CPS Migration and Upgrade Guide, Release 25.1.0

First Published: 2025-04-03

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

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED “AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: https://www.cisco.com/c/en/us/about/legal/trademarks.html. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2025 Cisco Systems, Inc. All rights reserved.

CONTENTS

PREFACE

Preface v

About This Guide v

Audience v

Additional Support vi

Conventions (all documentation) vi

Communications, Services, and Additional Information vii

Important Notes viii

CHAPTER 1

Migrate CPS 1

In-Service Migration to 25.1.0 1

Prerequisites 2

Overview 4

Check the System Health 5

Download the CPS ISO Image 5

Create a Backup of CPS 24.2.0 Cluster Manager 6

Migrate the Cluster Manager VM 6

Migrate CPS Set 1 VMs 10

Migrate CPS Set 2 VMs 18

Change Password 25

Change SSH Keys 25

Recover Replica-set Members from RECOVERING State 26

Change SSH Keys - GR Deployment 27

Geographic Redundant Deployment Migration 27

Migrate 3rd Site Arbiter 28

Disable Syncing Carbon Database and Bulk Stats Files 29

HAProxy Diagnostics Warnings 29

Troubleshooting 31
Migration Rollback 31
Rollback Considerations 31
Roll Back the Migration 32
Remove ISO Image 35

CHAPTER 2 Upgrade CPS 37
In-Service Software Upgrade 37

CHAPTER 3 Apply Patches to CPS 39
Apply a Patch 39
Rolling Restart of CPS VMs QNS Process (Odd Sides) 40
Rolling Restart of CPS VMs QNS Process (Even Sides) 41
Undo a Patch 42
Remove a Patch 42
List Applied Patches 43
CPS Installations using Custom Plug-in 43

Preface

- About This Guide, on page v
- Audience, on page v
- Additional Support, on page vi
- Conventions (all documentation), on page vi
- Communications, Services, and Additional Information, on page vii
- Important Notes, on page viii

About This Guide

Note The documentation set for this product strives to use bias-free language. For purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. While any existing biased terms are being substituted, exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on RFP documentation, or language that is used by a referenced third-party product.

This document is a part of the Cisco Policy Suite documentation set.

For information about available documentation, see the *CPS Documentation Map* for this release at Cisco.com.

Note The PATS/ATS, ANDSF, and MOG products have reached end of life and are not supported in this release. Any references to these products (specific or implied), their components or functions in this document are coincidental and are not supported. Full details on the end of life for these products are available at: https://www.cisco.com/c/en/us/products/wireless/policy-suite-mobile/eos-eol-notice-listing.html.

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.

Conventions	Indication
[ ]	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.

Communications, Services, and Additional Information

- To receive timely, relevant information from Cisco, sign up at Cisco Profile Manager.
- To get the business results you're looking for with the technologies that matter, visit Cisco Services.
- To submit a service request, visit Cisco Support.
- To discover and browse secure, validated enterprise-class apps, products, solutions and services, visit Cisco DevNet.
- To obtain general networking, training, and certification titles, visit Cisco Press.
- To find warranty information for a specific product or product family, access Cisco Warranty Finder.

Cisco Bug Search Tool

Cisco Bug Search Tool (BST) is a web-based tool that acts as a gateway to the Cisco bug tracking system that maintains a comprehensive list of defects and vulnerabilities in Cisco products and software. BST provides you with detailed defect information about your products and software.

Important Notes

Important Any feature or GUI functionality that is not documented may not be supported in this release or may be customer specific, and must not be used without consulting your Cisco Account representative.

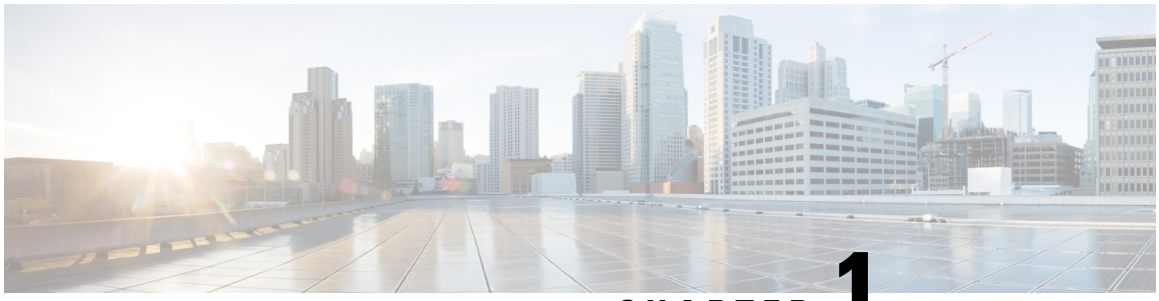

CHAPTER 1

Migrate CPS

- In-Service Migration to 25.1.0, on page 1
- Prerequisites, on page 2
- Overview, on page 4
- Check the System Health, on page 5
- Download the CPS ISO Image, on page 5
- Create a Backup of CPS 24.2.0 Cluster Manager, on page 6
- Migrate the Cluster Manager VM, on page 6
- Migrate CPS Set 1 VMs, on page 10
- Migrate CPS Set 2 VMs, on page 18
- Change Password, on page 25
- Change SSH Keys, on page 25
- Recover Replica-set Members from RECOVERING State, on page 26
- Geographic Redundant Deployment Migration, on page 27
- Migrate 3rd Site Arbiter, on page 28
- Disable Syncing Carbon Database and Bulk Stats Files, on page 29
- HAProxy Diagnostics Warnings, on page 29
- Troubleshooting, on page 31
- Migration Rollback, on page 31
- Remove ISO Image, on page 35

In-Service Migration to 25.1.0

This section describes the steps to perform an In-Service Software Migration (ISSM) of a CPS . This migration allows the traffic to continue running while the migration is being performed.

In-service software migrations to CPS 25.1. are supported only for Mobile (HA) and GR installations. Other CPS installation types cannot be migrated.

Note Before migration, you need to configure at least one Graphite/Grafana user. Grafana supports Graphite data source credential configuration capability. Graphite data source requires common data source credential to be configured using Grafana for Grafana user. Data source credential must be configured before migration. If you fail to add the user, then Grafana will not have an access to Graphite database and you will get continuous prompts for Graphite/Grafana credentials.

All Grafana users configured will be available after migration. However, you need to configure the graphite data source in Grafana UI.

Synchronize the Grafana information between the OAM (pcrfclient) VMs by running `grafana_sync.sh` script from pcrfclient01.

For more information on updating graphite data source, see *Configuring Graphite User Credentials in Grafana* in the *CPS Operations Guide*.

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.

Note During migration, the value of **Session Limit Overload Protection** under System configuration in Policy Builder can be set to 0 (default) which indefinitely accepts all the messages so that the traffic is not impacted but SNMP traps are raised. Once migration is complete, you must change the value as per the session capacity of the setup and publish it without restarting the Policy Server (QNS) process. For more information, contact your Cisco Account representative.

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 `/etc/broadhop/repositories` files before starting the Cluster Manager migration.

Note Before removing the repositories, contact your Cisco Technical Representative.

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 *CPS Installation Guide forVMware* for a list of supported hypervisors for this CPS release.

Note As CPS 25.1.0 supports ESXi 7.0 version until 7.0.3 and ESXi 8.0 version until 8.0.3.

Version 4.6.3 for VMware 7.0/8.0: VMware-ovftool-4.6.3-24031167-lin.x86_64.zip.

https://developer.broadcom.com/tools/open-virtualization-format-ovf-tool/4.6.3

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 Operation Guide* for more information.
7. Check the health of the CPS cluster as described in Check the System Health, on page 5.
8. The following logs must be enabled/set to debug before starting ISSM in `logback.xml` file.

   ```
   <logger name="com.broadhop.utilities.zmq.upgrade.ZMQInServiceUpgradeMgr" level="debug"/>
   ```

   Once ISSM is complete, remove the entry from `logback.xml` file.
9. 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.
10. If you are using IPv6 address, make sure the address you are using is in uncompressed format before starting the migration.

    For example,

    IPv6 in uncompressed format: 2345:f170:8306:8118:e0:208:0:100
11. If you are using Balance feature and Recurring Quota templates in the Policy Builder with **Recurrence Frequency** as **Bill Cycle (RFAmt ignored)**, refer to the *Recurring Quota not Working* section in the *CPS Troubleshooting Guide*.
12. Take the backup of the static route (i.e. route-ifname) and `route -n` collection files.
13. You can upgrade CPS 24.2.0 (using mongoDB version 6.0.14) to CPS 25.1.0 (using mongoDB version 7.0.14). Upgrade to MongoDB 7.0 is supported only from MongoDB 6.0.

    For example, if you are running a 5.0 series (CPS 24.1.0), you must first upgrade to MongoDB 6.0 (CPS 24.2.0) before you can upgrade to MongoDB 7.0 (CPS 25.1.0).

Note Any CPS versions prior to CPS 24.2.0 (using MongoDB version 6.0.14) such as CPS 24.1.0 (using MongoDB version 5.0.20), 23.1.0 or 23.2.0 (using MongoDB version 4.4.18), CPS 22.2.0 (using MongoDB version 4.2.20), or CPS 22.1.1 (using MongoDB version 4.0.27), and previous versions of CPS does not support direct upgrade to CPS 25.1.0 (using MongoDB version 7.0.14). For more information, consult your Cisco Technical Representative.

ISSM Considerations from CPS 24.2 to CPS 25.1

Due to security enhancements and configuration file changes, specific versions of the VMware OVF tool are required for each CPS release.

- **CPS 24.2:** Requires VMware OVF Tool **4.3.0**

  .
- **CPS 25.1:** Requires VMware OVF Tool **4.6.3**

Important

- VMware OVF tool 4.3.0 is **not compatible** with CPS 25.1.

  .
- VMware OVF Tool 4.6.3 is **not compatible** with CPS 24.2.

  .

Recommendation:

Install and use the correct VMware OVF Tool version on your Cluman VM that corresponds to the CPS release you are using. Using an incompatible version may result in errors or unexpected behavior.

Overview

The In-Service Software Migration (ISSM) 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 24.2.0 base image and the same IP address from the old Cluster Manager. The backed up data is then restored on the new Cluster Manager.

Note Graphite data source in Grafana needs to be updated to use configured Graphite/Grafana user credentials before upgrade/migrate start or after fresh installation. If you fail to add this, you will get continuous prompt for Graphite/Grafana credentials as Grafana does not have access to Graphite database.

Synchronize the Grafana information between the OAM (pcrfclient) VMs by running `grafana_sync.sh` script from pcrfclient01.

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 24.2.0 base image. The new VMs are then enabled and restored with the relevant data that was backed up.
4. Migrate CPS VMs Set 2 – After the first set of CPS VMs have been brought back up, the second set are then terminated and brought back up using the 24.2.0 base image. The new CPS VMs are then enabled and restored with the relevant data that was backed up.

Check the System Health

Procedure

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

Procedure

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

What to do next

Validate the md5sum checksum information against the checksum identified by Cisco for the software.

There should be no errors when you run `tar -tzf <file-name>`.

Create a Backup of CPS 24.2.0 Cluster Manager

Before migrating Cluster Manager to CPS 25.1.0, create a backup of the current Cluster Manager in case an issue occurs during migration.

Procedure

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 25.1.0.

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 29 for a workaround.

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 https://www.cisco.com/c/en/us/td/docs/wireless/quantum-policy-suite/R25-1-0/CPS-InstallationGuideforVMware/cps25-1-0installationguide-vmware.html .

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 5, on page 7 .

Procedure

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

```
umount /mnt/iso
```

Step 2 Mount the new CPS 25.1.0 ISO to the existing CPS Cluster Manager running the following command:

Note
You need to mount the existing 24.2 ISO because the system does not support python 3 ISO due to compatibility issues.

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

Step 3 Perform the prerequisite for the ISSM process as specified in the *Upgrade, Migrate, and Rollback Considerations* sectiion.

Step 4 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:

```
2018-07-14 02:39:07,842 INFO [backup.etc] Backup: etc
2018-07-14 02:39:07,878 INFO [__main__.run_recipe] Performing installation stage: Create backup Tar
2018-07-14 02:39:07,905 INFO [__main__.<module>] ======================
2018-07-14 02:39:07,905 INFO [__main__.<module>]  SUCCESS
2018-07-14 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 5 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_20180105_170515.tar.gz
Fetching /var/tmp/migrate_cluman_20170105_170515.tar.gz to migrate_20170105_170515.tar.gz
/var/tmp/migrate_cluman_20180105_170515.tar.gz
```

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

Note
When you move the `*.tar.gz` file to external location or control node, mark the file for easy identification in case you need to restore the configurations from the backup.

Step 6 For VMware, deploy the CPS 25.1.0 Cluster Manager VM following the instructions provided in the CPS Installation Guide for VMware depending on your deployment.

Note
Preserve the old Cluster Manager and create a new Cluster Manager with CPS 25.1.0 as new deployment. Deploy the CPS 25.1.0 Cluster Manager by referring to the following depending on your deployment.

- *Deploy the Cluster Manager VM* in the https://www.cisco.com/c/en/us/td/docs/wireless/quantum-policy-suite/R25-1-0/CPS-InstallationGuideforVMware/cps25-1-0installationguide-vmware.html
- *Installation* and *Orchestration API* chapters in the https://www.cisco.com/c/en/us/td/docs/wireless/quantum-policy-suite/R25-1-0/CPS-InstallationGuideforVMware/cps25-1-0installationguide-vmware.html.

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 7 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 8 Copy the migrate `tar.gz` file from the external location to the new CPS 25.1.0 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_20180720_200701.tar.gz on cluman.
cd /mnt/iso
./migrate.sh restore cluman /root/migrate_20180720_200701.tar.gz
```

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

```
2018-07-21 01:42:21,497 INFO [restore_cluman.restore_fingerprints] Restore fingerprint files.
2018-07-21 01:42:21,531 INFO [restore_cluman.restore_logs] Restoring and copying migrated logs to
archive directory.
2018-07-21 01:42:21,532 INFO [restore_cluman.restore_env_config] Restore cluman env_config files.
2018-07-21 01:42:22,441 INFO [restore_cluman.restore_config_br] Restore cluman config_br files.
2018-07-21 01:42:22,441 INFO [backup.handleRequest] Action Import
2018-07-21 01:42:22,443 INFO [backup.etc] Restore: etc
2018-07-21 01:42:22,544 INFO [__main__.<module>] ====================
2018-07-21 01:42:22,544 INFO [__main__.<module>]  SUCCESS
2018-07-21 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 4, on page 7) .

Step 9 To update the Grafana queries, run the following commands.

```
scp /var/qps/bin/support/grafana_update_query.sh
root@pcrfclient01:/var/qps/bin/support/grafana_update_query.sh
 ssh pcrfclient01 /var/qps/bin/support/grafana_update_query.sh
```

Note
This script execution is mandatory to perform the ISSM from CPS 24.1.0 on previous versions. For more information, see the **grafana_update_query.sh** section in the *CPS Operations Guide*.

Step 10 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 - 22.1.1
CPS Core Versions
--------------------------
         lb01: qns-1           (iomanager): 21.1.0.release
         lb01: qns-2   (diameter_endpoint): 21.1.0.release
         lb01: qns-3   (diameter_endpoint): 21.1.0.release
```

In the example, you can see that the CPS Installer Version was migrated to 25.1.0, 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 8.1.1911 (Core)
```

Note
As AIDO was not running in the older VM sets which were on previous release (for example, 24.2.0), you can observe some failures in diagnostics for AIDO service till all the VMs are migrated to 25.1.0. You can ignore these failures.

diagnostics output sample:

```
Checking AIDO status on all VMs...[PASS]
                    AIDO service is not installed on pcrfclient01, may be pre 21.1 build on
pcrfclient01
                    AIDO service is not installed on pcrfclient02, may be pre 21.1 build on
pcrfclient02
                    AIDO service is not installed on sessionmgr01, may be pre 21.1 build on
sessionmgr01
                    AIDO service is not installed on sessionmgr02, may be pre 21.1 build on
sessionmgr02
                    AIDO service is not installed on sessionmgr03, may be pre 21.1 build on
sessionmgr03
                    AIDO service is not installed on sessionmgr04, may be pre 21.1 build on
sessionmgr04
                    AIDO service is not installed on sessionmgr05, may be pre 21.1 build on
sessionmgr05
                    AIDO service is not installed on sessionmgr06, may be pre 21.1 build on
sessionmgr06
                    AIDO service is not installed on sessionmgr07, may be pre 21.1 build on
sessionmgr07
                    AIDO service is not installed on sessionmgr08, may be pre 21.1 build on
sessionmgr08
```

Step 11 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
```

Note
If you are disabling the carbon database and bulk statistics synchronization (i.e., by setting the SKIP_BLKSTATS=1 and SKIP_CARBONDB=1), then the old Grafana and bulk statistics data are not available on newly deployed CPS system (migrated CPS system).

Note
If the whisper files are of large size, system takes more time to synchronize the carbon database files from pcrfclient01 to pcrfclient02 which can increase the ISSM time.

For example, if `/var/lib/carbon/whisper` is 55 GB, it takes around 15 - 20 hours to synchronize the carbon database files depending on the network speed. To decrease the ISSM time, you can disable the carbon database files and bulk statistics files synchronization as explained in earlier note.

Before starting the migration, it is recommended to run the `/var/qps/install/current/puppet/modules/qps/templates/etc/whisper/clear_wsp_files.sh` utility once on the PCRF nodes to clear any WSP files older than 90 days.

Note
For more information on this step, contact your Cisco Technical Representative.

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 29 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 29.

Note In CPS 23.1.0/CPS 23.2.0/CPS 24.1.0/CPS 24.2.0/CPS 25.1.0, CentOS version 8.1 is replaced with Alma Linux 8.10 with latest rpm packages.

The updated corosync version is not compatible with the previous version corosync. Due to this, there is some traffic loss expected. Traffic loss scenario is only applicable if you are using lbvips for Diameter peering. This is transient and the system recovers automatically once VMs are upgraded to the new Corosync version

Procedure

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

```
/var/qps/install/current/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
```

Note
Before executing `create-cluster-sets.sh` script make sure `Hosts.csv` file has the VM host names in order.

Here is a `Hosts.csv` sample file for reference:

```
[root@localhost csv]# cat Hosts.csv
Hypervisor Name,Guest Name,Role,Alias,Datastore,Networks -->,Internal,Management,
esxi-host-3,lb01,lb01,lb01,datastore3,,192.168.105.13,10.197.98.61,
esxi-host-4,lb02,lb02,lb02,datastore4,,192.168.105.14,10.197.98.62,
esxi-host-3,sessionmgr01,sm,sessionmgr01,datastore3,,192.168.105.15,,
esxi-host-4,sessionmgr02,sm,sessionmgr02,datastore4,,192.168.105.16,,
esxi-host-3,qns01,qps,qns01,datastore3,,192.168.105.17,,
esxi-host-4,qns02,qps,qns02,datastore4,,192.168.105.18,,
esxi-host-3,qns03,qps,qns03,datastore3,,192.168.105.22,,
esxi-host-4,qns04,qps,qns04,datastore4,,192.168.105.25,,
esxi-host-3,pcrfclient01,pcrfclient01,pcrfclient01,datastore3,,192.168.105.19,,
esxi-host-4,pcrfclient02,pcrfclient02,pcrfclient02,datastore4,,192.168.105.20,,
```

Once `create-cluster-sets.sh` is executed, `cluster-upgrade-set-1.txt` and `cluster-upgrade-set-2.txt` are created.

Here are the same output files:

```
[root@localhost scripts]# cat /var/tmp/cluster-upgrade-set-1.txt
pcrfclient02
lb02
sessionmgr02
qns02
qns04
[root@localhost scripts]# cat /var/tmp/cluster-upgrade-set-2.txt
pcrfclient01
lb01
sessionmgr01
qns01
qns03
```

Note
It is recommended to review the `cluster-upgrade-set` files to verify that the database members of a replica-set doesn't belong to the same `cluster-upgrade-set-x`. If Yes, either move any one of the database member to the

other `/var/tmp/cluster-upgrade-set-y.txt` file or change the replica-set definition in `mongocfg.cfg` file. Refer to the *CPS Installation Guide for VMware* for defining a replica-set.

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 .

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.

Note
The `--provision` and `--useprovision` options must be updated to `--nossh`. To use `--nossh` for deploying VM, vCenter 6.5 must be deployed and all existing ESXi hosts must be mapped to the vCenter. For `--usecl`, you should use a customized user instead of vCenter administrator user for which you need to have basic privileges. For nossh feature, you must use `--nossh` and for usecl feature you must use `--nossh --usecl`. To use `--usecl`, content libraries must be pre-configured. The pre-configured content libraries are part of usecl only. For more information, see the *CPS Installation Guide for VMware*.

Add the vCenter hostname and admin credentials either at run time or in `Configuration.csv` file.

Here is a sample configuration:

```
vcenter_hostname,host.cisco.com,
vcenter_user,administrator@vsphere.local,
vcenter_passwd,cisc0123
```

If you are deploying the VMs using the `--nossh`:

- You have to map the ESXi to the vCenter. While mapping, the ESXi must have the same name as ESXi name given in the CPS configurations.
- The vCenter used for the deployment should maintain the unique data store names in the ESXi.

b) (Optional) For Single Cluster Setup, if you have corosync cluster between pcrfclient01 and pcrfclient02 and you want to keep the newly deployed cluster of corosync up. To do so, shutdown the older corosync cluster which hosts arbitervips by executing `monit stop corosync` on Set 2 pcrfclient (pcrfclient01).

When Set 2 gets deployed, pcrfclient01 joins the new cluster normally.

Note
During pcrfclient02 deployment, there is no active arbitervip.

c) (Optional) For Two Cluster Setup or if you have arbitervip between Cluster-A pcrfclient01 and Cluster-B pcrfclient01: Before deploying Set 2 VM's on Cluster-A, execute `monit stop corosync` on Cluster-B pcrfclient01.

Note
Do not start corosync on Cluster-B pcrfclient01 manually.

When Cluster-B Set 2 gets deployed, Cluster-B's pcrfclient01 will join the new cluster normally.

Note
Perform the above step only if you have arbitervips across clusters (Cluster-A and Cluster-B). During Cluster-A's pcrfclient01 deployment, there will be no active arbitervip.

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:

```
2018-07-21 01:53:49,894 INFO [__main__.extra_banner]
        ======================================================================
        | Backing up to file: /var/tmp/migrate_set-1_20180621_212456.tar.gz
        ======================================================================

2018-07-21 02:00:12,252 INFO [backup.handleRequest]
=====================================================
2018-07-21 02:00:12,253 INFO [backup.handleRequest] Archive
/var/tmp/migrate/set-1/config_other_br.tar.gz is created with requested backups.
2018-07-21 02:00:12,253 INFO [backup.handleRequest]
=====================================================
2018-07-21 02:00:12,253 INFO [__main__.run_recipe] Performing installation stage: Create backup Tar
2018-07-21 02:00:12,577 INFO [__main__.<module>] ======================
2018-07-21 02:00:12,578 INFO [__main__.<module>]  SUCCESS
2018-07-21 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_rep
```

Example output is shown below:

```
CPS Diagnostics HA Multi-Node Environment
-----------------------------
Checking replica sets...
|-------------------------------------------------------------------------------------------------------|
| Mongo:3.6.17          MONGODB REPLICA-SETS STATUS INFORMATION        Date : 2021-10-21 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 12, you can restart VM using provisioned vmdk image by running the following command and then go to Step 6, on page 14 .

```
/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 14.

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

Step 6 Re-deploy the Set 1 VMs.

Note
After redeploying, the system takes time to restore the monit services.

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.

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.

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
```

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
```

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 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. After copying static route files, restart the network services and monit processes on these VMs.

Step 8 Run the following command to enable Set 1 VMs:

`/mnt/iso/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_20210921_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:

```
2021-09-21 14:52:15,999 INFO [etc_passwd.parse_etc_passwd] Parsing
/var/tmp/migrate/pcrfclient02/etc/passwd file
2021-09-21 14:52:16,000 INFO [etc_group.parse_etc_group] Parsing
/var/tmp/migrate/pcrfclient02/etc/group file
2021-09-21 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.
2021-09-21 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

2021-10-21 02:45:48,950 INFO [__main__.<module>] ====================
2021-10-21 02:45:48,950 INFO [__main__.<module>]  SUCCESS
2021-10-21 02:45:48,951 INFO [__main__.<module>] ======== END ========
```

Step 9 Execute `diagnostics.sh --get_replica_status` to get the status of the replica-sets in the deployment.

a) Login to the sessionmgr that holds the ADMIN database replica-set using `ssh <replica_set_name>` command.
b) Login to the database using `mongo --port <port-number>`.
c) Execute `show dbs` command to display the database information.

Example:

```
set06:PRIMARY> show dbs
admin                 0.078GB
config                0.078GB
cpsAlarms_cluster-1
  0.078GB
diameter              0.078GB
keystore              0.078GB
local                 4.076GB
policy_trace          2.078GB
queueing              0.078GB
scheduler             0.078GB
sharding              0.078GB
```

d) Execute `use sharding` command.

Example:

```
set06:PRIMARY> use sharding
switched to db sharding
```

Step 10 Execute the following on primary member of ADMIN replica-set using sharding database.

Note
Perform the following steps on all the other admin database replica-sets.

a) Unsetting migrating shards.

```
>db.cache_config.updateMany({},{"$unset":{"migratingShards":1}})
```

b) Change the configuration version.

```
>db.config.update({"_id" : 1},{$inc : { version : 1}})
```

c) Remove versions and insert it back.

```
>db.versions.remove({})
>db.versions.insert({ "_id" : "cache_config", "version" : NumberInt(0), "previousVersion" :
NumberInt(0), "migrationStatus" : "COMPLETE" });
```

Step 11 Check the status of the SkRings and run `rebuildAllSkRings` to create entries by executing the following commands from any OSGi console of any Policy Server (QNS) VM:

Note
The following command is applicable only for memcache setup. You can ignore this step for SKDB set up.

```
>skRingRebuildStatus
>rebuildAllSkRings
```

Step 12 Verify that Set 1 VMs have been migrated by running `about.sh` command:

Example output is shown below:

```
CPS Core Versions
--------------------------
         lb01: qns-1           (iomanager): 25.1.0.release
         lb01: qns-2   (diameter_endpoint): 25.1.0.release
         lb01: qns-3   (diameter_endpoint): 25.1.0.release
         lb01: qns-4   (diameter_endpoint): 25.1.0.release
         lb02: qns-1           (iomanager): 25.1.0.release
         lb02: qns-2   (diameter_endpoint): 25.1.0.release
         lb02: qns-3   (diameter_endpoint): 25.1.0.release
         lb02: qns-4   (diameter_endpoint): 25.1.0.release
        qns01: qns-1               (pcrf): 25.1.0.release
        qns02: qns-1               (pcrf): 25.1.0.release
```

```
        qns03: qns-1                (pcrf): 25.1.0.release
        qns04: qns-1                (pcrf): 25.1.0.release
pcrfclient01: qns-1        (controlcenter): 25.1.0.release
pcrfclient01: qns-2                  (pb): 25.1.0.release
pcrfclient02: qns-1        (controlcenter): 25.1.0.release
pcrfclient02: qns-2                  (pb): 25.1.0.release
```

Step 13 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

2021-10-21 19:56:09,092 INFO [__main__.<module>] ======================
2021-10-21 19:56:09,092 INFO [__main__.<module>]  SUCCESS
2021-10-21 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 31.

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 26.

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 29 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 29.

Procedure

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:

```
2021-10-21 01:53:49,894 INFO [__main__.extra_banner]
        ======================================================================
```

```
        | Backing up to file: /var/tmp/migrate_set-2_20180621_015349.tar.gz
        ============================================================================

2021-10-21 02:00:12,252 INFO [backup.handleRequest]
=====================================================
2021-10-21 02:00:12,253 INFO [backup.handleRequest] Archive
/var/tmp/migrate/set-2/config_other_br.tar.gz is created with requested backups.
2021-10-21 02:00:12,253 INFO [backup.handleRequest]
=====================================================
2021-10-21 02:00:12,253 INFO [__main__.run_recipe] Performing installation stage: Create backup Tar
2021-10-21 02:00:12,577 INFO [__main__.<module>] ======================
2021-10-21 02:00:12,578 INFO [__main__.<module>]  SUCCESS
2021-10-21 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.

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.

Note
The `--provision` and `--useprovision` options must be updated to `--nossh`. To use `--nossh` feature for deploying VM, vCenter 6.5 must be deployed and all existing ESXi hosts must be mapped to the vCenter. For `–usecl` feature, you can use a customized user instead of vCenter administrator user for which you need to have basic privileges. For nossh feature, you must use `--nossh` and for usecl feature you must use `--nossh --usecl`. To use this feature, content libraries must be pre-configured. The pre-configured content libraries are part of usecl feature only. For more information, see the *CPS Installation Guide for VMware*.

Add the vCenter hostname and admin credentials either at run time or in `Configuration.csv` file.

Here is a sample configuration:

```
vcenter_hostname,host.cisco.com,
vcenter_user,administrator@vsphere.local,
vcenter_passwd,cisc0123
```

If you are deploying the VMs using the `--nossh` feature:

- You have to map the ESXi to the vCenter. While mapping, the ESXi must have the same name as ESXi name given in the CPS configurations.
- The vCenter used for the deployment should maintain the unique data store names in the ESXi.

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:3.6.17          MONGODB REPLICA-SETS STATUS INFORMATION        Date : 2021-10-21 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, you can restart VM using provisioned vmdk image by running the following command and then go to Step 5 .

```
/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 Step 5, on page 21

Step 5 Re-deploy the Set 2 VMs.

Note
After redeploying, the system takes time to restore the monit services.

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 the *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.

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
```

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
```

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 6 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. After copying static route files, restart the network services and monit processes on these VMs.

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

`/mnt/iso/migrate.sh 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:

```
2021-10-21 14:52:15,999 INFO [etc_passwd.parse_etc_passwd] Parsing
/var/tmp/migrate/pcrfclient02/etc/passwd file
2021-10-21 14:52:16,000 INFO [etc_group.parse_etc_group] Parsing
/var/tmp/migrate/pcrfclient02/etc/group file
2021-10-21 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.
2021-10-21 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

2021-10-21 20:46:01,621 INFO [__main__.<module>] ======================
2021-10-21 20:46:01,621 INFO [__main__.<module>]  SUCCESS
2021-10-21 20:46:01,621 INFO [__main__.<module>] ======== END ========
```

Step 8 Execute `diagnostics.sh --get_replica_status` to get the status of the replica-sets in the deployment.

a) Login to the sessionmgr that holds the ADMIN database replica-set using `ssh <replica_set_name>` command.

b) Login to the database using `mongo --port <port-number>`.

c) Execute `show dbs` command to display the database information.

Example:

```
set06:PRIMARY> show dbs
admin                 0.078GB
config                0.078GB
cpsAlarms_cluster-1
  0.078GB
diameter              0.078GB
keystore              0.078GB
local                 4.076GB
policy_trace          2.078GB
queueing              0.078GB
scheduler             0.078GB
sharding              0.078GB
```

d) Execute `use sharding` command.

Example:

```
set06:PRIMARY> use sharding
switched to db sharding
```

Step 9 Execute the following on primary member of ADMIN replica-set using sharding database.

a) Unsetting migrating shards.

```
>db.cache_config.updateMany({},{"$unset":{"migratingShards":1}})
```

b) Change the configuration version.

```
>db.config.update({"_id" : 1},{$inc : { version : 1}})
```

c) Remove versions and insert it back.

```
>db.versions.remove({})
>db.versions.insert({ "_id" : "cache_config", "version" : NumberInt(0), "previousVersion" :
NumberInt(0), "migrationStatus" : "COMPLETE" });
```

Step 10 Check the status of the SkRings and run `rebuildAllSkRings` to create entries by executing the following commands from OSGi console of any Policy Server (QNS) VM:

```
>skRingRebuildStatus
>rebuildAllSkRings
```

Step 11 Run diagnostics to verify that the replica set has all of the members back with the correct priorities.

Step 12 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:

```
2021-10-21 20:54:21,083 INFO [command.execute] (stdout): Stopping haproxy: [  OK  ]
Stopping haproxy: [  OK  ]

2021-10-21 20:54:21,083 INFO [__main__.<module>] ====================
2021-10-21 20:54:21,083 INFO [__main__.<module>]  SUCCESS
2021-10-21 20:54:21,083 INFO [__main__.<module>] ======== END ========
```

Step 13 Restart collectd service on pcrfclient01 and pcrfclient02. If you are using GR or dual cluster setups, then you need to execute the following commands on both site's pcrfclient VM. This collectd service restart is required to refresh the memory consumption on pcrfclient VMs which could be high in number because of the carbon stats copy operation.

Note
Make sure that the ISSM is successfully completed on HA or your respective GR/dual cluster sites before restarting collectd service.

The following is an example of how to restart the collectd on pcrfclient01 VM. You need to follow the same procedures on pcrfclient02 VM and on other site pcrfclient VM's if your setup is GR or a dual cluster.

a) SSH to pcrfclient01.

```
ssh root@pcrfclient01
```

b) Stop collectd service.

```
monit stop collectd
```

c) Confirm that the collectd service is no longer monitored and also the service is successfully stopped.

```
monsum | grep -i collectd
collectd                        Not monitored               Process

systemctl status collectd
```

d) Once the service is successfully stopped, bring the collectd service back on.

```
monit start collectd
```

e) Confirm that the service is up and running using `monsum` and `systemctl` commands.

```
monsum | grep -i collectd
systemctl status collectd
```

What to do next

Execute the following command from Cluster Manager to cleanup the backup which was been created at the time of provisioning:

```
/var/qps/install/current/scripts/deployer/support/deploy_all.py --cleanupbackup
```

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.

Caution Executing `restartall.sh` will cause messages to be dropped.

If some of the replica-set members are in RECOVERING state, refer to Recover Replica-set Members from RECOVERING State, on page 26.

Change Password

This section outlines changing password of all VMs.

Run the `change_passwd.sh` script on Cluster Manager to change the password of root, qns, qns-svn, qns-admin and qns-su users across the system.

The `change_passwd.sh` script changes the password on all the VMs temporarily. You also need to generate an encrypted password. To generate encrypted password, refer to *System Password Encryption* in *CPS Installation Guide for VMware*. The encrypted password must be added in the `Configuration.csv` spreadsheet. To make the new password persisent, execute `import_deploy.sh`. If the encrypted password is not added in the spreadsheet and `import_deploy.sh` is not executed, then after running `reinit.sh` script, the qns-svn user takes the existing default password from `Configuration.csv` spreadsheet.

Change SSH Keys

Note If you are using default SSH keys, you are required to change the SSH keys after migration. Make sure migration is completed successfully.

Before you begin

Before changing SSH keys, make sure diagnostics is clean and there is no alarm/warning.

Note It's important to change SSH keys at least once.

Procedure

Step 1 To generate new keys execute the following command on installer VM (Cluster Manager).

```
/var/qps/install/current/scripts/bin/support/manage_sshkey.sh --create
```

Step 2 Update keys on CPS VMs and installer VM (Cluster Manager).

```
/var/qps/install/current/scripts/bin/support/manage_sshkey.sh --update
```

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:

Procedure

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 the *CPS Troubleshooting Guide* to remove failed member and add the member back.

Change SSH Keys - GR Deployment

It is required to modify SSH keys once GR installation/migration is complete.

Before you begin

Make sure diagnostics is clean and there is no alarm/warning.

Procedure

Step 1 On Site 1's Cluster Manager execute the following steps:

a) Generate new keys.

```
/var/qps/install/current/scripts/bin/support/manage_sshkey.sh --create
```

b) Update keys on CPS VMs and Installer VM (Cluster Manager).

```
/var/qps/install/current/scripts/bin/support/manage_sshkey.sh --update
```

c) Backup new SSH keys.

```
cd /var/qps/auth; tar -cvf current_ssh_keys.tar current/
```

Step 2 Transfer `/var/qps/auth/current_ssh_keys.tar` to Site 2 Cluster Manager.

Step 3 On Site 2 execute the following steps.

a) Restore SSH keys.

```
mkdir -p /var/qps/auth/temp/{qns,root};
```

b) Copy tar file to `/var/qps/auth/ ; tar -xvf current_ssh_keys.tar.`

```
cp /var/qps/auth/current/root/* /var/qps/auth/temp/root/ ; cp /var/qps/auth/current/qns/*
/var/qps/auth/temp/qns/ ;
```

c) Update keys on CPS VMs and Installer VM (Cluster Manager).

```
/var/qps/install/current/scripts/bin/support/manage_sshkey.sh --update
```

Geographic Redundant Deployment Migration

Note In CPS 21.1.0/21.2.0/CPS 22.1.1/CPS 22.2.0/CPS 23.1.0/CPS 23.2.0/CPS 24.1.0/CPS 24.2.0 release, Centos version 8.8 is replaced with Alma Linux 8.9 with latest rpm packages.

The updated corosync version is not compatible with the previous version corosync. Due to this, there is some traffic loss expected. Traffic loss scenario is only applicable if you are using lbvips for Diameter peering. This is transient and the system recovers automatically once VMs are upgraded to the new Corosync version

Migrate 3rd Site Arbiter

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

Procedure

Step 1 Copy the new CPS 25.1.0 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 25.1.0 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:

```
2021-10-10 07:51:42,633 INFO [command.execute] Mongo port:27719 stopped successfully
2021-10-10 07:51:42,633 INFO [__main__.run_recipe] Performing installation stage:
ExtractInstallArtifacts
2021-10-10 07:51:42,634 INFO [extract_install_artifacts.extract_scripts] Extracting CPS scripts
2021-10-10 07:51:43,506 INFO [__main__.run_recipe] Performing installation stage: PrepareWorkingDir
2021-10-10 07:51:43,506 INFO [__main__.run_recipe] Performing installation stage: BackupArbiter
2021-10-10 07:52:25,715 INFO [__main__.run_recipe] Performing installation stage: Create backup Tar
2021-10-10 07:52:37,921 INFO [__main__.<module>] ====================
2021-10-10 07:52:37,921 INFO [__main__.<module>]  SUCCESS
2021-10-10 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.gz migrate_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 25.1.0 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 25.1.0 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:

```
2021-10-10 15:45:53,187 INFO [command.execute] child process started successfully, parent exiting
2021-10-10 15:45:53,188 INFO [command.execute] ^[[60G[^[[0;32m  OK  ^[[0;39m]
```

```
2021-10-10 15:45:53,189 INFO [command.execute]
2021-10-10 15:45:53,189 INFO [__main__.<module>] ======================
2021-10-10 15:45:53,189 INFO [__main__.<module>]  SUCCESS
2021-10-10 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) 2022. All rights reserved.

CPS Arbiter

CPS Installer Version - 22.1.1
```

The timezone on the arbiter must be changed to UTC as shown below:

```
cat /etc/*elease
CentOS Linux release 8.1.1911 (Core)
CentOS Linux release 8.1.1911 (Core)
CentOS Linux release 8.1.1911 (Core)
[root@site3-arbiter log]# date
Tue June 15 02:42:54 UTC 2022
```

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.

Procedure

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.

a. 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>
```

b. Enable the endpoints by running the following command:

```
clearExcludedEndpoints
```

c. 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
```

Troubleshooting

If an error is reported during the migration, the migration process is paused. In order to allow you to resolve the underlying issue, refer to the following sections:

Session Cache Database in UNKNOWN State

Issue: Session cache database is in UNKNOWN state. Session cache status can be displayed using `diagnostics.sh --get_replica_status`.

Symptoms: While migrating/upgrading, if there is a power outage or blade issue it impacts both the sets/clusters/sites resulting in inconsistent MongoDB versions between the sets/clusters/sites. For example, one set/cluster/site is upgraded to the latest MongoDB version, for example, 3.6.9 and another set/cluster/site is still on lower MongoDB version, for example, 3.4.16. In this case, auto-recovery for session cache/SK database does not work resulting in UNKNOWN state of the members.

Solution: You can either upgrade to the latest CPS version supporting latest MongoDB version or rollback to previous CPS version supporting older MongoDB version.

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 `/mnt/iso/migrate.sh enable set 1` command. The rollback works only after that command has been run.
- Run `diagnostics.sh --get_replica_status` to check which new Set 1 sessionmgr*XX* (even numbered) MongoDB processes are in RECOVERING state. If so, manually stop all those processes on respective session managers.

  Example:

```
|-----------------------------------------------------------------------------------------|
| Mongo:3.6.17          MONGODB REPLICA-SETS STATUS INFORMATION Date : 2019-07-23
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 MongoDB on that member and check its actual status.

Note If the restore step is already performed, during rollback, it is not neccessary to revert the grafana queries as the updated grafana graphs will work in the previous releases also.

Procedure

Step 1 For Single Cluster Setup, if you have corosync cluster between pcrfclient01 and pcrfclient02 and you want to keep the newly deployed cluster of corosync up. To do so, shutdown the older corosync cluster which hosts arbitervips by executing `monit stop corosync` on Set 2 pcrfclient (pcrfclient01).

When Set 2 gets deployed, pcrfclient01 joins the new cluster normally.

Note
During pcrfclient02 deployment, there will be no active arbitervip.

Step 2 (Optional) For Two Cluster Setup or if you have arbitervip between Cluster-A pcrfclient01 and Cluster-B pcrfclient01: Before deploying Set 2 VM's on Cluster-A, execute `monit stop corosync` on Cluster-B pcrfclient01.

Note
Do not start corosync on Cluster-B pcrfclient01 manually.

When Cluster-B Set 2 gets deployed, Cluster-B's pcrfclient01 will join the new cluster normally.

Note
Perform the above step only if you have arbitervips across clusters (Cluster-A and Cluster-B). During Cluster-A's pcrfclient01 deployment, there will be no active arbitervip.

Step 3 Perform the prerequisite for the ISSM process as specified in the *Upgrade, Migrate, and Rollback Considerations* sectiion.

Step 4 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:

```
2019-07-23 20:10:30,653 INFO [fabric_tasks.run] Stopping all services on remote VM qns04
2019-07-23 20:10:30,654 INFO [transport._log] Secsh channel 3 opened.
2019-07-23 20:10:40,745 INFO [transport._log] Secsh channel 4 opened.
2019-07-23 20:10:42,111 INFO [__main__.<module>] ======================
2019-07-23 20:10:42,111 INFO [__main__.<module>]  SUCCESS
2019-07-23 20:10:42,111 INFO [__main__.<module>] ======== END ========
```

Step 5 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 6 Restore the older Cluster Manager VM (for example, CPS 19.4.0) from the backup (snapshot/clone).

Note
If restoring Cluster Manager on OpenStack environment fails, refer to the *Failure when Restoring Cluster Manager during Rollback on Openstack Environment* section in the *CPS Troubleshooting Guide*.

Step 7 Create cluster sets for migration rollback by running the following command:

`/var/qps/install/current/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 8 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
```

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

Step 10 Mount the *CPS_*.release.iso* to the existing CPS Cluster Manager by running the following command. If the migration was attempted from *x*.iso to *y*.iso, and for the rollback, mount *x* .iso.

Note
Since the Python 2 systems does not support Python 3, you need to mount the current ISO running on Python 2.

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

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

For example, `mount -o loop CPS_20.2.0.release.iso /mnt/iso`

Step 11 Run the following command to enable Set 1 VMs. For example:

`/mnt/iso/migrate.sh 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
```

```
2019-07-23 02:45:48,950 INFO [__main__.<module>] ======================
2019-07-23 02:45:48,950 INFO [__main__.<module>]  SUCCESS
2019-07-23 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`.

Step 12 In the CPS 24.2.0 release, mongo is upgraded to 6. x, where mandatory prerequisites are automated to support in-service migration. To restore previous mongo settings, run the following command:

```
/mnt/iso/modules/mongo_parameters_rollback.sh
```

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

Procedure

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
```

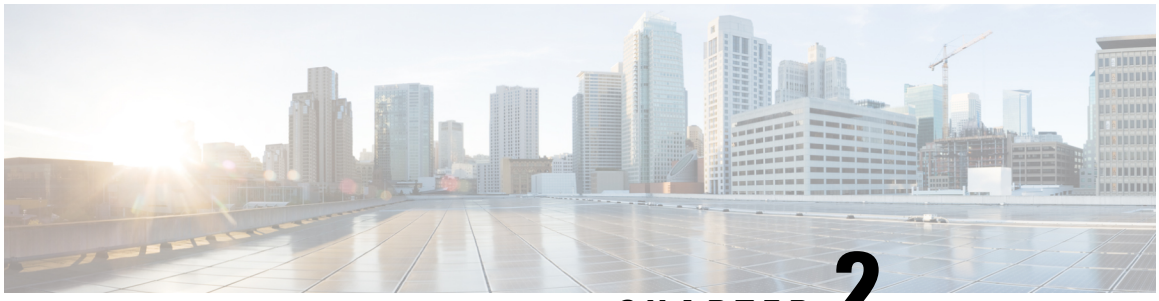

CHAPTER 2

Upgrade CPS

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

- In-Service Software Upgrade, on page 37

In-Service Software Upgrade

In-Service Software Upgrade (ISSU) is strictly **not supported** and **only** the In-Service Software Migration (ISSM) is **supported** for **any CPS** release.

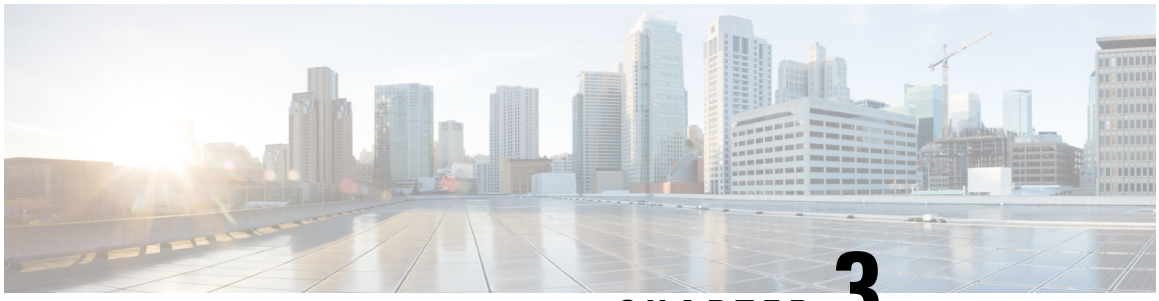

CHAPTER 3

Apply Patches to CPS

- Apply a Patch, on page 39
- Undo a Patch, on page 42
- Remove a Patch, on page 42
- List Applied Patches, on page 43
- CPS Installations using Custom Plug-in, on page 43

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.

Procedure

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 40 and Rolling Restart of CPS VMs QNS Process (Even Sides), on page 41 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.

Procedure

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.

Procedure

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 40 and Rolling Restart of CPS VMs QNS Process (Even Sides), on page 41 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 40 and Rolling Restart of CPS VMs QNS Process (Even Sides), on page 41 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.

Procedure

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.

a. 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
```

b. 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
```