo/gen/<original_xml_filename>
```
The PCRF cluster will spawn with the help of the above command, and then restore the cluster manager configurations from the backups taken with config_br.py restore, mongorestore from dump taken in backup

```
delete - nova boot --config-drive true --image "" --flavor "" --nic net-id=",v4-fixed-ip=" --nic
net-id="network_id,v4-fixed-ip=ip_address" --block-device-mapping "/dev/vdb=2edbac5e-55de-4d4c-
a427-ab24ebe66181:::0" --availability-zone "az-2:megh-os2-compute2.cisco.com" --security-groups
cps_secgrp "cluman"
```
If the Cluster Manager VM instance exists, use a nova rebuild command to rebuild the Cluman VM instance with the uploaded snapshot as shown:

```
nova rebuild <instance_name> <snapshot_image_name>
```
For example: nova rebuild cps-cluman-5f3tujqvbi67 cluman_snapshot

Step 5 List all the instances as shown and verify that the new cluster manager instance is created and running:

nova list

Figure 3: Example Output

```
+------------------------------------+--------+--------+------------+-------------+---------------------------------------------+
| ID                                 | Name   | Status | Task State | Power State | Networks                                    |
+------------------------------------+--------+--------+------------+-------------+---------------------------------------------+
| ac3d2dbc-7b0e-4df4-a690-7f84ca3032bd | cluman | ACTIVE | -          | Running     | management=172.20.67.34; internal=172.20.70.34 |
+------------------------------------+--------+--------+------------+-------------+---------------------------------------------+
```

Restore the latest patches on the system

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

> **Note**: The software components must all display Not Monitored as the current status) 8. Update the qns VMs with the new software using reinit.sh script: /var/qps/install/current/scripts/upgrade/reinit.sh 9. Restart all software components on the target VMs: runonall.sh sudo monit start all 10. Verify that the component is updated, run: about.sh

Restore the Cronjobs

1. Move the backed-up file from OSPD to the Cluman/Pcrfclient01

2. Run the command to activate the cronjob from backup
   #crontab Cron-backup
3. Check if the cronjobs have been activated by below command
       #crontab -l

Restore Individual VMs in the Cluster

To Redeploy the pcrfclient01 VM

Step 1. Log in to the Cluster Manager VM as the root user.

Step 2. Note the UUID of SVN repository using the following command:

```
svn info http://pcrfclient02/repos | grep UUID
```
The command will output the UUID of the repository.

For example: Repository UUID: ea50bbd2-5726-46b8-b807-10f4a7424f0e

Step 3. Import the backup Policy Builder configuration data on the Cluster Manager, as shown in the following example:

```
config_br.py -a import --etc-oam --svn --stats --grafanadb --auth-htpasswd --users
/mnt/backup/oam_backup_27102016.tar.gz
```

> **Note**: Many deploymentsrun a cron job that backs up configuration data regularly.SeeSubversion Repository Backup, for more details.

Step 4. To generate the VM archive files on the Cluster Manager using the latest configurations, execute the following command:

```
/var/qps/install/current/scripts/build/build_svn.sh
```
Step 5. To deploy the pcrfclient01 VM, perform one of the following:

In OpenStack, use the HEAT template or the Nova command to re-create the VM. For more information, see CPS Installation Guide for OpenStack.

Step 6. Re-establishSVN master/slave synchronization between the pcrfclient01 and pcrfclient02 with pcrfclient01 as the master as it runs these commands.

If SVN is already synchronized, do not issue these commands.

To check if SVN is in sync, run this command from pcrfclient02.

If a value is returned, then SVN is already in sync:

```
/usr/bin/svn propget svn:sync-from-url --revprop -r0 http://pcrfclient01/repos
```
Run this commands from pcrfclient01:

```
/bin/rm -fr /var/www/svn/repos

/usr/bin/svnadmin create /var/www/svn/repos

/usr/bin/svn propset --revprop -r0 svn:sync-last-merged-rev 0 http://pcrfclient02/repos-proxy-
sync

/usr/bin/svnadmin setuuid /var/www/svn/repos/ "Enter the UUID captured in step 2"
```

```
/etc/init.d/vm-init-client /var/qps/bin/support/recover_svn_sync.sh
```
Step 7. If pcrfclient01 is also the arbiter VM, then run these steps:

1. Create the mongodb start/stop scripts based on the system configuration. Not all deployments have all these databases configured.

> **Note**:Refer to /etc/broadhop/mongoConfig.cfg to determine which databases need to be set up.

```
cd /var/qps/bin/support/mongo

build_set.sh --session --create-scripts
build_set.sh --admin --create-scripts
build_set.sh --spr --create-scripts
build_set.sh --balance --create-scripts
build_set.sh --audit --create-scripts
build_set.sh --report --create-scripts
```

2. Start the mongo process:

```
/usr/bin/systemctl start sessionmgr-XXXXX
```

3. Wait for the arbiter to start, then run diagnostics.sh --get_replica_status to check the health of the replica set.

To redeploy the pcrfclient02 VM

Step 1. Log in to the Cluster Manager VM as the root user.

Step 2. To generate the VM archive files on the Cluster Manager using the latest configurations, run this command:

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

Step 3 To deploy the pcrfclient02 VM, perform one of the following:

In OpenStack, use the HEAT template or the Nova command to re-create the VM. For more information, see CPS Installation Guide for OpenStack.

Step 4 Secure shell to the pcrfclient01:

```
ssh pcrfclient01
```

Step 5 Run this script to recover the SVN repos from pcrfclient01:

```
/var/qps/bin/support/recover_svn_sync.sh
```

To redeploy a sessionmgr VM

Step 1. Login to the Cluster Manager VM as the root user

Step 2. To deploy the sessionmgr VM and replace the failed or corrupt VM, perform one of these:

In OpenStack, use the HEAT template or the Nova command to re-create the VM. For more information, see CPS Installation Guide for OpenStack

Step 3. Create the mongodb start/stop scripts based on the system configuration.

Not all deployments have all these databases configured. Refer to /etc/broadhop/mongoConfig.cfg to determine which databases need to be set up

```
cd /var/qps/bin/support/mongo

build_set.sh --session --create-scripts
build_set.sh --admin --create-scripts
build_set.sh --spr --create-scripts
build_set.sh --balance --create-scripts
build_set.sh --audit --create-scripts
build_set.sh --report --create-scripts
```

Step 4. Secure shell to the sessionmgr VM and start the mongo process:

```
ssh sessionmgrXX

/usr/bin/systemctl start sessionmgr-XXXXX
```

Step 5. Wait for the members to start and for the secondary members to synchronize, then run diagnostics.sh --get_replica_status to check the health of the database.

Step 6. To restore Session Manager database, use one of the following example commands depending on whether the backup was performed with --mongo-all or --mongo option:

- `config_br.py -a import --mongo-all --users /mnt/backup/Name of backup`

`or`

- `config_br.py -a import --mongo --users /mnt/backup/Name of backup`

To Redeploy the Policy Director (Load Balancer) VM

Step 1. Log in to the Cluster Manager VM as the root user.

Step 2. To import the backup Policy Builder configuration data on the Cluster Manager, run this command:

```
config_br.py -a import --network --haproxy --users /mnt/backup/lb_backup_27102016.tar.gz
```

Step 3 To generate the VM archive files on the Cluster Manager using the latest configurations, run this command:

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

Step 4. To deploy the lb01 VM, perform one of these:

In OpenStack, use the HEAT template or the Nova command to re-create the VM. For more

information, see CPS Installation Guide for OpenStack.

To Redeploy the Policy Server (QNS) VM

Step 1. Log in to the Cluster Manager VM as the root user.

Step 2. Import the backup Policy Builder configuration data on the Cluster Manager, as shown in this example:

```
config_br.py -a import --users /mnt/backup/qns_backup_27102016.tar.gz
```

Step 3. To generate the VM archive files on the Cluster Manager using the latest configurations, run this command:

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

Step 4 To deploy the qns VM, perform one of this:

In OpenStack, use the HEAT template or the Nova command to re-create the VM. For more information, see CPS Installation Guide for OpenStack

General Procedure for Database Restore

Step 1. Run this command to restore the database:

```
config_br.py -a import --mongo-all /mnt/backup/backup_$date.tar.gz where $date is the timestamp
when the export was made.
```
For example,

```
config_br.py -a import --mongo-all /mnt/backup/backup_27092016.tgz
```
Step 2. Login to the database and verify whether it is running and is accessible:

1. Log into session manager:

```
mongo --host sessionmgr01 --port $port
```
where $port is the port number of the database to check. For example, 27718 is the default Balance port.

2. Display the database by executing the following command:

```
show dbs
```
3. Switch the mongo shell to the database by executing the following command:

```
use $db
```
where $db is a database name displayed in the previous command.

The **use** command switches the mongo shell to that database.

For example,

```
use balance_mgmt
```

4. In order to display the collections, run this command:

```
show collections
```

5. In order to display the number of records in the collection, run this command:

```
db.$collection.count()
```

```
For example, db.account.count()
```

The above example will show the number of records in the collection "account" in the Balance database (balance_mgmt).

Subversion Repository Restore

To restore the Policy Builder Configuration Data from a backup, execute the following command:

```
config_br.py –a import --svn /mnt/backup/backup_$date.tgz where, $date is the date when the cron
created the backup file.
```

Restore Grafana Dashboard

You can restore Grafana dashboard using the following command:

```
config_br.py -a import --grafanadb /mnt/backup/
```

Validate the Restore

After you restore the data, verify the working system through this command:

/var/qps/bin/diag/diagnostics.sh

When ESC Fails to Start VM

- In some cases, ESC fails to start the VM due to an unexpected state. A workaround is to perform an ESC switchover by rebooting the Master ESC. The ESC switchover takes about a minute. Run **health.sh** on the new Master ESC to verify it is up. When the ESC becomes Master, ESC may fix the VM state and start the VM. Since this operation is scheduled, you must wait 5-7 minutes for it to complete.

- You can monitor **/var/log/esc/yangesc.log** and **/var/log/esc/escmanager.log**. If you do NOT see VM getting recovered after 5-7 minutes, the user would need to go and do the manual recovery of the impacted VM(s).

- If the Cluster is completely unavailable and only ESC is reachable, then restore has to be

performed from the latest backups which are taken from the scheduled backups done through Cronjobs. The procedure for recovery remains the same as mentioned in the MOP.