CPS Migration and Upgrade Guide, Release 18.3.0 (Restricted Release)

First Published: 2018-07-20

Last Modified: 2018-07-20

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

CONTENTS

Preface

About this Guide

This guide describes the migration and upgrade procedures for the Cisco Policy Suite (CPS) system.

Refer to the *CPS Installation Guide for VMware* for instructions to install a new CPS deployment in a VMware environment, or to the *CPS Installation Guide for OpenStack* to install a new CPS deployment in an OpenStack environment.

Audience

This guide is best used by these readers:

- Network administrators
- Network engineers
- Network operators
- System administrators

This document assumes a general understanding of network architecture, configuration, and operations.

Additional Support

For further documentation and support:

- Contact your Cisco Systems, Inc. technical representative.
- Call the Cisco Systems, Inc. technical support number.

- Write to Cisco Systems, Inc. at support@cisco.com.
- Refer to support matrix at https://www.cisco.com/c/en/us/support/index.html and to other documents related to Cisco Policy Suite.

Conventions (all documentation)

This document uses the following conventions.

Conventions	Indication
bold font	Commands and keywords and user-entered text appear in **bold** font.
italic font	Document titles, new or emphasized terms, and arguments for which you supply values are in *italic* font.
[ ]	Elements in square brackets are optional.
{x | y | z }	Required alternative keywords are grouped in braces and separated by vertical bars.
[x | y | z]	Optional alternative keywords are grouped in brackets and separated by vertical bars.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.
courier font	Terminal sessions and information the system displays appear in courier font.
< >	Nonprinting characters such as passwords are in angle brackets.
[ ]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

Note Means reader take note. Notes contain helpful suggestions or references to material not covered in the manual.

Caution Means reader be careful. In this situation, you might perform an action that could result in equipment damage or loss of data.

Warning IMPORTANT SAFETY INSTRUCTIONS.

Means danger. You are in a situation that could cause bodily injury. Before you work on any equipment, be aware of the hazards involved with electrical circuitry and be familiar with standard practices for preventing accidents. Use the statement number provided at the end of each warning to locate its translation in the translated safety warnings that accompanied this device.

SAVE THESE INSTRUCTIONS

Note Regulatory: Provided for additional information and to comply with regulatory and customer requirements.

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, using the Cisco Bug Search Tool (BST), submitting a service request, and gathering additional information, see What's New in Cisco Product Documentation.

To receive new and revised Cisco technical content directly to your desktop, you can subscribe to the . RSS feeds are a free service.

RESTRICTED RELEASE

Important This is a Short Term Support (STS) release with availability and use restrictions. Contact your Cisco Account or Support representatives for more information.

RESTRICTED RELEASE

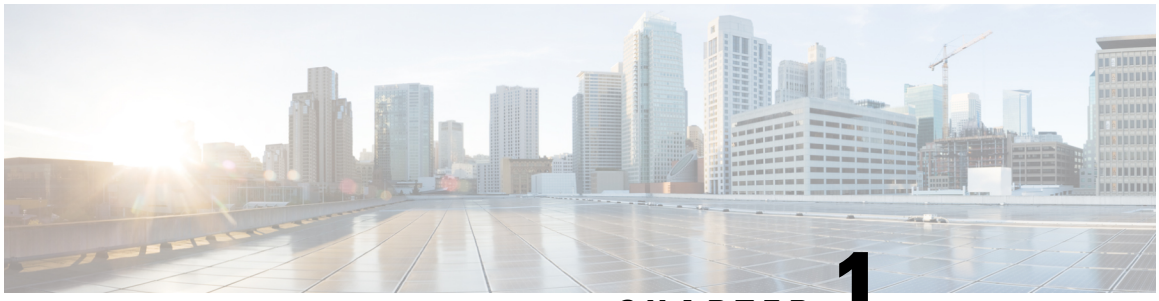

CHAPTER 1

Migrate CPS

Migration is supported from CPS 14.0.0, CPS 18.0.0 and CPS 18.1.0 to CPS 18.3.0.

In-Service Migration to 18.3.0

This section describes the steps to perform an in-service software migration of a CPS 14.0.0, CPS 18.0.0 and CPS 18.1.0 to CPS 18.3.0. This migration allows the traffic to continue running while the migration is being performed.

In-service software migrations to CPS 18.3.0 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 Check the System Health, on page 3.
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 Rollback Considerations, on page 22 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 18.3.0 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 18.3.0 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 14.0.0/18.0.0/18.1.0 Cluster Manager

Before migrating Cluster Manager to CPS 18.3.0, 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 After removing the files, reboot the Cluster Manager.

Step 3 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 18.3.0.

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*.

The following logback files are overwritten with latest files after ISSM. Any modification done to these files, needs to merge mannually after migration is complete:

```
/etc/broadhop/logback-debug.xml
/etc/broadhop/logback-netcut.xml
/etc/broadhop/logback-pb.xml
/etc/broadhop/logback.xml
```

```
/etc/broadhop/controlcenter/logback.xml
```

Backup of old `logback.xml` files is available at `/var/tmp/logback_backup` on newly deployed Cluster Manager VM after running `restore_cluman.py` script. Same files are also available in `migrate_cluman_*.tar.gz` generated in Step 4, on page 5.

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

```
umount /mnt/iso
```

Step 2 Mount the new CPS 18.3.0 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 18.3.0 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 18.3.0 as new deployment. Deploy the CPS 18.3.0 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 18.3.0 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 - 18.3.0
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 18.3.0, 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:

```
2018-06-21 01:53:49,894 INFO [__main__.extra_banner]
        ======================================================================
        | Backing up to file: /var/tmp/migrate_set-1_20180621_212456.tar.gz
        ======================================================================

2018-06-21 02:00:12,252 INFO [backup.handleRequest]
=====================================================
2018-06-21 02:00:12,253 INFO [backup.handleRequest] Archive
/var/tmp/migrate/set-1/config_other_br.tar.gz is created with requested backups.
2018-06-21 02:00:12,253 INFO [backup.handleRequest]
=====================================================
2018-06-21 02:00:12,253 INFO [__main__.run_recipe] Performing installation stage: Create backup Tar
2018-06-21 02:00:12,577 INFO [__main__.<module>] ======================
2018-06-21 02:00:12,578 INFO [__main__.<module>]  SUCCESS
2018-06-21 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:3.2.19          MONGODB REPLICA-SETS STATUS INFORMATION        Date : 2018-06-21 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, if you are using static route, copy static route files (i.e. route-ifname) to the VMs where they are configured and restart network service on these VMs.

Step 9 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_20180621_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:

```
2018-06-21 14:52:15,999 INFO [etc_passwd.parse_etc_passwd] Parsing
/var/tmp/migrate/pcrfclient02/etc/passwd file
2018-06-21 14:52:16,000 INFO [etc_group.parse_etc_group] Parsing
/var/tmp/migrate/pcrfclient02/etc/group file
2018-06-21 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.
2018-06-21 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

2018-06-21 02:45:48,950 INFO [__main__.<module>] ======================
2018-06-21 02:45:48,950 INFO [__main__.<module>]  SUCCESS
2018-06-21 02:45:48,951 INFO [__main__.<module>] ======== END ========
```

Step 10 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): 18.3.0.release
         lb02: qns-3   (diameter_endpoint): 18.3.0.release
         lb02: qns-4   (diameter_endpoint): 18.3.0.release
        qns01: qns-1                (pcrf): 18.0.0.release
```

```
        qns02: qns-1                (pcrf): 18.3.0.release
        qns03: qns-1                (pcrf): 18.0.0.release
        qns04: qns-1                (pcrf): 18.3.0.release
pcrfclient01: qns-1        (controlcenter): 18.0.0.release
pcrfclient01: qns-2                  (pb): 18.0.0.release
pcrfclient02: qns-1        (controlcenter): 18.3.0.release
pcrfclient02: qns-2                  (pb): 18.3.0.release
```

Step 11 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

2018-06-21 19:56:09,092 INFO [__main__.<module>] ======================
2018-06-21 19:56:09,092 INFO [__main__.<module>]  SUCCESS
2018-06-21 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:

```
2018-06-21 01:53:49,894 INFO [__main__.extra_banner]
        ========================================================================
        | Backing up to file: /var/tmp/migrate_set-2_20180621_015349.tar.gz
```

```
        ============================================================

2018-06-21 02:00:12,252 INFO [backup.handleRequest]
=====================================================
2018-06-21 02:00:12,253 INFO [backup.handleRequest] Archive
/var/tmp/migrate/set-2/config_other_br.tar.gz is created with requested backups.
2018-06-21 02:00:12,253 INFO [backup.handleRequest]
=====================================================
2018-06-21 02:00:12,253 INFO [__main__.run_recipe] Performing installation stage: Create backup Tar
2018-06-21 02:00:12,577 INFO [__main__.<module>] ======================
2018-06-21 02:00:12,578 INFO [__main__.<module>]  SUCCESS
2018-06-21 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 Step 3, on page 13.

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:3.2.19          MONGODB REPLICA-SETS STATUS INFORMATION        Date : 2018-06-21 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_20180621_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:

```
2018-06-21 14:52:15,999 INFO [etc_passwd.parse_etc_passwd] Parsing
/var/tmp/migrate/pcrfclient02/etc/passwd file
2018-06-21 14:52:16,000 INFO [etc_group.parse_etc_group] Parsing
/var/tmp/migrate/pcrfclient02/etc/group file
2018-06-21 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.
2018-06-21 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

2018-06-21 20:46:01,621 INFO [__main__.<module>] ======================
2018-06-21 20:46:01,621 INFO [__main__.<module>]  SUCCESS
2018-06-21 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:

```
2018-06-21 20:54:21,083 INFO [command.execute] (stdout): Stopping haproxy: [  OK  ]
Stopping haproxy: [  OK  ]

2018-06-21 20:54:21,083 INFO [__main__.<module>] ======================
2018-06-21 20:54:21,083 INFO [__main__.<module>]  SUCCESS
2018-06-21 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 14.0.0 to CPS 18.3.0). Each step shows the Mongo version and the CentOS version on the VM; for example, 3.2.13/7.4.

Note In the following table:

- SM = Session Manager
- S1 = Site 1
- S2 = Site 2
- S3 = Third Site
- 7.4 = CentOS 7.4
- R140 = CPS Release 14.0.0
- R183 = CPS Release 18.3.0

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

Step	SM02-site1	SM01-site1	SM02-site2	SM01-site2	Arbiter	Description
0	3.2.13/7.4	3.2.13/7.4	3.2.13/7.4	3.2.13/7.4	3.2.13/7.4	S1 - R140 S2 - R140 S3 - R140
1	3.2.19/7.4	3.2.13/7.4	3.2.13/7.4	3.2.13/7.4	3.2.13/7.4	-
2	3.2.19/7.4	3.2.19/7.4	3.2.13/7.4	3.2.13/7.4	3.2.13/7.4	S1 - R183 S2 - R140 S3 - R140
3	3.2.19/7.4	3.2.19/7.4	3.2.19/7.4	3.2.13/7.4	3.2.13/7.4	-
4	3.2.19/7.4	3.2.19/7.4	3.2.19/7.4	3.2.19/7.4	3.2.13/7.4	S1 - R183 S2 - R183 S3 - R140
5	3.2.19/7.4	3.2.19/7.4	3.2.19/7.4	3.2.19/7.4	3.2.19/7.4	S1 - R183 S2 - R183 S3 - R183
6	3.2.19/7.4	3.2.19/7.4	3.2.19/7.4	3.2.19/7.4	3.2.19/7.4	-
7	3.2.19/7.4	3.2.19/7.4	3.2.19/7.4	3.2.19/7.4	3.2.19/7.4	-

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 CPS 18.3.0.

Steps 4 and 5: Using HA migrate process, Migrate site 2 VMs to CPS 18.3.0.

Step 6: Using 3rd site arbiter migrate process, Migrate site 3 (3rd Site arbiter) to CPS 18.3.0.

Step 7: Verify all replica set members are running CPS 18.3.0 and Mongo 3.2.19.

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 18.3.0 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 18.3.0 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 18.3.0 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 18.3.0 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 - 18.3.0
```

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.

Note If a member is shown in an unknown state, it is likely that the member is not accessible from one of other members, mostly an arbiter. In that case, you must go to that member and check its connectivity with other members.

Also, you can login to mongo on that member and check its actual status.

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.

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
```

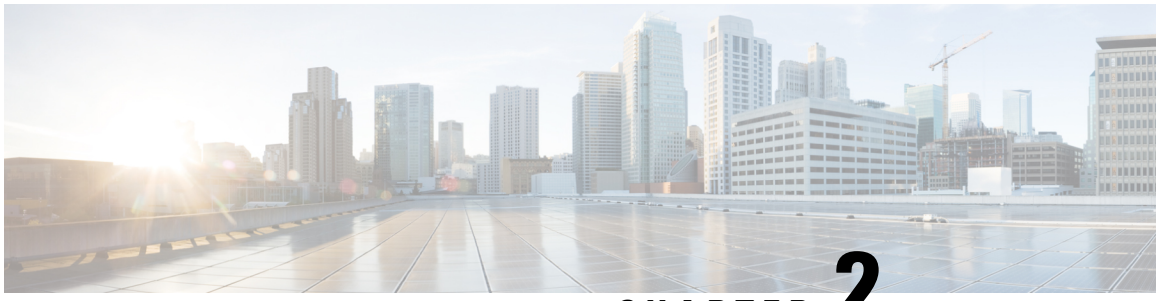

CHAPTER 2

Upgrade CPS

Refer to the *CPS Installation Guide for VMware* for instructions to install a new CPS deployment in a VMware environment, or the *CPS Installation Guide for OpenStack* to install a new CPS deployment in an OpenStack environment.

In-Service Software Upgrade to 18.3.0

This section describes the steps to perform an in-service software upgrade (ISSU) of an existing CPS 18.2.0 deployment to CPS 18.3.0. This upgrade allows traffic to continue running while the upgrade is being performed.

In-service software upgrade to 18.3.0 is supported **only** from CPS 18.2.0.

In-service software upgrade to 18.3.0 is supported **only** for Mobile installation. Other CPS installation types cannot be upgraded using ISSU.

Note During ISSU from CPS 18.2.0 to CPS 18.3.0, if the following issue is observed then one needs to reboot Cluster Manager and start ISSU again:

```
/dev/mapper/control: open failed: No such device
Failure to communicate with kernel device-mapper driver.
Check that device-mapper is available in the kernel.
Incompatible libdevmapper 1.02.140-RHEL7 (2017-05-03) and kernel driver (unknown version).
Command failed
```

The issue is observed only when the kernel is updated for the first time. In subsequent ISSU, the kernel issue is not observed.

Prerequisites

Important During the upgrade 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 upgrade has been successfully completed and properly validated.

Before beginning the upgrade:

1. Create a backup (snapshot/clone) of the Cluster Manager VM. If errors occur during the upgrade process, this backup is required to successfully roll back the upgrade.
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 directly will not be backed up and must be reapplied manually after the upgrade is complete.
3. Remove any previously installed patches. For more information on patch removal steps, refer to Remove a Patch.
4. If necessary, upgrade the underlying hypervisor before performing the CPS in-service software upgrade. 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 *CPS Installation Guide for VMware* or *CPS Installation Guide for OpenStack* 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 upgrade 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 Check the System Health, on page 29
8. The following files are overwritten with latest files after ISSU. Any modification done to these files, needs to merge manually after the upgrade.

   ```
   /etc/broadhop/logback-debug.xml
   /etc/broadhop/logback-netcut.xml
   /etc/broadhop/logback-pb.xml
   /etc/broadhop/logback.xml
   /etc/broadhop/controlcenter/logback.xml
   ```

   Refer to logback.xml Update, on page 47 for more details.

Refer also to Rollback Considerations, on page 47 for more information about the process to restore a CPS cluster to the previous version if an upgrade is not successful.

Overview

The in-service software upgrade is performed in increments:

1. Download and mount the CPS software on the Cluster Manager VM.
2. Divide CPS VMs in the system into two sets.
3. Start the upgrade (`install.sh`). The upgrade automatically creates a backup archive of the CPS configuration.
4. Manually copy the backup archive (`/var/tmp/issu_backup-<timestamp>.tgz`) to an external location.
5. Perform the upgrade on the first set while the second set remains operational and processes all running traffic. The VMs included in the first set are rebooted during the upgrade. After upgrade is complete, the first set becomes operational.
6. Evaluate the upgraded VMs before proceeding with the upgrade of the second set. If any errors or issues occurred, the upgrade of set 1 can be rolled back. Once you proceed with the upgrade of the second set, there is no automated method to roll back the upgrade.
7. Perform the upgrade on the second set while the first assumes responsibility for all running traffic. The VMs in the second set are rebooted during the upgrade.

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 to Download and Mount the CPS ISO Image.

Download and Mount the CPS ISO Image

Step 1 Download the Full Cisco Policy Suite Installation software package (ISO image) from software.cisco.com. Refer to the 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.

Step 3 Execute the following commands to mount the ISO image:

```
mkdir /mnt/iso

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

cd /mnt/iso
```

Step 4 Continue with Verify VM Database Connectivity, on page 30.

Verify VM Database Connectivity

Verify that the Cluster Manager VM has access to all VM ports. If the firewall in your CPS deployment is enabled, the Cluster Manager can not access the CPS database ports.

To temporarily disable the firewall, run the following command on each of the OAM (pcrfclient) VMs to disable IPTables:

IPv4: `service iptables stop`

IPv6: `service ip6tables stop`

The iptables service restarts the next time the OAM VMs are rebooted.

Create Upgrade Sets

The following steps divide all the VMs in the CPS cluster into two groups (upgrade set 1 and upgrade set 2). These two groups of VMs are upgraded independently in order allow traffic to continue running while the upgrade is being performed.

Step 1 Determine which VMs in your existing deployment should be in upgrade set 1 and upgrade set 2 by running the following command on the Cluster Manager:

```
/mnt/iso/platform/scripts/create-cluster-sets.sh
```

Step 2 This script outputs two files defining the 2 sets:

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

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

Step 3 Create the file backup-db at the location `/var/tmp`. This file contains backup-session-db (hot-standby) set name which is defined in `/etc/broadhop/mongoConfig.cfg` file (for example, SESSION-SETXX).

For example:

```
cat /var/tmp/backup-db

SESSION-SET23
```

Step 4 Review these files to verify that all VMs in the CPS cluster are included. Make any changes to the files as needed.

Step 5 Continue with Move the Policy Director Virtual IP to Upgrade Set 2, on page 31.

Move the Policy Director Virtual IP to Upgrade Set 2

Before beginning the upgrade of the VMs in upgrade set 1, you must transition the Virtual IP (VIP) to the Policy Director (LB) VM in Set 2.

Check which Policy Director VM has the virtual IP (VIP) by connecting to (ssh) to lbvip01 from the Cluster Manager VM. This connects you to the Policy Director VM which has the VIP either lb01 or lb02.

You can also run `ifconfig` on the Policy Director VMs to confirm the VIP assignment.

- If the VIP is already assigned to the Policy Director VM that is to be upgraded later (Set 2), continue with Upgrade Set 1, on page 31.
- If the VIP is assigned to the Policy Director VM that is to be upgraded now (Set 1), issue the following commands from the Cluster Manager VM to force a switchover of the VIP to the other Policy Director:

```
ssh lbvip01
```

```
service corosync stop
```

Continue with Upgrade Set 1, on page 31.

Upgrade Set 1

Important Perform these steps while connected to the Cluster Manager console via the orchestrator. This prevents a possible loss of a terminal connection with the Cluster Manager during the upgrade process.

The steps performed during the upgrade, including all console inputs and messages, are logged to `/var/log/install_console_<date/time>.log`.

Step 1 Run the following command to initiate the installation script:

```
/mnt/iso/install.sh
```

Step 2 When prompted for the install type, enter **mobile**.

```
Please enter install type [mobile|mog|pats|arbiter|andsf|escef]:
```

Note
- In-service software upgrade to CPS 18.3.0 is supported only for **mobile** installations.
- Currently, eSCEF is an EFT product and is for Lab Use Only. This means it is not supported by Cisco TAC and cannot be used in a production network. The features in the EFT are subject to change at the sole discretion of Cisco.

Step 3 When prompted to initialize the environment, enter **y**.

```
Would you like to initialize the environment... [y|n]:
```

Step 4 (Optional) You can skip Step 2, on page 31 and Step 3, on page 31 by configuring the following parameters in `/var/install.cfg` file:

```
INSTALL_TYPE
INITIALIZE_ENVIRONMENT
```

Example:

```
INSTALL_TYPE=mobile
INITIALIZE_ENVIRONMENT=yes
```

Step 5 When prompted for the type of installation, enter **3**.

```
Please select the type of installation to complete:
1) New Deployment
2) Upgrade to different build within same release (eg: 1.0 build 310 to 1.0 build 311)
   or Offline upgrade from one major release to another (eg: 1.0 to 2.0)
3) In-Service Upgrade from one major release to another (eg: 1.0 to 2.0)
```

Step 6 When prompted, open a second terminal session to the Cluster Manager VM and copy the backup archive to an external location. This archive is needed if the upgrade needs to be rolled back.

```
********** Action Required **********
  In a separate terminal, please move the file /var/tmp/issu_backup-<timestamp>.tgz
  to an external location.
  When finished, enter 'c' to continue:
```

After you have copied the backup archive, enter **c** to continue.

Step 7 When prompted to enter the SVN repository to back up the policy files, enter the Policy Builder data repository name.

```
Please pick a Policy Builder config directory to restore for upgrade [configuration]:
```

The default repository name is `configuration`. This step copies the SVN/policy repository from the pcrfclient01 and stores it in the Cluster Manager. After pcrfclient01 is upgraded, these SVN/policy files are restored.

Step 8 (Optional) If prompted for a user, enter `qns-svn`.

Step 9 (Optional) If prompted for the password for `qns-svn`, enter the valid password.

```
Authentication realm: <http://pcrfclient01:80> SVN Repos

Password for 'qns-svn':
```

Step 10 The upgrade proceeds on Set 1 until the following message is displayed:

Note If CPS detects that the kernel upgrade has already occurred, the next prompt you see is in Step 12, on page 33. If this is the case, skip to Step 12, on page 33.

For example:

```
================================================
Upgrading Set /var/tmp/cluster-upgrade-set-1.txt
================================================
Checking if reboot required for below hosts
pcrfclient02 lb02 sessionmgr02 qns02 <-- These VMs may differ on your deployment.
================================================
WARN - Kernel will be upgraded on below hosts from current set of hosts. To take the effect of new
 kernel below hosts will be rebooted.
Upgrading kernel packages on:
  pcrfclient02 lb02 sessionmgr02 qns02
================================================
```

Step 11 Enter **y** to proceed with the kernel upgrade.

Important The kernel upgrade is mandatory. If you enter **n** at the prompt, the upgrade process is aborted.

Step 12 (Optional) The upgrade proceeds until the following message is displayed:

```
All VMs in /var/tmp/cluster-upgrade-set-1.txt are Whisper READY.
Run 'diagnostics.sh --get_replica_status' in another terminal to check DB state.
Please ensure that all DB members are UP and are to correct state i.e. PRIMARY/SECONDARY/ARBITER.
Continue the upgrade for Next Step? [y/n]
```

Step 13 (Optional) Open a second terminal to the Cluster Manager VM and run the following command to check that all DB members are UP and in the correct state:

```
diagnostics.sh --get_replica_status
```

Step 14 (Optional) After confirming the database member state, enter **y** to continue the upgrade.

Step 15 The upgrade proceeds until the following message is displayed:

```
Please ensure that all the VMS from the /var/tmp/cluster-upgrade-set-1.txt have been upgraded and
restarted. Check logs for failures
If the stop/start for any qns process has failed, please manually start the same before continuing
 the upgrade.
Continue the upgrade? [y/n]
```

Note If you do not want to upgrade, enter **n** to rollback the upgrade and close the window.

If you can cancel the upgrade using any other command (for example, control+c) and start rollback the traffic is not recovered.

Step 16 If you have entered **y**, continue with Evaluate Upgrade Set 1, on page 33.

Evaluate Upgrade Set 1

At this point in the in-service software upgrade, the VMs in Upgrade Set 1 have been upgraded and all calls are now directed to the VMs in Set 1.

Before continuing with the upgrade of the remaining VMs in the cluster, check the health of the Upgrade Set 1 VMs. If any of the following conditions exist, the upgrade should be rolled back.

- Errors were reported during the upgrade of Set 1 VMs.
- Calls are not processing correctly on the upgraded VMs.
- `about.sh` does not show the correct software versions for the upgraded VMs (under CPS Core Versions section).

Note `diagnostics.sh` reports errors about haproxy that the Set 2 Policy Director (Load Balancer) diameter ports are down, because calls are now being directed to the Set 1 Policy Director. These errors are expected and can be ignored.

If clock skew is seen with respect to VM or VMs after executing `diagnostics.sh`, you need to synchronize the time of the redeployed VMs.

For example,

```
Checking clock skew for qns01...[FAIL]
  Clock was off from lb01 by 57 seconds.  Please ensure clocks are synced. See:
/var/qps/bin/support/sync_times.sh
```

Synchronize the times of the redeployed VMs by running the following command:

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

For more information on `sync_times.sh`, refer to *CPS Operations Guide*.

Important Once you proceed with the upgrade of Set 2 VMs, there is **no** automated method for rolling back the upgrade.

If any issues are found which require the upgraded Set 1 VMs to be rolled back to the original version, refer to Upgrade Rollback, on page 47.

To continue upgrading the remainder of the CPS cluster (Set 2 VMs), refer to Move the Policy Director Virtual IP to Upgrade Set 1, on page 34.

Move the Policy Director Virtual IP to Upgrade Set 1

Issue the following commands from the Cluster Manager VM to switch the VIP from the Policy Director (LB) on Set 1 to the Policy Director on Set 2:

```
ssh lbvip01
```

```
service corosync stop
```

If the command prompt does not display again after running this command, press **Enter**.

Continue with Upgrade Set 2, on page 34.

Upgrade Set 2

Step 1 In the first terminal, when prompted with the following message, enter **y** after ensuring that all the VMs in Set 1 are upgraded and restarted successfully.

```
Please ensure that all the VMs from the /var/tmp/cluster-upgrade-set-1.txt have been upgraded and
restarted.
Check logs for failures.
If the stop/start for any qns process has failed, please manually start the same before continuing
 the upgrade.
Continue the upgrade? [y/n]
```

Step 2 The upgrade proceeds on Set 2 until the following message is displayed:

Note If CPS detects that the kernel upgrade has already occurred, the next prompt you see is in Step 10, on page 35 . If this is the case, please skip to Step 10, on page 35.

For example:

```
================================================
Upgrading Set /var/tmp/cluster-upgrade-set-2.txt
================================================
Checking if reboot required for below hosts
```

```
pcrfclient02 lb02 sessionmgr02 qns02 <-- These VMs may differ on your deployment.
================================================
WARN - Kernel will be upgraded on below hosts from current set of hosts. To take the effect of new
 kernel below hosts will be rebooted.
Upgrading kernel packages on:
  pcrfclient02 lb02 sessionmgr02 qns02
================================================
```

Step 3 Enter **y** to proceed with the kernel upgrade.

Important The kernel upgrade is mandatory. If you enter **n** at the prompt, the upgrade process is aborted.

Step 4 (Optional) The upgrade proceeds until the following message is displayed:

```
All VMs in /var/tmp/cluster-upgrade-set-2.txt are Whisper READY.
Run 'diagnostics.sh --get_replica_status' in another terminal to check DB state.
Please ensure that all DB members are UP and are to correct state i.e. PRIMARY/SECONDARY/ARBITER.
Continue the upgrade for Next Step? [y/n]
```

Step 5 (Optional) In the second terminal to the Cluster Manager VM, run the following command to check the database members are UP and in the correct state:

```
diagnostics.sh --get_replica_status
```

Step 6 (Optional) After confirming the database member state, enter **y** on first terminal to continue the upgrade.

Step 7 (Optional) The upgrade proceeds until the following message is displayed:

```
rebooting pcrfclient01 VM now
pcrfclient01 VM is Whisper READY.
Run 'diagnostics.sh --get_replica_status' in another terminal to check DB state.
Please ensure that all DB members are UP and are to correct state i.e. PRIMARY/SECONDARY/ARBITER.
Continue the upgrade for the Next Step? [y/n]
```

Step 8 (Optional) In the second terminal to the Cluster Manager VM, run the following command to check the database members are UP and in the correct state:

```
diagnostics.sh --get_replica_status
```

Step 9 (Optional) After confirming the database member state, enter **y** on first terminal to continue the upgrade.

Step 10 The upgrade proceeds until the following message is displayed.

```
Please ensure that all the VMS from the /var/tmp/cluster-upgrade-set-2.txt have been upgraded and
restarted. Check logs for failures
If the stop/start for any qns process has failed, please manually start the same before continuing
 the upgrade.
Continue the upgrade? [y/n]
```

Step 11 Once you verify that all VMs in Set 2 are upgraded and restarted successfully, enter **y** to continue the upgrade.

Step 12 The upgrade proceeds until the message for Cluster Manager VM reboot is displayed.

Note This prompt is only shown when upgrading from a release prior to CPS 12.1.0.

Enter **y** to reboot the Cluster Manager VM.

Once the Cluster Manager VM reboots, the CPS upgrade is complete.

Step 13 Continue with Verify System Status, and Remove ISO Image.

Step 14 Any Grafana dashboards used prior to the upgrade must be manually migrated. Refer to *Migrate Existing Grafana Dashboards* in the *CPS Operations Guide* for instructions.

Offline Software Upgrade to 18.3.0

This section describes the steps to perform an offline software upgrade of an existing CPS 18.2.0 deployment to CPS 18.3.0. The offline procedure does not allow traffic to continue running while the upgrade is being performed.

Offline software upgrade to 18.3.0 is supported **only** from CPS 18.2.0.

Offline software upgrade to 18.3.0 is supported **only** for Mobile installations only.

Prerequisites

Important During the upgrade 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 upgrade has been successfully completed and properly validated.

Before beginning the upgrade:

1. Create a backup (snapshot/clone) of the Cluster Manager VM. If errors occur during the upgrade process, this backup is required to successfully roll back the upgrade.
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 directly will not be backed up and must be reapplied manually after the upgrade is complete.
3. Remove any previously installed patches. For more information on patch removal steps, refer to Remove a Patch.
4. If necessary, upgrade the underlying hypervisor before performing the CPS in-service software upgrade. 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 *CPS Installation Guide for VMware* or *CPS Installation Guide for OpenStack* 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 upgrade process.
6. Synchronize the Grafana information between the OAM (pcrfclient) VMs by running the following command from pcrfclient01:
   - When upgrading from CPS 18.2.0, run `/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 Check the System Health, on page 37

Overview

The offline software upgrade is performed in increments:

1. Download and mount the CPS software on the Cluster Manager VM.
2. By default, offline upgrade is performed on all the VMs in a single set.

Note If there is a kernel upgrade between releases, then upgrade is performed in two sets.

For kernel upgrade, if you still want upgrade is performed in a single set then run the following command:

```
/var/platform/platform/scripts/create-cluster-sets.sh 1
Created /var/tmp/cluster-upgrade-set-1.txt
```

3. Start the upgrade (`install.sh`).
4. Once you proceed with the offline upgrade, there is no automated method to roll back the upgrade.

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 to Download and Mount the CPS ISO Image.

Download and Mount the CPS ISO Image

Step 1 Download the Full Cisco Policy Suite Installation software package (ISO image) from software.cisco.com. Refer to the 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.

Step 3 Execute the following commands to mount the ISO image:

```
mkdir /mnt/iso

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

cd /mnt/iso
```

Step 4 Continue with Verify VM Database Connectivity, on page 38.

Verify VM Database Connectivity

Verify that the Cluster Manager VM has access to all VM ports. If the firewall in your CPS deployment is enabled, the Cluster Manager can not access the CPS database ports.

To temporarily disable the firewall, run the following command on each of the OAM (pcrfclient) VMs to disable IPTables:

IPv4: `service iptables stop`

IPv6: `service ip6tables stop`

The iptables service restarts the next time the OAM VMs are rebooted.

Offline Upgrade with Single Set

Important Perform these steps while connected to the Cluster Manager console via the orchestrator. This prevents a possible loss of a terminal connection with the Cluster Manager during the upgrade process.

By default, offline upgrade with single set is used when there is no kernel upgrade detected. For example, offline upgrade from CPS 18.2.0 to CPS 18.3.0.

The steps performed during the upgrade, including all console inputs and messages, are logged to `/var/log/install_console_<date/time>.log`.

Step 1 Run the following command to initiate the installation script:

```
/mnt/iso/install.sh
```

Step 2 (Optional) You can skip Step 3, on page 38 and Step 4, on page 39 by configuring the following parameters in `/var/install.cfg` file:

```
INSTALL_TYPE
INITIALIZE_ENVIRONMENT
```

Example:

```
INSTALL_TYPE=mobile
INITIALIZE_ENVIRONMENT=yes
```

Step 3 When prompted for the install type, enter **mobile**.

```
Please enter install type [mobile|mog|pats|arbiter|andsf|escef]:
```

Note Offline software upgrade to CPS 18.3.0 is supported only for **mobile** installations.

Currently, eSCEF is an EFT product and is for Lab Use Only. This means it is not supported by Cisco TAC and cannot be used in a production network. The features in the EFT are subject to change at the sole discretion of Cisco.

Step 4 When prompted to initialize the environment, enter **y**.

```
Would you like to initialize the environment... [y|n]:
```

Step 5 When prompted for the type of installation, enter **2**.

```
Please select the type of installation to complete:
1) New Deployment
2) Upgrade to different build within same release (eg: 1.0 build 310 to 1.0 build 311)
   or Offline upgrade from one major release to another (eg: 1.0 to 2.0)
3) In-Service Upgrade from one major release to another (eg: 1.0 to 2.0)
```

Step 6 When prompted to enter the SVN repository to back up the policy files, enter the Policy Builder data repository name.

```
Please pick a Policy Builder config directory to restore for upgrade [configuration]:
```

The default repository name is `configuration`. This step copies the SVN/policy repository from the pcrfclient01 and stores it in the Cluster Manager. After pcrfclient01 is upgraded, these SVN/policy files are restored.

Step 7 (Optional) If prompted for a user, enter `qns-svn`.

Step 8 (Optional) If prompted for the password for `qns-svn`, enter the valid password.

```
Authentication realm: <http://pcrfclient01:80> SVN Repos

Password for 'qns-svn':
```

Step 9 (Optional) If CPS detects that there need to be a kernel upgrade on VMs, the following prompt is displayed:

```
================================================
WARN - Kernel will be upgraded on below hosts from current set of hosts.
To take the effect of new kernel below hosts will be rebooted.

================================================
```

Step 10 The upgrade proceeds until the following message is displayed (when kernel upgrade is detected):

```
Please make sure all the VMs are up and running before continue..
If all above VMs are up and running, Press enter to continue..:
```

Offline Upgrade with Two Sets

Important Perform these steps while connected to the Cluster Manager console via the orchestrator. This prevents a possible loss of a terminal connection with the Cluster Manager during the upgrade process.

The steps performed during the upgrade, including all console inputs and messages, are logged to `/var/log/install_console_<date/time>.log`.

Step 1 Run the following command to initiate the installation script:

```
/mnt/iso/install.sh
```

Step 2 (Optional) You can skip Step 3, on page 40 and Step 4, on page 40 by configuring the following parameters in `/var/install.cfg` file:

```
INSTALL_TYPE
INITIALIZE_ENVIRONMENT
```

Example:

```
INSTALL_TYPE=mobile
INITIALIZE_ENVIRONMENT=yes
```

Step 3 When prompted for the install type, enter **mobile**.

```
Please enter install type [mobile|mog|pats|arbiter|andsf|escef]:
```

Note Offline software upgrade to CPS 18.3.0 is supported only for **mobile** installations.

Currently, eSCEF is an EFT product and is for Lab Use Only. This means it is not supported by Cisco TAC and cannot be used in a production network. The features in the EFT are subject to change at the sole discretion of Cisco.

Step 4 When prompted to initialize the environment, enter **y**.

```
Would you like to initialize the environment... [y|n]:
```

Step 5 When prompted for the type of installation, enter **2**.

```
Please select the type of installation to complete:
1) New Deployment
2) Upgrade to different build within same release (eg: 1.0 build 310 to 1.0 build 311)
   or Offline upgrade from one major release to another (eg: 1.0 to 2.0)
3) In-Service Upgrade from one major release to another (eg: 1.0 to 2.0)
```

Step 6 When prompted to enter the SVN repository to back up the policy files, enter the Policy Builder data repository name.

```
Please pick a Policy Builder config directory to restore for upgrade [configuration]:
```

The default repository name is `configuration`. This step copies the SVN/policy repository from the pcrfclient01 and stores it in the Cluster Manager. After pcrfclient01 is upgraded, these SVN/policy files are restored.

Step 7 (Optional) If prompted for a user, enter `qns-svn`.

Step 8 (Optional) If prompted for the password for `qns-svn`, enter the valid password.

```
Authentication realm: <http://pcrfclient01:80> SVN Repos

Password for 'qns-svn':
```

Step 9 If CPS detects that there need to be a kernel upgrade on VMs, the following prompt is displayed:

```
================================================
WARN - Kernel will be upgraded on below hosts from current set of hosts.
To take the effect of new kernel below hosts will be rebooted.

================================================
```

Step 10 The upgrade set 2 proceeds until the following message is displayed:

```
Please make sure all the VMs are up and running before proceeding for Set2 VMs.
If all above VMs are up and running, Press enter to proceed for Set2 VMs:
```

To evaluate the upgrade set 1, refer to Evaluate Upgrade Set 1, on page 41.

Step 11 The upgrade proceeds until the following message is displayed:

```
Please make sure all the VMs are up and running before continue..
If all above VMs are up and running, Press enter to continue..:
```

Evaluate Upgrade Set 1

At this point in the offline software upgrade, the VMs in Upgrade Set 1 have been upgraded.

Before continuing with the upgrade of the remaining VMs in the cluster, check the health of the Upgrade Set 1 VMs. If any of the following conditions exist, the upgrade should be stopped.

- Errors were reported during the upgrade of Set 1 VMs.
- `about.sh` does not show the correct software versions for the upgraded VMs (under CPS Core Versions section).
- All database members (PRIMARY/SECONDARY/ARBITER) are in good state.

Note `diagnostics.sh` reports errors about haproxy that the Set 2 Policy Director (Load Balancer) diameter ports are down, because calls are now being directed to the Set 1 Policy Director. These errors are expected and can be ignored.

If clock skew is seen with respect to VM or VMs after executing `diagnostics.sh`, you need to synchronize the time of the redeployed VMs.

For example,

```
Checking clock skew for qns01...[FAIL]
  Clock was off from lb01 by 57 seconds.  Please ensure clocks are synced. See:
/var/qps/bin/support/sync_times.sh
```

Synchronize the times of the redeployed VMs by running the following command:

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

For more information on `sync_times.sh`, refer to *CPS Operations Guide*.

Verify System Status

The following commands can be used to verify that all CPS components were successfully upgraded and that the system is in a fully operational state:

- **about.sh** - This command displays the updated version information of all components.

- **diagnostics.sh** - This command runs a set of diagnostics and displays the current state of the system. If any components are not running red failure messages will be displayed.

After confirming that CPS has been upgraded and all processes are running normally:

- Reapply any non-standard customizations or modifications to the system that you backed up prior to the upgrade.
- Reapply any patches, if necessary.

Remove ISO Image

Step 1 (Optional) After the upgrade 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
```

Configure Redundant Arbiter (arbitervip) between pcrfclient01 and pcrfclient02

After the upgrade is complete, if the user wants a redundant arbiter (ArbiterVIP) between pcrfclient01 and pcrfclient02, perform the following steps:

Currently, this is only supported for HA setups.

Step 1 Update the `AdditionalHosts.csv` and `VLANs.csv` files with the redundant arbiter information:

- **Update AdditionalHosts.csv:**

  Assign one internal IP for Virtual IP (arbitervip).

  Syntax:

  <alias for Virtual IP>,<alias for Virtual IP>,<IP for Virtual IP>

  For example,

  `arbitervip,arbitervip,` *< IP for Virtual IP>*

- **Update VLANs.csv:**

Add a new column **Pcrfclient VIP Alias** in the `VLANs.csv` file to configure the redundant arbiter name for the pcrfclient VMs:

Figure 1: VLANs.csv

	VLAN Name	Network Target Name	Netmask	Gateway	VIP Alias	Pcrfclient VIP Alias	guestNic
2	Internal	VM Network	255.255.255.0	NA	lbvip02	arbitervip	eth0
3	Management	VLAN 94	255.255.255.0	NA	lbvip01		eth1
4	Gx	VM Network	255.255.255.0	NA	lbvip03		eth2
5							

Execute the following command to import csv files into the Cluster Manager VM:

```
/var/qps/install/current/scripts/import/import_deploy.sh
```

This script converts the data to JSON format and outputs it to `/var/qps/config/deploy/json/`.

Step 2 SSH to the pcrfclient01 and pcrflclient02 VMs and run the following command to create arbitervip:

```
/etc/init.d/vm-init-client
```

Step 3 Synchronize `/etc/hosts` files across VMs by running the following command the Cluster Manager VM:

```
/var/qps/bin/update/synchosts.sh
```

Moving Arbiter from pcrfclient01 to Redundant Arbiter (arbitervip)

In this section we are considering the impacts to a session database replica set when the arbiter is moved from the pcrfclient01 VM to a redundant arbiter (arbitervip). The same steps need to be performed for SPR/balance/report/audit/admin databases.

Step 1 Remove pcrfclient01 from replica set (set01 is an example in this step) by executing the following command from Cluster Manager:

To find the replica set from where you want to remove pcrfclient01, refer to your `/etc/broadhop/mongoConfig.cfg` file.

```
build_set.sh --session --remove-members --setname set01
```

This command asks for member name and port number. You can find the port number from your `/etc/broadhop/mongoConfig.cfg` file.

```
Member:Port   --------> pcrfclient01:27717
pcrfclient01:27717
Do you really want to remove [yes(y)/no(n)]: y
```

Step 2 Verify whether the replica set member has been deleted by executing the following command from Cluster Manager:

```
diagnostics.sh --get_replica_status

|-------------------------------------------------------------------------------------------|
| SESSION:set01                                                                             |
```

```
|  Member-1  -  27717 : 221.168.1.5  -  PRIMARY   -  sessionmgr01 - ON-LINE  -  --------  -  1 |
|  Member-2  -  27717 : 221.168.1.6  -  SECONDARY -  sessionmgr02 - ON-LINE  -   0 sec    -  1 |
|-----------------------------------------------------------------------------------------------|
```

The output of `diagnostics.sh --get_replica_status` should not display pcrfclient01 as the member of replica set (set01 in this case).

Step 3 Change arbiter member from pcrfclient01 to redundant arbiter (arbitervip) in the `/etc/broadhop/mongoConfig.cfg` file by executing the following command from Cluster Manager:

```
vi /etc/broadhop/mongoConfig.cfg
[SESSION-SET1]
SETNAME=set01
OPLOG_SIZE=1024
ARBITER=pcrfclient01:27717             <-- change pcrfclient01 to arbitervip
ARBITER_DATA_PATH=/var/data/sessions.1
MEMBER1=sessionmgr01:27717
MEMBER2=sessionmgr02:27717
DATA_PATH=/var/data/sessions.1
[SESSION-SET1-END]
```

Step 4 Add a new replica set member by executing the following command from Cluster Manager:

```
build_set.sh --session --add-members --setname set01
             Replica-set Configuration
-------------------------------------------------------------------------------
REPLICA_SET_NAME: set01
  Please select your choice for replica sets
            sharded (1) or non-sharded (2): 2


  Adding members to replica set             [  Done  ]
```

The progress of this script can be monitored in the `build_set` log file. For example, `/var/log/broadhop/scripts/build_set_23102017_220554.log`

Step 5 Verify whether the replica set member has been created by executing the following command from Cluster Manager:

```
diagnostics.sh --get_replica_status

|-----------------------------------------------------------------------------------------------|
| SESSION:set01                                                                                 |
|  Member-1  -  27717 : 221.168.1.5   -  PRIMARY    -  sessionmgr01  - ON-LINE  -  --------  -  1 |
|  Member-2  -  27717 : 221.168.1.6   -  SECONDARY  -  sessionmgr02  - ON-LINE  -   0 sec    -  1 |
|  Member-3  -  27717 : 221.168.1.9   -  ARBITER    -  arbitervip    - ON-LINE  -  --------  -  1 |
|-----------------------------------------------------------------------------------------------|
```

The output of `diagnostics.sh --get_replica_status` should now display arbitervip as the member of replica set (set01 in this case).

Troubleshooting

If an error is reported during the upgrade, the upgrade process is paused in order to allow you to resolve the underlying issue.

No Cluster Set Files Found

If you did not run the following script before starting the in-service upgrade:

```
/mnt/iso/platform/scripts/create-cluster-sets.sh
```

You will receive the following error:

```
WARNING: No cluster set files detected.
In a separate terminal, run the create-cluster-sets.sh before continuing.
See the upgrade guide for the location of this script.
After running the script, enter 'y' to continue or 'n' to abort. [y/n]:
```

Run the `create-cluster-sets.sh` script in a separate terminal, then enter **y** to continue the upgrade.

The location of the script depends on where the iso is mounted. Typically it is mounted to `/mnt/iso/platform/scripts`.

VMs Not in Ready State

Whisper is a process used by the CPS cluster to monitor the status of individual VMs. If the Whisper process itself does not start properly or shows status errors for one or more VMs, then the upgrade cannot proceed. In such a case, you may receive the following error:

```
The following VMs are not in Whisper READY state:
 pcrfclient02
 See log file for details:  /var/log/puppet_update_2016-03-07-1457389293.log
 WARNING: One or more VMs are not in a healthy state. Please address the failures before
continuing.
After addressing failures, hit 'y' to continue or 'n' to abort. [y/n]:
```

Whisper Not Running on All VMs

In a separate terminal, log in to the VM that is not in Whisper READY state and run the following command:

```
monit summary | grep whisper
```

If Whisper shows that it is not "Running", attempt to start the Whisper process by running the following command:

```
monit start whisper
```

Run `monit summary | grep whisper` again to verify that Whisper is now "Running".

Verify Puppet Scripts Have Completed Successfully

Check the `/var/log/puppet.log` file for errors.

Run the puppet scripts again on the VM by running the following command

```
/etc/init.d/vm-init-client
```

If the above steps resolve the issue, then proceed with the upgrade by entering **y** at the prompt.

Cannot Set Mongo Priorities

You will receive the following error if the upgrade process cannot reconfigure the Mongo database priorities during the upgrade of Set 1 or Set 2 VMs.

```
WARNING: Mongo re-configuration failed for databases in /var/tmp/cluster-upgrade-set-1.txt.

Please investigate. After addressing the issue, enter 'y' to continue or 'n' to abort.
```

```
[y/n]:
```

Verify that the Cluster Manager VM has connectivity to the Mongo databases and the Arbiter VM. The most common cause is that the firewall on the pcrfclient01 VM was not disabled before beginning the upgrade. Refer to Verify VM Database Connectivity, on page 30 for more information.

Once the connectivity is restored, enter **y** to re-attempt to set the priorities of the Mongo database in the upgrade set.

Cannot Restore Mongo Priorities

You will receive the following error if the upgrade process cannot restore the Mongo database priorities following the upgrade of Set 1 or Set 2 VMs:

```
WARNING: Failed to restore the priorities of Mongo databases in
/var/tmp/cluster-upgrade-set-1.txt. Please address the issue in a separate terminal and
then select one of the following options [1-3]:
   [1]: Continue upgrade. (Restore priorities manually before choosing this option.)
   [2]: Retry priority restoration.
   [3]: Abort the upgrade.
```

Select one of the options, either **1** or **2** to proceed with the upgrade, or **3** to abort the upgrade. Typically there will be other console messages which give indications of the source of this issue.

Note Option 1 does **not** retry priority restoration. Before selecting option 1, you must resolve the issue and restore the priorities manually. The upgrade will not recheck the priorities if you select Option 1.

Timezone Reset

If the timezone was set manually on the CPS VMs using the `/etc/localtime` file, the timezone may be reset on CPS VMs after the upgrade. During the CPS upgrade, the glibc package is upgraded (if necessary) and resets the `localtime` file. This is a known glibc package issue. Refer to https://bugzilla.redhat.com/show_bug.cgi?id=858735 for more information.

As a workaround, in addition to changing the timezone using `/etc/localtime`, also update the Zone information in `/etc/sysconfig/clock`. This will preserve the timezone change during an upgrade.

Error Determining qns Count

During an ISSU, all qns processes are stopped on the CPS VMs. If the upgrade cannot determine the total number of qns processes to stop on a particular VM, you will receive a message similar to the following:

```
Attempting to stop qns-2 on pcrfclient02
Performed monit stop qns-2 on pcrfclient02
Error determining qns count on lb02
Please manually stop qns processes on lb02 then continue.
Continue the upgrade ? [y/n]
```

In a separate terminal, ssh to the VM and issue the following command to manually stop each qns process:

```
/usr/bin/monit stop qns-<instance id>
```

Use the `monit summary` command to verify the list of qns processes which need to be stopped.

logback.xml Update

If the `/etc/broadhop/logback.xml` or `/etc/broadhop/controlcenter/logback.xml` files have been manually modified on the Cluster Manager, the modifications may be overwritten during the upgrade process. A change in logback.xml is necessary during upgrade because certain upgraded facilities require changes to their respective configurations in logback.xml as the facility evolves.

During an upgrade, the previous version of logback.xml is saved as `logback.xml-preupgrade-<date and timestamp>`. To restore any customizations, the previous version can be referenced and any customizations manually applied back to the current logback.xml file. To apply the change to all the VMs, use the `copytoall.sh` utility. Additional information about `copytoall.sh` can be found in the *CPS Operations Guide*.

about.sh Reports Different Versions for the Same Component after the Update

If after running `about.sh`, CPS returns different versions for the same component, run the `restartall.sh` command again to make sure all of the Policy Server (qns) instances on each node have been restarted.

`restartall.sh` performs a rolling restart that is not service impacting. Once the rolling restart is complete, re-run `about.sh` to see if the CPS versions reflect the updated software.

Upgrade 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 Upgrade (ISSU) is not progressing correctly or needs to be abandoned after evaluation of the new version.

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

Rollback Considerations

- You must have a valid Cluster Manager VM backup (snapshot/clone) which you took prior to starting the upgrade.
- You must have the backup archive which was generated at the beginning of the ISSU.
- The upgrade 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 are not backed up or restored.
- Rollback is not supported with a `mongoConfig.cfg` file that has sharding configured.
- Before doing rollback, check the OPLOG_SIZE entry in `/etc/broadhop/mongoConfig.cfg` file.

  If the entry is not there and you have a default `--oplogSize = 1024` value (run `ps -eaf | grep oplog` command from Session Mgr), then add `OPLOG_SIZE=1024` entry in your `/etc/broadhop/mongoConfig.cfg` file for all the replica-sets. Use the value from the output of the `ps` command.

  Example:

```
[SESSION-SET1]
SETNAME=set01
OPLOG_SIZE=1024
ARBITER1=pcrfclient01:27717
```

```
ARBITER_DATA_PATH=/var/data/sessions.1/set01
MEMBER1=sessionmgr01:27717
MEMBER2=sessionmgr02:27717
DATA_PATH=/var/data/sessions.1/set01
```

Once you have updated `mongoConfig.cfg` file, run `/var/qps/install/current/scripts/build/build_etc.sh` script to update the image on Cluster Manager.

Run the following commands to copy the updated `mongoConfig.cfg` file to pcrfclient01/02.

```
scp /etc/broadhop/mongoConfig.cfg pcrfclient01:/etc/broadhop/mongoConfig.cfg

scp /etc/broadhop/mongoConfig.cfg pcrfclient02:/etc/broadhop/mongoConfig.cfg
```

- For deployments using an arbiter VIP, the arbiter VIP must be set to point to the pcrfclient01 before beginning the ISSU or Rollback.
- For replica sets, a rollback does not guarantee that the primary member of the replica set remains the same after a rollback is complete. For example, if sessionmgr02 starts off as the primary, then an ISSU can 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 remains as secondary. In this case, you must manually reconfigure the primary after the rollback, if desired.

Rollback the Upgrade

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

Step 1 Log in to the Cluster Manager VM.

Step 2 Run the following command to prepare the Upgrade Set 1 VMs for removal:

```
/var/qps/install/current/scripts/modules/rollback.py -l <log_file> -a quiesce
```

Specify the log filename by replacing the <*log_file*> variable.

After the rollback.py script completes, the console will display output similar to the following:

```
INFO Host pcrfclient02 status.............................[READY]
INFO Host lb02         status.............................[READY]
INFO Host sessionmgr02 status.............................[READY]
INFO Host qns02        status.............................[READY]
INFO Host qns04        status.............................[READY]
INFO VMs in set have been successfully quiesced
```

Refer to Rollback Troubleshooting, on page 49 if any errors are reported.

Step 3 Take a backup of the log file created by the rollback.py script.

Step 4 If no errors are reported, revert the Cluster Manager VM back to the older version that was taken before the upgrade was started.

Step 5 After reverting the Cluster Manager VM, run `about.sh` to check the VM connectivity with the other VMs in the CPS cluster.

Step 6 Delete (remove) the Upgrade Set 1 VMs using your hypervisor.

Step 7 Redeploy the original Upgrade Set 1 VMs:

- **VMware:** Issue the following command from the Cluster Manager VM to deploy each VM individually. Refer to the *Manual Deployment* section of the *CPS Installation Guide for VMware* for more information about this command.

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

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

- **OpenStack:** Refer to the *Create CPS VMs using Nova Boot Commands* or *Create CPS VMs using Heat* sections of the *CPS Installation Guide for OpenStack* for instructions to redeploy the VMs in an OpenStack environment.

Note After redeployment of Set 1 VMs, traffic is handled by the Set1 VMs immediately.

Step 8 After the Cluster Manager VM is reverted, copy the ISSU backup archive to the reverted Cluster Manager VM. It should be copied to `/var/tmp/issu_backup-<timestamp>.tgz`.

Step 9 Extract the ISSU backup archive:

```
tar -zxvf issu_backup-<timestamp>.tgz
```

Step 10 After the original VMs are redeployed, run the following command to enable these VMs within the CPS cluster:

```
/var/tmp/rollback.py -l <log_file> -a enable
```

Specify the log filename by replacing the <*log_file*> variable.

Note This step adds the member to mongo replica-sets for redeployed VMs, synchronize the statistics, synhcronize the grafana database and so on.

Step 11 During the enablement phase of the rollback, the following prompt appears several times (with different messages) as the previous data and configurations are restored. Enter **y** to proceed each time.

```
Checking options and matching against the data in the archive...
 --svn : Policy Builder configuration data will be replaced
 Is it OK to proceed? Please remember that data will be lost if it has not been properly backed up
 [y|n]:
```

Step 12 When the command prompt returns, confirm that the correct software version is reported for all VMs in the CPS cluster by running `about.sh`.

Step 13 Manually replace any customizations after performing the rollback.

Step 14 Run **diagnostics.sh** to check the health of the CPS cluster.

After the VMs have been redeployed and enabled, follow any repair actions suggested by **diagnostics.sh** before proceeding further.

Refer to Rollback Troubleshooting, on page 49 if any errors are reported.

Rollback Troubleshooting

The following sections describe errors which can occur during an upgrade rollback.

Failures During Backup Phase

During the phase where the ISSU backup archive is created, you may see the following error:

```
INFO     Performing a system backup.
 ERROR    Not enough diskspace to start the backup.
 ERROR: There is not enough diskspace to backup the system.
 In a separate terminal, please clear up at least
 10G and enter 'c' to continue or 'a' to abort:
```

The creation of the ISSU backup archive requires at least 10 GB of free disk space.

If you see this error, open a separate terminal and free up disk space by removing unneeded files. Once the disk space is freed, you can enter **c** to continue.

The script will perform the disk space check again and will continue if it now finds 10 GB of free space. If there is still not enough disk space, you will see the prompt again.

Alternatively, you can enter **a** to abort the upgrade.

Failures During the Quiesce Phase

During the quiesce phase where the upgraded set 1 VMs are taken out of service, you may see the following errors:

```
INFO     Host pcrfclient02    status..............................[READY]
 INFO     Host lb02            status..............................[READY]
 INFO     Host sessionmgr02    status..............................[FAIL]
 INFO          Could not stop Mongo processes. May already be in stopped state
 INFO          Could not remove from replica sets
 INFO     Host qns02           status..............................[READY]
 INFO     Host qns04           status..............................[FAIL]
 INFO          Graceful shutdown failed
 INFO     VMs in set have been quiesced, but there were some failures.
 INFO     Please investigate any failures before removing VMs.
```

These may also be accompanied with other error messages in the console. Since the quiesce phase is expected to occur during a possible failed upgrade, it may be ok for there to be failures. You should investigate the failures to make sure they are not severe. If the failures will not affect the rollback, then they may be ignored. Here are some things to look at for each failure:

Could Not Stop Mongo Processes

If this happens, you can run **diagnostics.sh** to see the state of the session managers. If the mongo processes are already stopped, then no action is necessary. If the session managers in set 1 have been removed from the replica set, then no action is necessary and you can continue with the rollback.

If the mongo processes are not stopped, log onto the session manager and try to stop the mongo processes manually by running this command:

```
/etc/init.d/sessionmgr-<port> stop
```

Run this for each port that has a mongo replica set. The mongo configuration file in `/etc/broadhop/mongoConfig.cfg` will tell you what the ports should be, as well as the output of **diagnostics.sh**.

Could Not Remove From Replica Sets

If the session managers have not been removed from the replica set, then this will need to be done manually before continuing the rollback.

This can be done by logging in to the primary of each replica set and using the mongo commands to remove the session managers in set 1 from each replica set.

If the session manager that is in set 1 happens to be the primary, it needs to step down first. You should not attempt to continue the rollback until all session managers in set 1 have been completely removed from the replica sets.

Graceful Shutdown Failed

If the VMs in set 1 are in a failed state, it is possible that the rollback script will be unable to shut down their monit processes. To investigate, you can ssh into the failed VMs and try to stop all monit processes manually by running this command:

```
monit stop all
```

If the monit processes are already stopped, then no action is necessary. If the VM is in such a failed state that monit processes are stuck or the VM has become unreachable or unresponsive, then there is also no action necessary. You will be removing these VMs anyway, so redeploying them should fix these issues.

Failures in Enable Phase

During this phase the restored VMs are enabled within the CPS cluster. The enable phase consists of the following steps:

1. Add session managers to replica sets.
2. Synchronize statistics from pcrfclient01 to pcrfclient02.
3. Synchronize grafana database from pcrfclient01 to pcrfclient02.
4. Restore SVN repository.
5. Restore `/etc` configuration files.
6. Restore users.
7. Restore authentication information.

Add Session Managers to Replica Sets

If a failure occurs when adding Session Managers to the mongo replica sets, the following message will be displayed:

```
ERROR: Adding session manager VMs to mongo failed.
Please try to manually resolve the issue before continuing.

Enter 'c' to continue or 'a' to abort:
```

There are many conditions which could cause this step to fail. To resolve this issue, manually remove the Upgrade Set 1 session managers from the replica set and then re-add the session managers, as follows:

1. Stop the Mongo processes on the Upgrade Set 1 session manager VMs.

   ```
   service sessionmgr-<port> stop
   ```

2. Remove the session managers from the replica sets. Execute the follow command for each replica set member in set 1.

   ```
   /var/qps/install/current/scripts/build/build_set.sh --<replica set id> --remove-members
   ```

   Note: The replica set id "REPORTING" must be entered as "report" for the replica set id option.

3. Add the session managers back to the replica sets. Repeat the following command for each replica set listed in /etc/broadhop/monogConfig.cfg.

   ```
   /var/qps/install/current/scripts/build/build_set.sh --<replica set id> --add-members
   --setname <replica set name>
   ```

Note The replica set id "REPORTING" must be entered as "report" for the replica set id option.

The replica set information is stored in the `/etc/broadhop/mongoConfig.cfg` file on the Cluster Manager VM. Consult this file for replica set name, member hosts/ports, and set id.

Mongo Priorities

If you receive errors from Mongo, the database priorities may not be set as expected. Run the following command to correct the priorities:

```
/var/qps/install/current/scripts/bin/support/mongo/set_priority.sh
```

Synchronize Statistics from pcrfclient01 to pcrfclient02

If the statistics fail to synchronize from pcrfclient01 to pcrfclient02, the following message will be displayed:

```
ERROR: rsync stats from pcrfclient01 to pcrfclient02 failed.
Please try to manually resolve the issue before continuing.

Enter 'c' to continue or 'a' to abort:
```

To resolve this error, **ssh** to the pcrfclient02 VM and run the following command:

```
rsync -avz pcrfclient01:/var/broadhop/stats /var/broadhop
```

Take note of any errors and try to resolve the root cause, such as not sufficient disk space on the pcrfclient01 VM.

Synchronize Grafana Database from pcrfclient01 to pcrfclient02

If the grafana database fails to synchronize from pcrfclient01 to pcrfclient02, the following message will be displayed:

```
ERROR: rsync grafana database from pcrfclient01 to pcrfclient02 failed.
Please try to manually resolve the issue before continuing.

Enter 'c' to continue or 'a' to abort:
```

To resolve this error, **ssh** to the pcrfclient02 VM and rsync the grafana database from pcrfclient01 using the appropriate command:

CPS 8.1.0 and later:

```
rsync -avz pcrfclient01:/var/lib/grafana/grafana.db /var/lib/grafana
```

CPS versions earlier than 8.1.0:

```
rsync -avz pcrfclient01:/var/lib/elasticsearch /var/lib
```

Resolve any issues that arise.

Restore SVN Repository

If the restoration of the SVN repository fails, the following message will be displayed:

```
ERROR: import svn failed.
Please try to manually resolve the issue before continuing.

Enter 'c' to continue or 'a' to abort:
```

To manually restore the SVN repository, **cd** to the directory where the issu_backup file was unpacked and execute the following command:

```
/var/qps/install/current/scripts/bin/support/env/env_import.sh --svn env_backup.tgz
```

Resolve any issues that arise.

Restore /etc Configuration Files

If the restoration of the configuration files fails, the following message will be displayed:

```
ERROR: import configuration failed.
Please try to manually resolve the issue before continuing.

Enter 'c' to continue or 'a' to abort:
```

To manually restore the configuration files, **cd** to the directory where the issu_backup file was unpacked and execute the following command:

```
/var/qps/install/current/scripts/bin/support/env/env_import.sh --etc=pcrfclient env_backup.tgz
```

Rename the following files:

```
CONF_DIR=/var/qps/current_config/etc/broadhop
TSTAMP=$(date +"%Y-%m-%d-%s")

mv $CONF_DIR/qns.conf $CONF_DIR/qns.conf.rollback.$TSTAMP
mv $CONF_DIR/qns.conf.import $CONF_DIR/qns.conf

mv $CONF_DIR/authentication-provider.xml
$CONF_DIR/authentication-provider.xml.rollback.$TSTAMP
mv $CONF_DIR/authentication-provider.xml.import $CONF_DIR/authentication-provider.xml

mv $CONF_DIR/logback.xml $CONF_DIR/logback.xml.rollback.$TSTAMP
mv $CONF_DIR/logback.xml.import $CONF_DIR/logback.xml

mv $CONF_DIR/pb/policyRepositories.xml $CONF_DIR/pb/policyRepositories.xml.rollback.$TSTAMP
mv $CONF_DIR/pb/policyRepositories.xml.import $CONF_DIR/pb/policyRepositories.xml

mv $CONF_DIR/pb/publishRepositories.xml $CONF_DIR/pb/publishRepositories.xml.rollback.$TSTAMP
mv $CONF_DIR/pb/publishRepositories.xml.import $CONF_DIR/pb/publishRepositories.xm

unset CONF_DIR
unset TSTAMP
```

Resolve any issues that arise.

Restore Users

If the restoration of users fails, the following message will be displayed:

```
ERROR: import users failed.
Please try to manually resolve the issue before continuing.

Enter 'c' to continue or 'a' to abort:
```

To manually restore users, **cd** to the directory where the issu_backup file was unpacked and execute the following command:

```
/var/qps/install/current/scripts/bin/support/env/env_import.sh --users env_backup.tgz
```

Resolve any issues that arise.

Restore Authentication Information

If the restoration of authentication information fails, the following message will be displayed:

```
ERROR: authentication info failed.
Please try to manually resolve the issue before continuing.

Enter 'c' to continue or 'a' to abort:
```

To manually restore authentication info, **cd** to the directory where the issu_backup file was unpacked and execute the following command:

```
/var/qps/install/current/scripts/bin/support/env/env_import.sh --auth --reinit env_backup.tgz
```

Resolve any issues that arise.

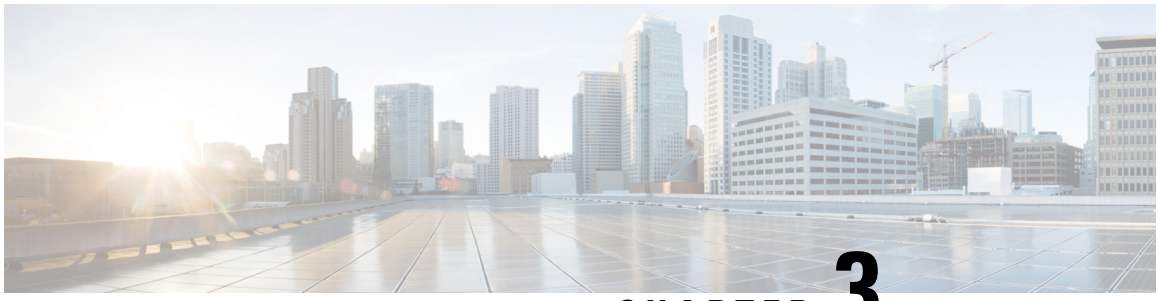

CHAPTER 3

Apply Patches to CPS

Apply a Patch

This section describes the general process to apply a patch to CPS.

Patches must be applied during a maintenance window. This section includes instructions for stopping all CPS components before applying the patch and restarting the components after the patch has been applied.

Note Only one patch can be applied to CPS at a time. If you have already applied a patch, you must Undo and then Remove the existing patch before applying the new patch. Refer to Undo a Patch and Remove a Patch for more information. To determine if a patch is currently applied to the system refer to List Applied Patches.

Step 1 Run **patch -u** and **patch -r** to remove any applied patches from the Cluster Manager before proceeding. For more information, refer to Undo a Patch and Remove a Patch.

Step 2 Download the latest patch file from a location provided by your Cisco representative to the Cluster Manager VM.

Step 3 Log in to the Cluster Manager as a root user.

Step 4 Download the patch file to the Cluster Manager VM. For example:

```
wget http://siteaddress/xxx.tar.gz
```

where,

`siteaddress` is the link to the website from where you can download the patch file.

`xxx.tar.gz` is the name of the patch file.

Step 5 Run the **patch -a** command to apply the patch:

```
/var/qps/install/current/scripts/patch/patch -a filename.tar.gz
```

where *filename* is the path and filename of the downloaded patch file.

For example:

```
/var/qps/install/current/scripts/patch/patch -a /tmp/CPS701_1234.tar.gz
```

Step 6 Run the following command to restore the Policy Builder configurations.

```
/var/qps/install/current/scripts/setup/restorePolicyRepositories.sh
```

Step 7 Run **build_all.sh** script to create updated CPS packages. This builds updated VM images on the Cluster Manager with the new patch applied.

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

Step 8 Update the VMs with the new software using **reinit.sh** script. This triggers each CPS VM to download and install the updated VM images from the Cluster Manager:

```
/var/qps/install/current/scripts/upgrade/reinit.sh
```

Step 9 Refer to section Rolling Restart of CPS VMs QNS Process (Odd Sides) , on page 56 and Rolling Restart of CPS VMs QNS Process (Even Sides), on page 57 for further steps.

Step 10 Run **about.sh** to verify that the component is updated:

```
about.sh
```

What to do next

After applying a patch in HA deployment, run the following command from Cluster Manager:

```
puppet apply --logdest=/var/log/cluman/puppet-custom-run.log
--modulepath=/opt/cluman/puppet/modules --config=/opt/cluman/puppet/puppet.conf
/opt/cluman/puppet/nodes/node_repo.pp
```

Note Manually enter `puppet apply` command in your system.

After applying the `puppet apply` command, run the following command from Cluster Manager to update the `/etc/httpd/conf/httpd.conf` file on all VMs:

```
/var/qps/install/current/scripts/modules/update_httpd_conf.py
```

Rolling Restart of CPS VMs QNS Process (Odd Sides)

Important The commands mentioned in the steps must be entered manually.

Step 1 Stop Policy Server (qns) process:

```
for vmName in `hosts.sh | sort | sed -n 'p;n'`; do echo $vmName; ssh $vmName "service monit stop";
ssh $vmName "service qns stop"; echo; done
```

Step 2 Verify whether the Policy Server (qns) process has stopped:

```
for vmName in `hosts.sh | sort | sed -n 'p;n'`; do echo $vmName; ssh $vmName "service qns status";
echo; done
```

Step 3 Start Policy Server (qns) process:

```
for vmName in `hosts.sh | sort | sed -n 'p;n'`; do echo $vmName; ssh $vmName "service qns start";
ssh $vmName "service monit start"; echo; done
```

Step 4 Verify that the Policy Server (qns) process has started:

```
for vmName in `hosts.sh | sort | sed -n 'p;n'`; do echo $vmName; ssh $vmName "service qns status";
echo; done
```

Step 5 Verify the CPS health status using the `diagnostics.sh` script.

Rolling Restart of CPS VMs QNS Process (Even Sides)

Important The commands mentioned in the steps must be entered manually.

Step 1 Stop Policy Server (qns) process:

```
for vmName in `hosts.sh | sort | sed -n 'n;p'`; do echo $vmName; ssh $vmName "service monit stop";
ssh $vmName "service qns stop"; echo; done
```

Step 2 Verify whether the Policy Server (qns) process has stopped:

```
for vmName in `hosts.sh | sort | sed -n 'n;p'`; do echo $vmName; ssh $vmName "service qns status";
echo; done
```

Step 3 Start Policy Server (qns) process:

```
for vmName in `hosts.sh | sort | sed -n 'n;p'`; do echo $vmName; ssh $vmName "service qns start";
ssh $vmName "service monit start"; echo; done
```

Step 4 Verify that the Policy Server (qns) process has started:

```
for vmName in `hosts.sh | sort | sed -n 'n;p'`; do echo $vmName; ssh $vmName "service qns status";
echo; done
```

Step 5 Verify the CPS health status using the `diagnostics.sh` script.

Undo a Patch

The following steps disables the currently applied CPS patch, and reverts the system to the base software version. For example, if a patch 7.5.0.xx is installed on the system, this command reverts the software to the base version 7.5.0.

Note If you have custom plug-ins installed in your system, refer to CPS Installations using Custom Plug-in before executing the `patch -u` command.

To undo the applied patch, execute the following command on the Cluster Manager:

```
/var/qps/install/current/scripts/patch/patch -u
```

After undoing the applied patch execute the following commands in Cluster Manager to re-build the CPS system and push the changes to VMs:

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

```
/var/qps/install/current/scripts/upgrade/reinit.sh
```

After undoing a patch, qns processes need to be restarted. Refer to Rolling Restart of CPS VMs QNS Process (Odd Sides) , on page 56 and Rolling Restart of CPS VMs QNS Process (Even Sides), on page 57 for further steps.

Remove a Patch

Execute the following command on the Cluster Manager to completely remove a patch and all related items from the Cluster Manager. This deletes the patch file from the `/var/qps/.tmp/patches` directory of the Cluster Manager:

```
/var/qps/install/current/scripts/patch/patch -r patch_name
```

where, *patch_name* is the name of patch you want to remove.

Example,

/var/qps/install/current/scripts/patch/patch -r Patch_1_11.9.9

Note Currently, CPS supports only one patch at a time. You must remove any existing patches before applying a new patch.

After removing a patch, qns processes need to be restarted. Refer to Rolling Restart of CPS VMs QNS Process (Odd Sides) , on page 56 and Rolling Restart of CPS VMs QNS Process (Even Sides), on page 57 for further steps.

List Applied Patches

Execute the following command on Cluster Manager to list the applied patches installed in the system:

```
/var/qps/install/current/scripts/patch/patch -l
```

The `about.sh` command also displays if any patch is applied on the current CPS system or not.

CPS Installations using Custom Plug-in

CPS provides several methods to patch baseline release functionality. One method utilizes the “repositories” configuration file to specify the location of additional software on the CPS Cluster Manger. As such, the current patch utilities aide in removing all repositories. However, CPS Custom plug-in software also uses the “repositories” configuration file to specify the location of custom software. Therefore an additional manual step is required to reconfigure CPS custom plug-in code after patches are removed.

Step 1 From the CPS Cluster Manager, undo the patches:

Note While the patch utility logs that it is removing the repositories configuration file, it actually renames it, at the same path location, as “repositories.back”.

```
/var/qps/install/current/scripts/patch/patch -u
```

The following messages show the progress of the patch -u command:

```
undo the patches
copy puppets from /var/qps/patches backup to /var/qps/install/current/puppet
copy scripts from /var/qps/patches backup to /var/qps/install/current/scripts
remove /etc/broadhop/repositories
patch undone successfully, please run build_all.sh and reinit.sh to push the changes to VMs
```

Step 2 For CPS installations utilizing custom plug-ins, the following step is required before software upgrade.

1. From the CPS Cluster Manager, restore the “repositories” configuration file, without patch references.

   Copy the repositories backup to the original location:

   ```
   cp /etc/broadhop/repositories.back /etc/broadhop/repositories
   ```

2. Remove references to software patch locations, and leave references to custom plugin code:

   In the example below, leave the first line (file:///var/qps/.tmp/plugin1) as it is, and remove the second line (file:///var/qps/.tmp/patch1) before continuing with the software upgrade process.

   ```
   file:///var/qps/.tmp/plugin1

   file:///var/qps/.tmp/patch1
   ```