Migrate CPS

For release 18.2.0, the Cisco Policy Suite (CPS) is built on a newer version of CentOS. Previous versions of the CPS platform used CentOS 6.8; however CPS 18.2.0 uses CentOS 7.4. Because of this change, an in-service software upgrade (ISSU) is not possible. If customers want to move to CPS 18.2.0, they must perform an in-service migration, which has been designed to migrate their system with minimal disruption of traffic.

In-Service Migration to 18.2.0

This section describes the steps to perform an in-service software migration of a CPS 12.1.0, CPS 13.x.x, CPS 14.0.0, and CPS 18.0.0 to CPS 18.2.0. This migration allows the traffic to continue running while the migration is being performed.

In-service software migrations to 18.2.0 are supported only for Mobile (HA) and GR installations. Other CPS installation types cannot be migrated.

Prerequisites

☞

Important During the migration process, do not make policy configuration changes, CRD table updates, or other system configuration changes. These type of changes should only be performed after the migration has been successfully completed and properly validated.

Before beginning the migration:

1. Create a backup (snapshot/clone) of the Cluster Manager VM following the guidelines of the prior release. If errors occur during the migration process, this backup is required to successfully roll back the migration. For more information refer to *CPS Backup and Restore Guide*.

2. Back up any nonstandard customizations or modifications to system files. Only customizations which are made to the configuration files on the Cluster Manager are backed up. Refer to the *CPS Installation Guide for VMware* for an example of this customization procedure. Any customizations which are made directly to the CPS VMs must be reapplied manually after the migration is complete.

3. Remove any previously installed patches.

4. If necessary, upgrade the underlying hypervisor before performing the CPS in-service software migration. The steps to upgrade the hypervisor or troubleshoot any issues that may arise during the hypervisor upgrade is beyond the scope of this document. Refer to the appropriate CPS installation guide for a list of supported hypervisors for this CPS release.

5. Verify that the Cluster Manager VM has at least 10 GB of free space. The Cluster Manager VM requires this space when it creates the backup archive at the beginning of the migration process.

6. Synchronize the Grafana information between the OAM (pcrfclient) VMs by running the following command from pcrfclient01:

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

   Also verify that the `/var/broadhop/.htpasswd` files are the same on pcrfclient01 and pcrfclient02 and copy the file from pcrfclient01 to pcrfclient02 if necessary.

   Refer to *Copy Dashboards and Users to pcrfclient02* in the *CPS Operations Guide* for more information.

7. Check the health of the CPS cluster as described in

8. The contents of `logback.xml` file are overwritten during an upgrade or a migration. Make sure to update the `logback.xml` file as per your requirements after an upgrade or a migration.

Refer also to for more information about the process to restore a CPS cluster to the previous version if the migration is not successful.

✎

Note Cisco Smart Licensing is supported for CPS 10.0.0 and later releases. For information about Smart Licensing and how to enable it for CPS, refer to the *CPS Operations Guide*.

Overview

The in-service software migration is performed in the following general steps:

1. Download and mount the CPS software on the Cluster Manager VM.

2. Migrate the Cluster Manager VM – Relevant data is backed up from the old Cluster Manager VM in addition to the other CPS VMs, and stored in a tar file. Then the old Cluster Manager can be terminated and brought back up with the new 18.2.0 base image and the same IP address from the old Cluster Manager. The backed up data is then be restored on the new Cluster Manager.

3. Migrate CPS VMs Set 1 – The rest of the CPS VMs are split in half. The first set of CPS VMs, Set 1, can then be terminated and brought back up with the new 18.2.0 base image. The new VMs are then enabled and restored with the relevant data that was backed up.

4. Migrate CPS VMs Set 2 – After the first set of CPS VMs have been brought back up, the second set are then terminated and brought back up using the 18.2.0 base image. The new CPS VMs are then enabled and restored with the relevant data that was backed up.

Check the System Health

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

Step 2 Check the health of the system by running the following command:

```
diagnostics.sh
```

Clear or resolve any errors or warnings before proceeding.

Download the CPS ISO Image

Step 1 Download the Full Cisco Policy Suite Installation software package (ISO image) from software.cisco.com. Refer to *CPS Release Notes* for the download link.

Step 2 Load the ISO image on the Cluster Manager.

For example:

```
wget http://linktoisomage/CPS_x.x.x.release.iso
```

where,

`linktoisoimage` is the link to the website from where you can download the ISO image.

`CPS_x.x.x.release.iso` is the name of the Full Installation ISO image.

Create a Backup of CPS 12.1.0/CPS 13.x.x/14.0.0/18.0.0 Cluster Manager

Before migrating Cluster Manager to CPS 18.2.0, create a backup of the current Cluster Manager in case an issue occurs during migration.

Step 1 On Cluster Manager, remove the following files if they exist:

```
* /etc/udev/rules.d/65-cps-ifrename.rules
* /etc/udev/rules.d/70-persistent-net.rules
```

Step 2 Create a backup (snapshot/clone) of Cluster Manager. For more information, refer to the *CPS Backup and Restore Guide*.

Migrate the Cluster Manager VM

This section describes how to migrate the Cluster Manager VM to CPS 18.2.0.

Note Diagnostic fails during migration. This is normal since NTP may be converging, mongo replica sets are not synced, and so on. If you see HAProxy diagnostics warnings about Diameter endpoints being down, see HAProxy Diagnostics Warnings, on page 21 for a workaround.

Before you begin

For VMware based setup, check `Configuration.csv` under `/var/qps/config/deploy/csv/` and confirm whether `db_authentication_enabled` parameter is present in the file. For migration to succeed, `db_authentication_enabled,FALSE,` must be configured in `Configuration.csv` file.

- The migration succeeds:

    - If `db_authentication_enabled` is disabled as `db_authentication_enabled,FALSE,` OR the parameter is enabled as `db_authentication_enabled,TRUE,`

    - db_authentication_admin_passwd,<xxxxxxx>,

    - db_authentication_readonly_passwd,<xxxxx>,

- If the parameter `db_authentication_enabled` is not present in the file, you need to configure it as `db_authentication_enabled,FALSE,` for migration to succeed.

This is mongo authentication related feature. For more information, see *CPS Installation Guide for VMware*.

Step 1 Unmount the old CPS ISO by running the following command:

```
umount /mnt/iso
```

Step 2   Mount the new CPS 18.2.0 ISO to the existing CPS Cluster Manager running the following command:

```
mount -o loop CPS_x.x.x.release.iso /mnt/iso
```

Step 3   Back up the Cluster Manager by running the following command:

```
/mnt/iso/migrate.sh backup cluman
```

After the backup has run successfully, you should see messages like the following:

```
2017-07-14 02:39:07,842 INFO [backup.etc] Backup: etc
2017-07-14 02:39:07,878 INFO [__main__.run_recipe] Performing installation stage: Create backup Tar
2017-07-14 02:39:07,905 INFO [__main__.<module>] =====================
2017-07-14 02:39:07,905 INFO [__main__.<module>]   SUCCESS
2017-07-14 02:39:07,905 INFO [__main__.<module>] ======== END ========
```

> **Important**   Back up any nonstandard customization or modifications to system files and configuration files that are not a part of the default configuration (/etc/broadhop/).

Step 4   After the Cluster Manager data has been backed up, copy the tar.gz file to an external location or control node as shown in the following example:

For example:

```
sftp root@172.16.2.19
sftp> get migrate_cluman_20170105_170515.tar.gz
Fetching /var/tmp/migrate_cluman_20170105_170515.tar.gz to migrate_20170105_170515.tar.gz
/var/tmp/migrate_cluman_20170105_170515.tar.gz
```

In this example, 172.16.2.19 is the internal IP address of the Cluster Manager VM.

Step 5   For VMware, deploy the CPS 18.2.0 Cluster Manager VM following the instructions provided in the *CPS Installation Guide for VMware* or *CPS Installation Guide for OpenStack* depending on your deployment.

> **Note**   Preserve the old Cluster Manager and create a new Cluster Manager with CPS 18.2.0 as new deployment. Deploy the CPS 18.2.0 Cluster Manager by referring to *CPS Installation Guide for VMware* or *CPS Installation Guide for OpenStack* depending on your deployment.

> **Important**   The VM is rebooted in rescue mode for the first time for CentOS to adjust disk/hardware to the new version. Subsequent reboots if necessary is a normal operation.

For Openstack, it is mandatory to delete the previously deployed Cluster Manager in order to deploy the new Cluster Manager. If the previously deployed Cluster Manager is not deleted, new Cluster Manager deployment fails.

Step 6   After the deployment has been completed, check its status using the Status API.

For example:

```
URL : http://<<cluster-ip>>:8458/api/system/status/cluman
Eg:http://172.18.11.151:8458/api/system/status/cluman
Header: Content-Type:application/json
Success Message: {
 "status": "ready"
}
```

Step 7 Copy the migrate `tar.gz` file from the external location to the new CPS 18.2.0 Cluster Manager, and run the `/mnt/iso/migrate.sh restore cluman <full_path>/migrate_<date_and_time>.tar.gz` command as shown in the following example.

```
sftp> put migrate_cluman_20170120_200701.tar.gz on cluman.
cd /mnt/iso
./migrate.sh restore cluman /root/migrate_20170120_200701.tar.gz
```

When the restore has completed, you should see messages like the following:

```
2017-01-21 01:42:21,497 INFO [restore_cluman.restore_fingerprints] Restore fingerprint files.
2017-01-21 01:42:21,531 INFO [restore_cluman.restore_logs] Restoring and copying migrated logs to
archive directory.
2017-01-21 01:42:21,532 INFO [restore_cluman.restore_env_config] Restore cluman env_config files.
2017-01-21 01:42:22,441 INFO [restore_cluman.restore_config_br] Restore cluman config_br files.
2017-01-21 01:42:22,441 INFO [backup.handleRequest] Action Import
2017-01-21 01:42:22,443 INFO [backup.etc] Restore: etc
2017-01-21 01:42:22,544 INFO [__main__.<module>] =====================
2017-01-21 01:42:22,544 INFO [__main__.<module>]   SUCCESS
2017-01-21 01:42:22,544 INFO [__main__.<module>] ======== END ========
```

> **Important** After restoring Cluster Manager, manually reapply any nonstandard customizations or modifications that were done previously; for example, system files/configuration files (which were backed up in Step 3, on page 5).

Step 8 Run the `about.sh` and `diagnostics.sh` scripts to verify that Cluster Manager is able to communicate with other VMs. For example:

```
about.sh
Cisco Policy Suite - Copyright (c) 2015. All rights reserved.

CPS Multi-Node Environment

CPS Installer Version - 18.2.0
CPS Core Versions
--------------------------
         lb01: qns-1            (iomanager): 18.0.0.release
         lb01: qns-2    (diameter_endpoint): 18.0.0.release
         lb01: qns-3    (diameter_endpoint): 18.0.0.release
```

In the example, you can see that the CPS Installer Version was migrated to 18.2.0, but the VMs still have the old version, since they have not yet been migrated.

You can also verify the time zone and the CentOS version as shown in the following example:

```
cat /etc/redhat-release
CentOS Linux release 7.4.1708 (Core)
```

Step 9 You can disable syncing of carbon database and bulk statistics files to decrease the ISSM time by adding the following parameters in `/var/install.cfg` file.

- SKIP_BLKSTATS

- SKIP_CARBONDB

Example to disable:

```
SKIP_BLKSTATS=1
SKIP_CARBONDB=1
```

Migrate CPS Set 1 VMs

Once Cluster Manager has been migrated, the migration of the CPS VMs can be started. To do this, the CPS cluster must be divided into two sets: Set 1 and Set 2 (similar to what is done during an ISSU). Set 1 is migrated first, as described in this section. After the migration of Set 1, if there are no call drops, you can continue with the migration of Set 2 VMs. However, if there is a failure after migrating Set 1, you must perform a migration rollback.

Note Diagnostic fails during migration. This is normal since NTP may be converging, mongo replica sets are not synced, and so on. If you see HAProxy diagnostics warnings about Diameter endpoints being down, see HAProxy Diagnostics Warnings, on page 21 for a workaround.

You can disable syncing of carbon database and bulk statistics files to decrease the ISSM time. For more information, refer to Disable Syncing Carbon Database and Bulk Stats Files, on page 20.

Step 1 Run the create-cluster-sets command to create the cluster sets for migration:

/var/platform/platform/scripts/create-cluster-sets.sh

You should see the following output:

```
Created /var/tmp/cluster-upgrade-set-1.txt
Created /var/tmp/cluster-upgrade-set-2.txt
```

Step 2 (Optional) You can reduce the migration time by provisioning the VMs. If you do not want to provision the VMs, go to Step 3, on page 7.

Note The VM provisioning requires extra disk space for each VM. Provisioning can be done only for VMware environment setups.

a) Open a separate terminal and run the following command to provision Set 1 VMs:

```
/var/qps/install/current/scripts/deployer/support/deploy_all.py --provision --vms
/var/tmp/cluster-upgrade-set-1.txt
```

This command can be run in parallel to disabling Set 1.

Note Manually enter `deploy_all.py` command in your system.

Step 3 Run the following command to disable Set 1 VMs:

/mnt/iso/migrate.sh disable set 1

When Set 1 has been disabled, you should see messages like the following:

```
2017-01-21 01:53:49,894 INFO [__main__.extra_banner]
        ====================================================================
        | Backing up to file: /var/tmp/migrate_set-1_20170120_212456.tar.gz
```

```
                =====================================================================
2017-01-21 02:00:12,252 INFO [backup.handleRequest]
==================================================
2017-01-21 02:00:12,253 INFO [backup.handleRequest] Archive
/var/tmp/migrate/set-1/config_other_br.tar.gz is created with requested backups.
2017-01-21 02:00:12,253 INFO [backup.handleRequest]
==================================================
2017-01-21 02:00:12,253 INFO [__main__.run_recipe] Performing installation stage: Create backup Tar
2017-01-21 02:00:12,577 INFO [__main__.<module>] ====================
2017-01-21 02:00:12,578 INFO [__main__.<module>]   SUCCESS
2017-01-21 02:00:12,578 INFO [__main__.<module>] ======== END ========
```

Step 4   Confirm that the Set 1 VMs' sessionmgrs are removed from the replica sets by running the following command:

diagnostics.sh --get

Example output is shown below:

```
CPS Diagnostics HA Multi-Node Environment
---------------------------
Checking replica sets...
|----------------------------------------------------------------------------------------------|
| Mongo:3.2.13            MONGODB REPLICA-SETS STATUS INFORMATION        Date : 2017-01-21 02:00:27
|
|----------------------------------------------------------------------------------------------|
|  SET NAME  -  PORT  : IP ADDRESS - REPLICA STATE - HOST NAME  -  HEALTH  - LAST SYNC - PRIORITY
|
|----------------------------------------------------------------------------------------------|
| ADMIN:set06
|
|  Member-1  -  27721 : 172.16.2.100  - ARBITER  - arbitervip  - ON-LINE  -  --------  -   1
|
|  Member-2  -  27721 : 172.16.2.22   - PRIMARY  - sessionmgr01 - ON-LINE  - No Sync    -   3
|
|----------------------------------------------------------------------------------------------|
| AUDIT:set05
|
|  Member-1  -  27017 : 172.16.2.100  - ARBITER  - arbitervip  - ON-LINE  -  --------  -   1
|
|  Member-2  -  27017 : 172.16.2.22   - PRIMARY  - sessionmgr01 - ON-LINE  - No Sync    -   3
|
|----------------------------------------------------------------------------------------------|
| BALANCE:set02
|
|  Member-1  -  27718 : 172.16.2.100  - ARBITER  - arbitervip   - ON-LINE  -  --------  -   1
|
|  Member-2  -  27718 : 172.16.2.22   - PRIMARY  - sessionmgr01 - ON-LINE  - No Sync    -   3
|
|----------------------------------------------------------------------------------------------|
| REPORTING:set03
|
|  Member-1  -  27719 : 172.16.2.100  - ARBITER  - arbitervip   - ON-LINE  -  --------  -   1
|
|  Member-2  -  27719 : 172.16.2.22   - PRIMARY  - sessionmgr01 - ON-LINE  - No Sync    -   3
|
|----------------------------------------------------------------------------------------------|
| SESSION:set01
|
|  Member-1  -  27717 : 172.16.2.100  - ARBITER  - arbitervip   - ON-LINE  -  --------  -   1
|
|  Member-2  -  27717 : 172.16.2.22   - PRIMARY  - sessionmgr01 - ON-LINE  - No Sync    -   3
|
|----------------------------------------------------------------------------------------------|
| SPR:set04
```

```
|
|  Member-1  -  27720 : 172.16.2.100  -  ARBITER  -  arbitervip   - ON-LINE  -   --------  -   1
|
|  Member-2  -  27720 : 172.16.2.22   -  PRIMARY  -  sessionmgr01 - ON-LINE  -  No Sync    -   3
|
|------------------------------------------------------------------------------------------|
```

Step 5 If you have provisioned VMs using Step 2, on page 7, you can restart VM using provisioned vmdk image by running the following command and then go to Step 7, on page 10:

```
/var/qps/install/current/scripts/deployer/support/deploy_all.py --useprovision --vms
/var/tmp/cluster-upgrade-set-1.txt
```

> **Note** If you have not provisioned VMs, go to Step 6, on page 9.

> **Note** Manually enter `deploy_all.py` command in your system.

Step 6 Re-deploy the Set 1 VMs.

> **Note** Delete Set 1 VMs before re-deploying them with the new base.vmdk.

> **Note** To install the VMs using shared or single storage, you must use
> `/var/qps/install/current/scripts/deployer/deploy.sh` *$host* command.
>
> For more information, refer to *Manual Deployment* section in *CPS Installation Guide for VMware*.

For VMware: `/var/qps/install/current/scripts/deployer/support/deploy_all.py --vms`
`/var/tmp/cluster-upgrade-set-1.txt`

> **Note** Manually enter `deploy_all.py` command in your system.

For OpenStack: Use nova boot commands or Heat templates. For more information, refer to *CPS Installation Guide for OpenStack*.

Example deploying Set 1 with Openstack using nova boot command: The commands given below are for reference purpose only. The user must type the commands manually.

```
nova boot --config-drive true --user-data=pcrfclient02-cloud.cfg --image "new_base_vm" --flavor
"pcrfclient02"  --nic net-id="c3df93c2-c2bb-4143-8bb6-b4ec6df65e53,v4-fixed-ip=172.16.2.21" --nic
net-id="4759babe-491a-4c1a-a028-ec4daefa1662,v4-fixed-ip=172.18.11.153" --block-device-mapping
"/dev/vdb=50914841-70e5-44c1-9be6-019f96a3b9fe:::0" "pcrfclient02" --availability-zone
az-2:os8-compute-2.cisco.com

nova boot --config-drive true --user-data=sessionmgr02-cloud.cfg  --image "new_base_vm" --flavor
"sm"  --nic net-id="c3df93c2-c2bb-4143-8bb6-b4ec6df65e53,v4-fixed-ip=172.16.2.23" --nic
net-id="4759babe-491a-4c1a-a028-ec4daefa1662,v4-fixed-ip=172.18.11.158" --block-device-mapping
"/dev/vdb=73436f2b-2c93-4eb1-973c-8490015b41b5:::0" "sessionmgr02" --availability-zone
az-2:os8-compute-2.cisco.com

nova boot --config-drive true --user-data=lb02-cloud.cfg --image "new_base_vm" --flavor "lb02" --nic
 net-id="c3df93c2-c2bb-4143-8bb6-b4ec6df65e53,v4-fixed-ip=172.16.2.202" --nic
net-id="4759babe-491a-4c1a-a028-ec4daefa1662,v4-fixed-ip=172.18.11.155" --nic
net-id="392b72f6-b8f1-47b2-ae5f-e529f69866bc,v4-fixed-ip=192.168.2.202" "lb02"  --availability-zone
 az-2:os8-compute-2.cisco.com

nova boot --config-drive true --user-data=qns02-cloud.cfg --image "new_base_vm" --flavor "qps" --nic
 net-id="c3df93c2-c2bb-4143-8bb6-b4ec6df65e53,v4-fixed-ip=172.16.2.25" "qns02" --availability-zone
```

```
 az-2:os8-compute-2.cisco.com
```

```
nova boot --config-drive true --user-data=qns04-cloud.cfg --image "new_base_vm" --flavor "qps" --nic
 net-id="c3df93c2-c2bb-4143-8bb6-b4ec6df65e53,v4-fixed-ip=172.16.2.27" "qns04" --availability-zone
 az-2:os8-compute-2.cisco.com
```

Important After deployment of load balancer VM, verify monit service status by executing the following command on deployed Load Balancer (lb) VM:

```
/bin/systemctl status monit.service
```

If monit service on load balancer VM is not running, then execute the following command on that VM to start it:

```
/bin/systemctl start monit.service
```

Step 7 If you are using OpenStack, assign:

- arbitervip to pcrfclient02 internal IP

- lbvip01 to lb02 management IP

- lbvip02 to lb02 internal IP

- Gx VIP to lb02 Gx IP

Example Assigning VIPs to Set 1 VMs using neutron port command: The commands given below are for reference purpose only. The user must type the commands manually.

```
[root@os8-control cloud(keystone_core)]# neutron port-list | grep "172.16.2.21"
| 3d40e589-993c-44b5-bb0a-0923a4abbfc0 |                                          |
fa:16:3e:5e:24:48 | {"subnet_id": "106db79e-da5a-41ea-a654-cffbc6928a56", "ip_address": "172.16.2.21"}
  |
[root@os8-control cloud(keystone_core)]# neutron port-update 3d40e589-993c-44b5-bb0a-0923a4abbfc0
--allowed-address-pairs type=dict list=true ip_address=172.16.2.100
Updated port: 3d40e589-993c-44b5-bb0a-0923a4abbfc0
[root@os8-control cloud(keystone_core)]# neutron port-list | grep "172.18.11.155"
| ca9ece72-794c-4351-b7b8-273ec0f81a98 |                                          |
fa:16:3e:9e:b9:fa | {"subnet_id": "641276aa-245f-46db-b326-d5017915ccf7", "ip_address":
"172.18.11.155"} |
[root@os8-control cloud(keystone_core)]# neutron port-update ca9ece72-794c-4351-b7b8-273ec0f81a98
--allowed-address-pairs type=dict list=true ip_address=172.18.11.156
Updated port: ca9ece72-794c-4351-b7b8-273ec0f81a98
[root@os8-control cloud(keystone_core)]# neutron port-list | grep "172.16.2.202"
| 2294991c-22a6-43c6-b846-2ec9c75c6bf8 |                                          |
fa:16:3e:0b:8c:b0 | {"subnet_id": "106db79e-da5a-41ea-a654-cffbc6928a56", "ip_address":
"172.16.2.202"}  |
[root@os8-control cloud(keystone_core)]# neutron port-update 2294991c-22a6-43c6-b846-2ec9c75c6bf8
--allowed-address-pairs type=dict list=true ip_address=172.16.2.200
Updated port: 2294991c-22a6-43c6-b846-2ec9c75c6bf8
[root@os8-control cloud(keystone_core)]# neutron port-list | grep "192.168.2.202"
| d6c82358-4755-47f4-bc64-995accbe0ea6 |                                          |
fa:16:3e:6c:47:a6 | {"subnet_id": "263ba6d1-31b0-450a-9a2d-30418f3476f9", "ip_address":
"192.168.2.202"} |
[root@os8-control cloud(keystone_core)]# neutron port-update d6c82358-4755-47f4-bc64-995accbe0ea6
--allowed-address-pairs type=dict list=true ip_address=192.168.2.200
Updated port: d6c82358-4755-47f4-bc64-995accbe0ea6
```

For more information, refer to *CPS Installation Guide for OpenStack*.

Important   The VMs are rebooted in rescue mode for the first time for CentOS to adjust disk/hardware to the new version. Subsequent reboots if necessary is a normal operation.

Step 8   Once the VMs are Powered ON, run the following command to enable Set 1 VMs:

migrate.sh enable set 1 /var/tmp/migrate-set-1-<*timestamp*>.tar.gz

For example:

```
/mnt/iso/migrate.sh enable set 1 /var/tmp/migrate_set-1_20170120_212456.tar.gz
```

Note   The migration does not restore users created with `adduser.sh` due to potential gid/uid conflicts. Check the migrate enable log for entries that indicate users that are not being migrated, and then manually recreate them using `adduser.sh`. An example log is shown below:

```
2017-02-13 14:52:15,999 INFO [etc_passwd.parse_etc_passwd] Parsing
/var/tmp/migrate/pcrfclient02/etc/passwd file
2017-02-13 14:52:16,000 INFO [etc_group.parse_etc_group] Parsing
/var/tmp/migrate/pcrfclient02/etc/group file
2017-02-13 14:52:16,000 WARNING [restore_vm.restore_vms] On Host:pcrfclient02 User/Group
mongoreadonly/mongoreadonly is not being migrated and must be manually created using
adduser.sh.
2017-02-13 14:52:16,000 WARNING [restore_vm.restore_vms] On Host:pcrfclient02 User/Group
admin/admin is not being migrated and must be manually created using adduser.sh.
```

After the script has run, you should see information like the following:

```
WARNING  Mongo Server trying to reconnect while pushing config. Attempt #1
INFO     Priority set operation is completed for SPR-SET1
INFO     Priority set to the Database members is finished
INFO     Validating if Priority is set correctly for Replica-Set: SPR-SET1
INFO     Validated Priority is set correctly for Replica-Set: SPR-SET1

2017-01-21 02:45:48,950 INFO [__main__.<module>] ====================
2017-01-21 02:45:48,950 INFO [__main__.<module>]  SUCCESS
2017-01-21 02:45:48,951 INFO [__main__.<module>] ======== END ========
```

Step 9   Verify that Set 1 VMs have been migrated by running `about.sh` command:

about.sh

Example output is shown below:

```
CPS Core Versions
--------------------------
        lb01: qns-1          (iomanager): 18.0.0.release
        lb01: qns-2   (diameter_endpoint): 18.0.0.release
        lb01: qns-3   (diameter_endpoint): 18.0.0.release
        lb01: qns-4   (diameter_endpoint): 18.0.0.release
        lb02: qns-1          (iomanager): 18.2.0.release
        lb02: qns-2   (diameter_endpoint): 18.2.0.release
        lb02: qns-3   (diameter_endpoint): 18.2.0.release
        lb02: qns-4   (diameter_endpoint): 18.2.0.release
       qns01: qns-1               (pcrf): 18.0.0.release
       qns02: qns-1               (pcrf): 18.2.0.release
       qns03: qns-1               (pcrf): 18.0.0.release
       qns04: qns-1               (pcrf):
  pcrfclient01: qns-1      (controlcenter): 18.0.0.release
  pcrfclient01: qns-2                 (pb): 18.0.0.release
  pcrfclient02: qns-1      (controlcenter): 18.2.0.release
  pcrfclient02: qns-2                 (pb): 18.2.0.release
```

Step 10 Migrate traffic swap by running the following command: Check for call traffic to determine if you can proceed with the migration of Set 2 VMs.

```
/mnt/iso/migrate.sh traffic swap
```

After the traffic swap has run, you should see information like the following:

```
Creating MD5 Checksum...
Redis config updated on installer

2017-01-23 19:56:09,092 INFO [__main__.<module>] ====================
2017-01-23 19:56:09,092 INFO [__main__.<module>]  SUCCESS
2017-01-23 19:56:09,092 INFO [__main__.<module>] ======== END ========
```

If the script ran successfully, you can proceed with the migration of Set 2 VMs. If not, you must roll back Set 1 as described in Migration Rollback, on page 22.

What to do next

If some of the replica-set members are in RECOVERING state, refer to Recover Replica-set Members from RECOVERING State, on page 17.

Migrate CPS Set 2 VMs

After you have successfully migrated the CPS Set 1 VMs, you can migrate the Set 2 VMs as described in this section.

Note Diagnostic fails during migration. This is normal since NTP may be converging, mongo replica sets are not synced, and so on. If you see HAProxy diagnostics warnings about Diameter endpoints being down, see HAProxy Diagnostics Warnings, on page 21 for a workaround.

You can disable syncing of carbon database and bulk statistics files to decrease the ISSM time. For more information, refer to Disable Syncing Carbon Database and Bulk Stats Files, on page 20.

Step 1 Run the following command to disable the Set 2 VMs:

```
/mnt/iso/migrate.sh disable set 2
```

After the script has run, you should see information like the following:

```
2017-01-21 01:53:49,894 INFO [__main__.extra_banner]
       ===================================================================
       | Backing up to file: /var/tmp/migrate_set-2_20170121_015349.tar.gz
       ===================================================================

2017-01-21 02:00:12,252 INFO [backup.handleRequest]
=======================================================
2017-01-21 02:00:12,253 INFO [backup.handleRequest] Archive
/var/tmp/migrate/set-2/config_other_br.tar.gz is created with requested backups.
2017-01-21 02:00:12,253 INFO [backup.handleRequest]
=======================================================
2017-01-21 02:00:12,253 INFO [__main__.run_recipe] Performing installation stage: Create backup Tar
```

```
2017-01-21 02:00:12,577 INFO [__main__.<module>] ====================
2017-01-21 02:00:12,578 INFO [__main__.<module>]  SUCCESS
2017-01-21 02:00:12,578 INFO [__main__.<module>] ======== END ========
```

Note   Grafana view (GUI) does not display any information till pcrfclient01 is deleted. As soon as pcrfclient01 is deleted, Grafana GUI display comes up. After recreating pcrfclient01, Grafana view (GUI) does not show till Set 2 VMs (where pcrfclient01 is present) is enabled. Data is not lost, only Grafana view (GUI) is not displayed.

Step 2   (Optional) You can reduce the migration time by provisioning the VMs. If you do not want to provision the VMs, go to .

Note   The VM provisioning requires extra disk space for each VM. Provisioning can be done only for VMware environment setups.

a)   After provisioning Set 1 VMs, you can provision Set 2 VMs by running the following command:

```
/var/qps/install/current/scripts/deployer/support/deploy_all.py --provision --vms
/var/tmp/cluster-upgrade-set-2.txt
```

Note   Manually enter `deploy_all.py` command in your system.

Step 3   Confirm that the Set 2 VMs' sessionmgrs are removed from the replica sets by running the following command:

`diagnostics.sh --get`

Example output is shown below:

```
CPS Diagnostics HA Multi-Node Environment
-------------------------
Checking replica sets...
|---------------------------------------------------------------------------------------------|
| Mongo:3.2.13          MONGODB REPLICA-SETS STATUS INFORMATION      Date : 2017-01-21 03:13:58 |
|---------------------------------------------------------------------------------------------|
|  SET NAME  -  PORT  : IP ADDRESS - REPLICA STATE - HOST NAME  -   HEALTH - LAST SYNC - PRIORITY |
|---------------------------------------------------------------------------------------------|
| ADMIN:set06                                                                                 |
|  Member-1  -  27721 : 172.16.2.100  -  ARBITER  -  arbitervip  - ON-LINE - --------   -   0 |
|  Member-2  -  27721 : 172.16.2.23   -  PRIMARY  - sessionmgr02 - ON-LINE -  No Sync   -   2 |
|---------------------------------------------------------------------------------------------|
| AUDIT:set05                                                                                 |
|  Member-1  -  27017 : 172.16.2.100  -  ARBITER  -  arbitervip  - ON-LINE  - --------   -   0 |
|  Member-2  -  27017 : 172.16.2.23   -  PRIMARY  - sessionmgr02 - ON-LINE  -  No Sync   -   2 |
|---------------------------------------------------------------------------------------------|
| BALANCE:set02                                                                               |
|  Member-1  -  27718 : 172.16.2.100  -  ARBITER  -  arbitervip   - ON-LINE - --------   -   0 |
|  Member-2  -  27718 : 172.16.2.23   -  PRIMARY  - sessionmgr02 - ON-LINE  -  No Sync   -   2 |
|---------------------------------------------------------------------------------------------|
| REPORTING:set03                                                                             |
|  Member-1  -  27719 : 172.16.2.100  -  ARBITER  -  arbitervip   - ON-LINE - --------   -   0 |
|  Member-2  -  27719 : 172.16.2.23   -  PRIMARY  - sessionmgr02 - ON-LINE  -  No Sync   -   2 |
|---------------------------------------------------------------------------------------------|
| SESSION:set01                                                                               |
|  Member-1  -  27717 : 172.16.2.100  -  ARBITER  -  arbitervip   - ON-LINE - --------   -   0 |
|  Member-2  -  27717 : 172.16.2.23   -  PRIMARY  - sessionmgr02 - ON-LINE  -  No Sync   -   2 |
|---------------------------------------------------------------------------------------------|
| SPR:set04                                                                                   |
|  Member-1  -  27720 : 172.16.2.100  -  ARBITER  -  arbitervip   - ON-LINE  -  --------  -   0 |
|  Member-2  -  27720 : 172.16.2.23   -  PRIMARY  - sessionmgr02 - ON-LINE  -  No Sync   -   2 |
|---------------------------------------------------------------------------------------------|
```

Step 4 If you have provisioned VMs using Step 2, on page 13, you can restart VM using provisioned vmdk image by running the following command and then go to Step 6, on page 15:

```
/var/qps/install/current/scripts/deployer/support/deploy_all.py --useprovision --vms
/var/tmp/cluster-upgrade-set-2.txt
```

Note If you have not provisioned the VMs, go to Step 5, on page 14.

Step 5 Re-deploy the Set 2 VMs.

Note Delete Set 2 VMs before redeploying them with the new base.vmdk.

Note To install the VMs using shared or single storage, you must use
`/var/qps/install/current/scripts/deployer/deploy.sh` *$host* command.

For more information, refer to *Manual Deployment* section in *CPS Installation Guide for VMware*.

For VMware: `/var/qps/install/current/scripts/deployer/support/deploy_all.py --vms`
`/var/tmp/cluster-upgrade-set-2.txt`

Note Manually enter `deploy_all.py` command in your system.

For OpenStack: Use nova boot commands or Heat templates. For more information, refer to *CPS Installation Guide for OpenStack*.

Example Deploying Set 2 with Openstack using nova boot command: The commands given below are for reference purpose only. The user must type the commands manually.

```
nova boot --config-drive true --user-data=pcrfclient01-cloud.cfg --image "new_base_vm" --flavor
"pcrfclient01" --nic net-id="c3df93c2-c2bb-4143-8bb6-b4ec6df65e53,v4-fixed-ip=172.16.2.20" --nic
net-id="4759babe-491a-4c1a-a028-ec4daefa1662,v4-fixed-ip=172.18.11.152" --block-device-mapping
"/dev/vdb=ef2ec05b-c5b2-4ffe-92cb-2e7c60b6ed9e:::0" "pcrfclient01" --availability-zone
az-1:os8-compute-1.cisco.com

nova boot --config-drive true --user-data=sessionmgr01-cloud.cfg --image "new_base_vm" --flavor "sm"
  --nic net-id="c3df93c2-c2bb-4143-8bb6-b4ec6df65e53,v4-fixed-ip=172.16.2.22" --nic
net-id="4759babe-491a-4c1a-a028-ec4daefa1662,v4-fixed-ip=172.18.11.157" --block-device-mapping
"/dev/vdb=04eaed49-2459-44eb-9a8b-011a6b4401aa:::0" "sessionmgr01" --availability-zone
az-1:os8-compute-1.cisco.com

nova boot --config-drive true --user-data=lb01-cloud.cfg --image "new_base_vm" --flavor "lb01" --nic
 net-id="c3df93c2-c2bb-4143-8bb6-b4ec6df65e53,v4-fixed-ip=172.16.2.201" --nic
net-id="4759babe-491a-4c1a-a028-ec4daefa1662,v4-fixed-ip=172.18.11.154" --nic
net-id="392b72f6-b8f1-47b2-ae5f-e529f69866bc,v4-fixed-ip=192.168.2.201" "lb01" --availability-zone
az-1:os8-compute-1.cisco.com

nova boot --config-drive true --user-data=qns01-cloud.cfg --image "new_base_vm" --flavor "qps" --nic
 net-id="c3df93c2-c2bb-4143-8bb6-b4ec6df65e53,v4-fixed-ip=172.16.2.24" "qns01" --availability-zone
az-1:os8-compute-1.cisco.com

nova boot --config-drive true --user-data=qns03-cloud.cfg --image "new_base_vm" --flavor "qps" --nic
 net-id="c3df93c2-c2bb-4143-8bb6-b4ec6df65e53,v4-fixed-ip=172.16.2.26" "qns03" --availability-zone
az-1:os8-compute-1.cisco.com
```

Important After deployment of load balancer VM, verify monit service status by executing the following command on deployed Load Balancer (lb) VM:

```
/bin/systemctl status monit.service
```

If monit service on load balancer VM is not running, then execute the following command on that VM to start it:

```
/bin/systemctl start monit.service
```

Step 6 If you are using OpenStack, assign:

- arbitervip to pcrfclient01 internal IP

- lbvip01 to lb01 management IP

- lbvip02 to lb01 internal IP

- Gx VIP to lb01 Gx IP

Example Assigning VIPs to Set 2 VMs using neutron port command: The commands given below are for reference purpose only. The user must type the commands manually.

```
[root@os8-control cloud(keystone_core)]# neutron port-list | grep "172.16.2.20"
| 19678c2d-3efc-4523-ac0b-dd25734e241a |       | fa:16:3e:a5:7c:16 | {"subnet_id":
"106db79e-da5a-41ea-a654-cffbc6928a56", "ip_address": "172.16.2.20"}    |
[root@os8-control cloud(keystone_core)]#  neutron port-update 19678c2d-3efc-4523-ac0b-dd25734e241a
--allowed-address-pairs type=dict list=true ip_address=172.16.2.100
Updated port: 19678c2d-3efc-4523-ac0b-dd25734e241a
[root@os8-control cloud(keystone_core)]#  neutron port-list | grep "172.16.2.201"
| ac12d0ae-4de6-4d15-b5de-b0140d895be8 |       | fa:16:3e:99:e3:7b | {"subnet_id":
"106db79e-da5a-41ea-a654-cffbc6928a56", "ip_address": "172.16.2.201"}    |
[root@os8-control cloud(keystone_core)]# neutron port-update ac12d0ae-4de6-4d15-b5de-b0140d895be8
--allowed-address-pairs type=dict list=true ip_address=172.16.2.200
Updated port: ac12d0ae-4de6-4d15-b5de-b0140d895be8
[root@os8-control cloud(keystone_core)]# neutron port-list | grep "172.18.11.154"
| adab87ae-6d00-4ba0-a139-a9522c881a07 |       | fa:16:3e:8a:d4:47 | {"subnet_id":
"641276aa-245f-46db-b326-d5017915ccf7", "ip_address": "172.18.11.154"} |
[root@os8-control cloud(keystone_core)]# neutron port-update adab87ae-6d00-4ba0-a139-a9522c881a07
--allowed-address-pairs type=dict list=true ip_address=172.18.11.156
Updated port: adab87ae-6d00-4ba0-a139-a9522c881a07
[root@os8-control cloud(keystone_core)]#  neutron port-list | grep "192.168.2.201"
| 2e0f0573-7f6f-4c06-aee1-e81608e84042 |       | fa:16:3e:c2:28:6b | {"subnet_id":
"263ba6d1-31b0-450a-9a2d-30418f3476f9", "ip_address": "192.168.2.201"} |
[root@os8-control cloud(keystone_core)]# neutron port-update 2e0f0573-7f6f-4c06-aee1-e81608e84042
--allowed-address-pairs type=dict list=true ip_address=192.168.2.200
Updated port: 2e0f0573-7f6f-4c06-aee1-e81608e84042
```

For more information, refer to *CPS Installation Guide for OpenStack*.

Important The VMs are rebooted in rescue mode for the first time for CentOS to adjust disk/hardware to the new version. Subsequent reboots if necessary is a normal operation.

Step 7 Run the following command to enable Set 2 VMs:

migrate enable set 2 /var/tmp/migrate-set-2-<*timestamp*>.tar.gz

For example:

```
/mnt/iso/migrate.sh enable set 2 /var/tmp/migrate_set-2_20170120_212456.tar.gz
```

Note The migration does not restore users created with `adduser.sh` due to potential gid/uid conflicts. Check the migrate enable log for entries that indicate users that are not being migrated, and then manually recreate them using `addusers.sh`. An example log is shown below:

```
2017-02-13 14:52:15,999 INFO [etc_passwd.parse_etc_passwd] Parsing
/var/tmp/migrate/pcrfclient02/etc/passwd file
2017-02-13 14:52:16,000 INFO [etc_group.parse_etc_group] Parsing
/var/tmp/migrate/pcrfclient02/etc/group file
2017-02-13 14:52:16,000 WARNING [restore_vm.restore_vms] On Host:pcrfclient02 User/Group
mongoreadonly/mongoreadonly is not being migrated and must be manually created using
adduser.sh.
2017-02-13 14:52:16,000 WARNING [restore_vm.restore_vms] On Host:pcrfclient02 User/Group
admin/admin is not being migrated and must be manually created using adduser.sh.
```

After the script has run, you should see information like the following:

```
INFO      Priority set operation is completed for SPR-SET1
INFO      Priority set to the Database members is finished
INFO      Validating if Priority is set correctly for Replica-Set: SPR-SET1
INFO      Validated Priority is set correctly for Replica-Set: SPR-SET1

2017-01-23 20:46:01,621 INFO [__main__.<module>] =====================
2017-01-23 20:46:01,621 INFO [__main__.<module>]  SUCCESS
2017-01-23 20:46:01,621 INFO [__main__.<module>] ======== END ========
```

Step 8   Run diagnostics to verify that the replica set has all of the members back with the correct priorities.

Step 9   Restore the traffic by running the following command:

`/mnt/iso/migrate.sh traffic restore`

After the script has run, you should see information like the following:

```
2017-01-23 20:54:21,083 INFO [command.execute] (stdout): Stopping haproxy: [  OK  ]
Stopping haproxy: [  OK  ]

2017-01-23 20:54:21,083 INFO [__main__.<module>] =====================
2017-01-23 20:54:21,083 INFO [__main__.<module>]  SUCCESS
2017-01-23 20:54:21,083 INFO [__main__.<module>] ======== END ========
```

What to do next

Note As the change in the replica-sets is not complete at the time of restart, sometimes non-functional impacting errors are listed in the logs. Therefore, for each site, run `restartall.sh` from the Cluster Manager to do a rolling restart of all the nodes at the end of the migration process.

If some of the replica-set members are in RECOVERING state, refer to Recover Replica-set Members from RECOVERING State, on page 17.

Recover Replica-set Members from RECOVERING State

If the migration is performed with live traffic on CPS, there is a possibility that after the migration replica-set members (for huge size databases members like, BALANCE, SPR, REPORTING and so on) can go into RECOVERING state. This is due to the oplog (operation log) size configured which holds the database operation on PRIMARY which might get rolled-over.

To recover the replica-set members from RECOVERY state, you need to perform the steps described in this section:

Execute `rs.printReplicationInfo()` command on PRIMARY database for replica-set whose members went into RECOVERING state to get the configured oplog size and log length start to end information:

```
mongo sessionmgr01:27718
set02:PRIMARY> rs.printReplicationInfo()
configured oplog size:   5120MB
log length start to end: 600secs (0.16hrs)
oplog first event time:  Fri Feb 24 2017 19:51:25 GMT+0530 (IST)
oplog last event time:   Mon Feb 27 2017 22:14:17 GMT+0530 (IST)
now:                     Mon Feb 27 2017 22:14:25 GMT+0530 (IST)
set02:PRIMARY>
```

`rs.printSlaveReplicationInfo` shows the replication lag time (how much secondary is behind the primary member). If you see that this lag is increasing and not catching-up with primary, then this indicates that oplog is getting rolled-over.

```
mongo sessionmgr01:27718
set02:PRIMARY> rs.printSlaveReplicationInfo()
source: sessionmgr02:27718
syncedTo: Mon Feb 27 2017 22:13:17 GMT+0530 (IST)

10 secs (0 hrs) behind the primary
```

What to do next

If the migrated members are still stuck in RECOVERING state, then:

1. Stop the process manually.

2. Refer to *Recovery using Remove/Add members Option* section in *CPS Troubleshooting Guide* to remove failed member and add the member back.

Geographic Redundant Deployment Migration

This section describes the process for performing a migration in a Geographic Redundant deployment. The following example is a Geo replica case involving a replica set containing five members: two members on site 1, two members on site 2, and one arbiter member on site 3 (migration from CPS 13.1.0 to CPS 18.2.0). Each step shows the Mongo version and the CentOS version on the VM; for example, 3.2.10/7.4.

> **Note** In the following table:
>
> - SM = Session Manager
> - S1 = Site 1
> - S2 = Site 2
> - S3 = Third Site
> - 6.8 = CentOS 6.8
> - 7.4 = CentOS 7.4
> - R131 = CPS Release 13.1.0
> - R182 = CPS Release 18.2.0

Table 1: GR Deployment with Site1, Site 2, and 3rd Site Arbiter

Step	SM02-site1	SM01-site1	SM02-site2	SM01-site2	Arbiter	Description
0	3.2.10/6.5	3.2.10/6.5	3.2.10/6.5	3.2.10/6.5	3.2.10/6.5	S1 - R131 S2 - R131 S3 - R131
1	3.2.13/7.4	3.2.10/6.5	3.2.10/6.5	3.2.10/6.5	3.2.10/6.5	-
2	3.2.13/7.4	3.2.13/7.4	3.2.10/6.5	3.2.10/6.5	3.2.10/6.5	S1 - R182 S2 - R131 S3 - R131
3	3.2.13/7.4	3.2.13/7.4	3.2.13/7.4	3.2.10/6.5	3.2.10/6.5	-
4	3.2.13/7.4	3.2.13/7.4	3.2.13/7.4	3.2.13/7.4	3.2.10/6.5	S1 - R182 S2 - R182 S3 - R131
5	3.2.13/7.4	3.2.13/7.4	3.2.13/7.4	3.2.13/7.4	3.2.13/7.4	S1 - R182 S2 - R182 S3 - R182
6	3.2.13/7.4	3.2.13/7.4	3.2.13/7.4	3.2.13/7.4	3.2.13/7.4	-
7	3.2.13/7.4	3.2.13/7.4	3.2.13/7.4	3.2.13/7.4	3.2.13/7.4	-

Detailed steps are shown below:

Step 0: Before starting the migration.

Step 1: Disable monitoring database script by commenting out configured sets from mon_db configs on pcrfclient01/pcrfclient02/cluman in the following files:

```
/etc/broadhop/mon_db_for_callmodel.conf
/etc/broadhop/mon_db_for_lb_failover.conf
/var/qps/install/current/scripts/build/build_etc.sh
```

Steps 2 and 3: Using HA migrate process, Migrate site 1 VMs to CentOS 7.4.

Steps 4 and 5: Using HA migrate process, Migrate site 2 VMs to CentOS 7.4.

Step 6: Using 3rd site arbiter migrate process, Migrate site 3 (3rd Site arbiter) to CentOS 7.4.

Step 7: Verify all replica set members are running Mongo 3.2.10.

Step 8: Enable monitoring by uncommenting configured sets from mon_db configs on pcrfclient01/pcrfclient02/cluman in the following files:

```
/etc/broadhop/mon_db_for_callmodel.conf
/etc/broadhop/mon_db_for_lb_failover.conf
/var/qps/install/current/scripts/build/build_etc.sh
```

Migrate 3rd Site Arbiter

Migrate Site 1 and Site 2, and then migrate 3rd Site Arbiter.

Step 1   Copy the new CPS 18.2.0 ISO to the existing CPS arbiter.

Step 2   Unmount the old CPS ISO by running the following command:

`unmount /mnt/iso`

Step 3   Mount the new CPS 18.2.0 ISO to the arbiter by running the following command:

`mount -o loop cps-arbiter-x.x.x.iso /mnt/iso`

Step 4   Disable the arbiter by running the following command:

`/mnt/iso/migrate.sh disable arbiter`

This command creates the following file:

/var/tmp/migrate_arbiter_<*date_&_time*>.tar.gz

After the command has run successfully, you should see messages like the following:

```
2017-02-10 07:51:42,633 INFO [command.execute] Mongo port:27719 stopped successfully
2017-02-10 07:51:42,633 INFO [__main__.run_recipe] Performing installation stage:
ExtractInstallArtifacts
2017-02-10 07:51:42,634 INFO [extract_install_artifacts.extract_scripts] Extracting CPS scripts
2017-02-10 07:51:43,506 INFO [__main__.run_recipe] Performing installation stage: PrepareWorkingDir
2017-02-10 07:51:43,506 INFO [__main__.run_recipe] Performing installation stage: BackupArbiter
2017-02-10 07:52:25,715 INFO [__main__.run_recipe] Performing installation stage: Create backup Tar
2017-02-10 07:52:37,921 INFO [__main__.<module>] ====================
2017-02-10 07:52:37,921 INFO [__main__.<module>]  SUCCESS
2017-02-10 07:52:37,921 INFO [__main__.<module>] ======== END ========
```

Step 5   Back up the `tar.gz` file to an external location using commands like the following:

For example:

```
sftp root@172.16.2.39
sftp> get migrate_arbiter_date_time.tar.gz
Fetching /var/tmp/migrate_arbiter_date_time.tar.gzmigrate_arbiter_date_time.tar.gz to
migrate_arbiter_20170210_075135.tar.gz
/var/tmp/migrate_arbiter_date_time.tar.gz
```

In this example, 172.16.2.39 is the internal IP address of the arbiter.

Step 6   Deploy the new arbiter using the CPS 18.2.0 ISO and the new_base_vm as the new deployment. To do this, use the instructions provided in the *CPS Geographic Redundancy Guide* for your operating system.

Step 7   Copy the migrate tar.gz file from the external location to the new CPS 18.2.0 arbiter, and run the following command:

`/mnt/iso/migrate.sh enable arbiter <full_path>/migrate_arbiter_<date_and_time>.tar.gz`

After the migration has run successfully, you should see messages like the following:

```
2017-02-13 15:45:53,187 INFO [command.execute] child process started successfully, parent exiting
2017-02-13 15:45:53,188 INFO [command.execute] ^[[60G[^[[0;32m  OK  ^[[0;39m]
2017-02-13 15:45:53,189 INFO [command.execute]
2017-02-13 15:45:53,189 INFO [__main__.<module>] =====================
2017-02-13 15:45:53,189 INFO [__main__.<module>]   SUCCESS
2017-02-13 15:45:53,189 INFO [__main__.<module>] ======== END ========
```

Step 8   Run **`about.sh`** and verify the time zone and CentOS version on the arbiter. You should see output like the following:

```
about.sh
Cisco Policy Suite - Copyright (c) 2015. All rights reserved.

CPS Arbiter

CPS Installer Version - 18.2.0
```

The timezone on the arbiter should have changed to UTC and the CentOS to version 7.4 as shown below:

```
cat /etc/*elease
CentOS release 7.4 (Final)
CentOS release 7.4 (Final)
CentOS release 7.4 (Final)
[root@site3-arbiter log]# date
Tue Feb 21 02:42:54 UTC 2017
```

Disable Syncing Carbon Database and Bulk Stats Files

To disable syncing of carbon database and bulk statistics files, add the following parameters in /var/install.cfg file:

- SKIP_BLKSTATS

- SKIP_CARBONDB

Example to disable:

```
SKIP_BLKSTATS=1
SKIP_CARBONDB=1
```

HAProxy Diagnostics Warnings

Traffic swapping/restoring is accomplished on a silo basis by turning Diameter endpoints up or down. During migration, there is a chance that endpoints might not recover. If this happens, HAProxy diagnostics warnings indicate that Diameter endpoints are down. This section provides a workaround for enabling the endpoints manually if these errors occur.

Step 1   Display any HAProxy diagnostics warnings by running the following command:

`diagnostics.sh --ha_proxy`

Example warnings that Diameter endpoints are down are shown below:

```
Checking HAProxy statistics and ports...
[WARN]
  HA Proxy Diameter is displaying some services as down or with errors.  If services are restarting,
 this is normal.
  Please wait up to a minute after restart is successful to ensure services are marked up.
  Services marked DOWN 1/2 are coming up (1 success in last 2 tries).  Services marked UP 1/3 are
going down.
  Go to the following url for complete HA Proxy status information:
http://lbvip01:5540/haproxy-diam?stats
---------------------------------
diameter-int1-vip-lb02-A DOWN L4CON
 Sessions (current,max,limit): 0,1, Rate (sessions,max,limit): 0,1, Last Status change (seconds):
63027
diameter-int1-vip-lb02-B DOWN L4CON
 Sessions (current,max,limit): 0,1, Rate (sessions,max,limit): 0,1, Last Status change (seconds):
63025
diameter-int1-vip-lb02-C DOWN L4CON
 Sessions (current,max,limit): 0,1, Rate (sessions,max,limit): 0,1, Last Status change (seconds):
63024
---------------------------------
```

In each load balancer, there are four java processes running (iomgr, diameter_endpoint_1, diameter_endpoint_2, diameter_endpoint_3). Each one of the diameter endpoints have a different OSGI port (9092, 9093, 9094).

Step 2   To disable endpoints, you need to run commands like those shown in the example series below.

Make sure you choose the proper load balancer node. If this is being done to enable the diameter endpoint for Set-1, then use lb02. If it is being done for Set-2, then use lb01. The example is for set-2, and thus uses lb01.

1. Log in to the OSGi console and run the excludeEndpoints command as shown in the following example:

```
telnet localhost 9092
Trying 127.0.0.1...
Connected to localhost.
Escape character is '^]'.

osgi> excludeEndpoints
osgi>
```

2. Enable the endpoints by running the following command:

`clearExcludedEndpoints`

3. Leave the OSGi console (without killing the process) by running the disconnect command as shown below:

```
osgi> disconnect
Disconnect from console? (y/n; default=y) y
Connection closed by foreign host.
```

Step 3 After you run clearExcludedEndpoints, it will take a minute and then HAProxy will pick it up. If it does not, then restart the processes as follows:

```
monit restart qns-1
monit restart qns-2
monit restart qns-3
monit restart qns-4
```

Migration Rollback

The following steps describe the process to restore a CPS cluster to the previous version when it is determined that an In Service Software Migration is not progressing correctly or needs to be abandoned after evaluation of the new version.

Migration rollback using the following steps can only be performed after Migration Set 1 is completed. These migration rollback steps cannot be used if the entire CPS cluster has been migrated.

Rollback Considerations

- The automated rollback process can only restore the original software version.

- You must have a valid Cluster Manager VM backup (snapshot/clone) which you took prior to starting the migration.

- The migration rollback should be performed during a maintenance window. During the rollback process, call viability is considered on a best effort basis.

- Rollback is only supported for deployments where Mongo database configurations are stored in `mongoConfig.cfg` file. Alternate methods used to configure Mongo will not be backed up or restored.

- Rollback is not supported with a `mongoConfig.cfg` file that has sharding configured.

- For replica sets, a rollback does not guarantee that the primary member of the replica set will remain the same after a rollback is complete. For example, if sessionmgr02 starts off as the primary, then a migration will demote sessionmgr02 to secondary while it performs an upgrade. If the upgrade fails, sessionmgr02 may remain in secondary state. During the rollback, no attempt is made to reconfigure the primary, so sessionmgr02 will remain secondary. In this case, you must manually reconfigure the primary after the rollback, if desired.

Roll Back the Migration

The following steps describe how to roll back the migration for Set 1 VMs.

Before you begin

- Check for call traffic.

- Make sure that you have the run `migrate.sh enable set 1` command. The rollback will work only after that command has been run.

- Run `diagnostics.sh --get_replica_status` to check which new Set 1 sessionmgr*XX* (even numbered) mongo processes are in RECOVERING state. If so, manually stop all those processes on respective session managers.

Example:

```
|-------------------------------------------------------------------------------------|
| Mongo:3.2.10        MONGODB REPLICA-SETS STATUS INFORMATION Date : 2017-02-24 10:44:20|
|-------------------------------------------------------------------------------------|
|  SET NAME - PORT  : IP ADDRESS - REPLICA STATE - HOSTNAME -HEALTH - LASTSYNC - PRIORITY|
|-------------------------------------------------------------------------------------|
| ADMIN:set06
   |
| Member-1 -  27721 : 172.20.35.25 - PRIMARY   - sessionmgr01 - ON-LINE - No Primary -
 3 |
| Member-2 -  27721 : 172.20.35.34 - ARBITER   - arbitervip   - ON-LINE - --------   -
 1 |
| Member-3 -  27721 : 172.20.35.26 - SECONDARY - sessionmgr02 - ON-LINE -   12 min   -
 1 |
|-------------------------------------------------------------------------------------|
| BALANCE:set02
        |
| Member-1 -  27718 : 172.20.35.34 - ARBITER   - arbitervip   - ON-LINE - --------  -
 1 |
| Member-2 -  27718 : 172.20.35.25 - PRIMARY   - sessionmgr01 - ON-LINE - --------  -
 3 |
| Member-3 -  27718 : 172.20.35.26 - RECOVERING- sessionmgr02 - ON-LINE - 12 min    -
 1 |
|-------------------------------------------------------------------------------------|
```

As you can see in the example, sessionmgr02 from `Balance: set02` is in RECOVERING state, you need to manually stop process for 27718.

`/usr/bin/systemctl stop sessionmgr-27718`

☞

Important   Make sure the process has been stopped properly by running the command: `ps -ef|grep 27718`. If it has not stopped, then manually kill the process.

Step 1   Start the rollback of Set 1 VMs by running the following command:

/mnt/iso/migrate.sh rollback

After the script has run, you should see information like the following:

```
2017-02-03 20:10:30,653 INFO [fabric_tasks.run] Stopping all services on remote VM qns04
2017-02-03 20:10:30,654 INFO [transport._log] Secsh channel 3 opened.
2017-02-03 20:10:40,745 INFO [transport._log] Secsh channel 4 opened.
2017-02-03 20:10:42,111 INFO [__main__.<module>] ====================
2017-02-03 20:10:42,111 INFO [__main__.<module>]  SUCCESS
```

```
2017-02-03 20:10:42,111 INFO [__main__.<module>] ======== END ========
```

Step 2 Save the Set 1 backup tar file (`migrate_set-1*.tar.gz`) to an external location.

> **Note** This file was created by the `migrate disable set 1` command that was run when Set 1 VMs were disabled.

Step 3 Restore the older Cluster Manager VM (for example, CPS 11) from the backup (snapshot/clone).

Step 4 Create cluster sets for migration rollback by running the following command:

`/var/platform/platform/scripts/create-cluster-sets.sh`

You should see the following output:

```
Created /var/tmp/cluster-upgrade-set-1.txt
Created /var/tmp/cluster-upgrade-set-2.txt
```

Step 5 Delete and redeploy Set 1 VMs on the original CPS/basevm (CntOS 6.5).

For VMware, run the following command to redeploy the Set 1 VMs:

`/var/qps/install/current/scripts/deployer/deploy.sh $host`

where, *$host* is the short alias name and not the full host name.

For example:

```
./deploy.sh qns02
```

Step 6 If you are using OpenStack, assign arbitervip, lbvip01, lbvip02 and gx vip to pcrfclient02 internal ip, lb02 management ip, lb02 internal ip, and lb02 gx ip respectively.

Step 7 Copy the migrate_set-1_* file from the external location to the Cluster Manager VM.

Step 8 Mount the *CPS_*.*release.iso* to the existing CPS Cluster Manager by running the following command:

```
mount -o loop CPS_*.release.iso /mnt/iso
```

where, * is the release number to which you have migrated.

Step 9 Run the following command to enable Set 1 VMs. For example:

`/mnt/iso/migrate enable set 1 /root/migrate-set-1-<timestamp>.tar.gz file`

For example:

```
/mnt/iso/migrate.sh enable set 1 /root/migrate_set-1_20170120_212456.tar.gz
```

After the script has run, you should see information like the following:

```
WARNING  Mongo Server trying to reconnect while pushing config. Attempt #1
INFO     Priority set operation is completed for SPR-SET1
INFO     Priority set to the Database members is finished
INFO     Validating if Priority is set correctly for Replica-Set: SPR-SET1
INFO     Validated Priority is set correctly for Replica-Set: SPR-SET1

2017-01-21 02:45:48,950 INFO [__main__.<module>] ====================
2017-01-21 02:45:48,950 INFO [__main__.<module>]   SUCCESS
2017-01-21 02:45:48,951 INFO [__main__.<module>] ======== END ========
```

Note   Corosync may disable the admin arbiter (mongod) on the active arbitervip. If so, re-run `/mnt/iso/migrate.sh enable set 1`.

What to do next

If after rollback is completed and few members are still stuck in RECOVERY state, then:

1. Stop the process manually.

2. Refer to the *Recovery using Remove/Add members Option* section in the *CPS Troubleshooting Guide* to remove failed member and add the member back.

Remove ISO Image

Step 1   (Optional) After the migration is complete, unmount the ISO image from the Cluster Manager VM. This prevents any "device is busy" errors when a subsequent upgrade is performed.

```
cd /root

umount /mnt/iso
```

Step 2   (Optional) After unmounting the ISO, delete the ISO image that you loaded on the Cluster Manager to free the system space.

```
rm -rf /<path>/CPS_x.x.x.release.iso
```

Remove ISO Image