

CPS Installation Guide for VMware, Release 26.1.0

First Published: 2026-04-23

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)

© 2026 Cisco Systems, Inc. All rights reserved.



CONTENTS

PREFACE

Preface ix

About This Guide ix

Audience ix

Additional Support x

Conventions (all documentation) x

Communications, Services, and Additional Information xi

Important Notes xii

CHAPTER 1

Overview 1

Planning the CPS Deployment 2

CPS Dimensioning Evaluation 2

Hardware Requirements 2

Virtual Machine Requirements 3

High Availability Deployment 3

Deployment Examples 7

Platform WSP File Sizing Calculation 9

Sample Customer Deployment 10

Application KPI Metrics Sizing Calculation 11

Sample Customer Deployment 12

Install and Configure VMware 14

Install VMware vSphere Hypervisor (ESXi) 14

Prerequisites 14

Installation 15

Enable SSH 15

Configure VMware ESXi Timekeeping 16

Collect Virtualization Parameters for CPS Installation 17

CHAPTER 2

CPS Installation 19
Obtain the CPS Software 19
Cluster Manager VM 20
Overview 20
Directory Structure 21
Puppet Overview 21
Deploy the Cluster Manager VM 23
Configure Cluster Manager VM 25
Common Steps 25
HA Installation 27
Install VMware OVF tool 4.6.3 28
Change Password 29
Configure System Parameters for Deployment 30
Definitions Configuration 30
VMSpecifications Configuration 31
VLANs Configuration 33
guestNic 34
Hosts Configuration 35
Additional Hosts Configuration 37
NTP Configuration 38
Configuration based on Diameter Endpoints Interface 38
General Configuration 42
System Password Encryption 60
Redis Authentication 61
SCTP Configuration 63
MongoDB Authentication 64
LDAP SSSD Configuration 66
Disable Zing 70
MongoDB Replication Health Monitoring 70
VIP Proxy Configuration 72
Secure Configuration 73
DSCP Configuration 75
Critical File Monitoring Configuration 76

Finish and Save 78
Import the Excel Information into the Cluster Manager VM 78
Save the csv Files 78
Copy the csv Files into Cluster Manager VM 80
Import the csv Files into the Cluster Manager VM 80
Validate Imported Data 81
Update System Parameters 81
Customize Features in the Deployment 81
LDAP Feature Installation 83
Enable LDAP on HA Deployment 83
Subscriber Lookup Feature Installation 84
Enable Subscriber Lookup on HA Deployment 84
License Generation and Installation 84
License Generation 84
License Installation 85
Validate Installed License 86
Upgrade License 87
SSL Certificates 88
Create SSL Certificates 88
Replace SSL Certificates 89
Enable Custom Puppet to Configure Deployment 90
Installing Platform Scripts for MongoDB Health Monitoring - VMware 92

CHAPTER 3 **Deploy CPS VMs** 95
Deploy the VMs 95
Build VM Images 95
Manual Deployment 96
Automatic Deployment of All CPS VMs in Parallel 96
Automatic Deployment of Selective CPS VMs in Parallel 97
Update Default Credentials 98
Initialize SVN Synchronization 99
External Port Matrix 100
Memory Reservation on VMs 100
Session Manager Configuration for Data Replication 100

Guidelines for Choosing MongoDB Ports 100
Supported Databases 101
Prerequisites 101
Script Usage 103
Guidelines for Adding Replica-sets 104
Defining a Replica-set 105
Example of Replica set Creation 107
Guidelines to Configure More than Seven Replica-set Members 108
Session Cache Scaling 109
Service Restart 109
Create Session Shards 109
Verify CPS Sanity 110
Validate VM Deployment 111
Virtual Interface Validation 111
Basic Networking 111
Diagnostics and Status Check 111
diagnostics.sh 111
about.sh 113
list_installed_features.sh 113
statusall.sh 114
Web Application Validation 114
Supported Browsers 115

CHAPTER 4 **Post Installation Processes** 117
Post Installation Configurations 117
Configure Control Center Access 117
Configure NTP on Cluster Manager 117
Change SSH Keys 117
IPv6 Support - VMware 118
Enable IPv6 Support 118
Set Up IPv4 and IPv6 Addresses for VMs 118
Converting IPv4 to IPv6 on Policy Director External Interfaces 119
Synchronize Time Between Nodes 120
Update the VM Configuration without Re-deploying VMs 120

Reserving Memory on the Virtual Machines (VMs) 121
Configure Custom Route 121
TACACS+ 122
TACACS+ Configuration Parameters 122
Arbiter Configuration for TACACS+ 123
TACACS+ Enabler 123
Configure Multiple Redis Instances 124
Configure Redis Instances for Keystore 126
Modify Configuration Files 127
Scaling Existing Installation 127
Adding Member to Existing Replica Set 129
Configure Balance Shards 129
Prerequisites 129
Shard Configuration 129
Add Shards to Balance Database 129
Remove Shards from Balance Database 130
Secondary Key Ring Configuration 131
Why it is Required 131
Key Ring Commands 132
Creating a New Ring 132
Adding a New Endpoint 132
Removing an Endpoint 132
Removing a Ring 132
Triggering a Ring Rebuild 133
Single Cluster Configuration 133
Multi-Cluster Configuration 133
GR Configuration with Session Replication Across Sites 134
Configuring SK DB 134
Upgrading SK DB Manually 134
Upgrading SK DB with Auto – Upgrade 136



Preface

- About This Guide, on page ix
- Audience, on page ix
- Additional Support, on page x
- Conventions (all documentation), on page x
- Communications, Services, and Additional Information, on page xi
- Important Notes, on page xii

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

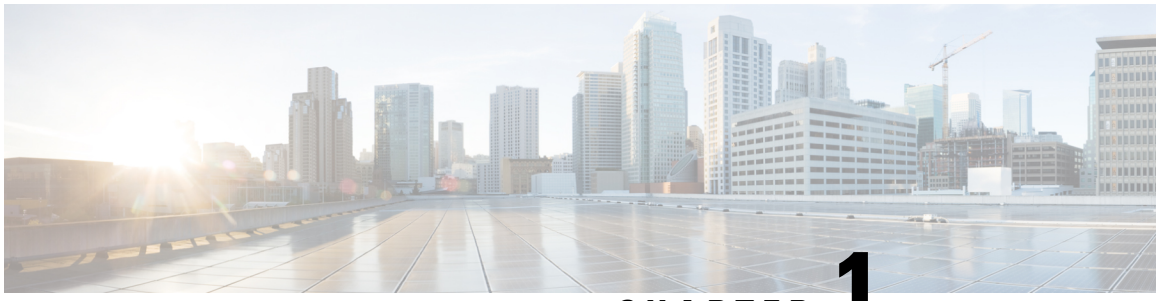
- To receive timely, relevant information from Cisco, sign up at Cisco Profile Manager.
- To get the business results you're looking for with the technologies that matter, visit Cisco Services.
- To submit a service request, visit Cisco Support.
- To discover and browse secure, validated enterprise-class apps, products, solutions and services, visit Cisco Marketplace.
- 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

Overview

Cisco Policy Suite (CPS) is a scalable software-based solution that requires installation of multiple virtual machines prior to the configuration phase.

The preceeding steps outline the basic process for a new installation of CPS:

Chapter 1:

1. Review physical hardware and virtual machine requirements.
2. Install and Configure VMware®.
3. Plan and collect information prior to installation.

The step by step procedures to download, configure and intstall the CPS software are listed in the given topics:

Chapter 2:

1. Download CPS software
2. Deploy Cluster Manager VM
3. Populate CPS Deployment Template file with information for deployment
4. Configure Cluster Manager VM
5. Configure and import the CPS Deployment Template information into Cluster Manager VM
6. Enable any Custom Features
7. Install CPS license
8. Replace Default SSL Certificates

The information on deployment and validation of CPS VMs are covered as a part of preceeding chapter:

Chapter 3:

1. Deploy all other CPS VMs
2. Update Default Credentials
3. Initialize SVN Synchronization
4. Configure Session Manager for Database Replication
5. Validate VM deployment

• Planning the CPS Deployment, on page 2
• Install and Configure VMware, on page 14
• Collect Virtualization Parameters for CPS Installation , on page 17

Planning the CPS Deployment

CPS Dimensioning Evaluation

With assistance from Cisco Technical Representatives, a dimensioning evaluation must be performed for each CPS deployment. This dimensioning evaluation uses customer-specific information such as call model, product features to be used, and traffic profiles to determine the specific requirements for your deployment, including:

- Hardware specifications (number and type of blades, memory, etc.)
- VM information (number, type and resource allocation)

The requirements established in the dimensioning evaluation must be met or exceeded.

The Hardware Requirements, on page 2 and Virtual Machine Requirements, on page 3 sections provide minimum guidelines for a typical CPS deployment.

Hardware Requirements

CPS is optimized for standard Commercial Off-The-Shelf (COTS) blade servers.

The given table provides a summary of the minimum requirements for a typical single-site High Availability (HA) CPS deployment:

Table 1: Hardware Requirements

Minimum Hardware Requirements (Blade Server)	
Memory	The total size of memory for a blade server should be sufficient to meet the memory needs for all the Virtual Machines (VMs) installed in the blade. Refer to the Virtual Machine Requirements, on page 3 section for the amount of memory needed for each VMs. Also consider the memory needed by the Hypervisor. For VMware 5.x it is recommended to reserve 8 GB memory.
Storage	Two (2) 400 GB Enterprise Performance SSD Drives Supporting hardware RAID 1 with write-back cache
Interconnect	Dual Gigabit Ethernet ports
Virtualization	Must be listed in the VMware Compatibility Guide at: https://www.vmware.com/resources/compatibility/search.php

Minimum Hardware Requirements (Blade Server)	
Minimum Hardware Requirements (Chassis)	
Device Bays	A minimum of 4 is required for HA deployments
Interconnect	Redundant interconnect support
Power	Redundant AC or DC power supplies (as required by the service provider)
Cooling	Redundant cooling support

Virtual Machine Requirements

High Availability Deployment

Here is the list of minimum CPU RAM and disk space requirements for each type of CPS virtual machine (VM) in a typical deployment (4 blade single-site high availability):



Important The requirements mentioned in the table is based on:

- Hyper-threading: Enabled (Default)
- CPU Pinning: Disabled
- CPU Reservation: Yes (if allowed by hypervisor)
- Memory Reservation: Yes
- Hard Disk (in GB): 100

Table 2: HA Virtual Machine Requirements - Chassis Architecture

Physical Cores / Blade	VM Type	Memory (in GB)	Hard Disk (in GB)	vCPU	Configuration
Blade with 16 CPUs	Policy Server VMs (QNS)	16	100	12	Threading = 200 Mongo per host = 10 Criss-cross Mongo for Session Cache = 2 on each VM
Blade with 16 CPUs	Session Manager VMs	128	100	6	
Blade with 16 CPUs	Control Center (OAM) VMs	16	100	6	
Blade with 16 CPUs	Policy Director VMs (LB)	32	100	12	
Blade with 16 CPUs	Cluster Manager	12	-	2	

Physical Cores / Blade	VM Type	Memory (in GB)	Hard Disk (in GB)	vCPU	Configuration
Blade with 24 CPUs	Policy Server VMs (QNS)	16	100	10	Threading = 100 Mongo per host = 10 Criss-cross Mongo for Session Cache = 2 on each VM
Blade with 24 CPUs	Session Manager VMs	80	100	8	
Blade with 24 CPUs	Control Center (OAM) VMs	16	100	12	
Blade with 24 CPUs	Policy Director VMs (LB)	32	100	12	
Blade with 24 CPUs	Cluster Manager	12	-	2	

Table 3: HA Virtual Machine Requirements - Cloud Architecture

Physical Cores / Blade	VM Type	Memory (in GB)	Hard Disk (in GB)	vCPU	Configuration
Blade with 16 CPUs	Policy Server VMs	16	100	12+	Threading = 200 Mongo per host = 10 Criss-cross Mongo for Session Cache = 2 on each VM
Blade with 16 CPUs	Session Manager VMs	128	100	6+	
Blade with 16 CPUs	Control Center (OAM) VMs	16	100	6+	
Blade with 16 CPUs	Policy Director VMs	32	100	8+	
Blade with 16 CPUs	Cluster Manager	12	-	2+	
Blade with 24 CPUs	Policy Server VMs	16	100	10+	Threading = 100 Mongo per host = 10 Criss-cross Mongo for Session Cache = 2 on each VM
Blade with 24 CPUs	Session Manager VMs	80	100	8+	
Blade with 24 CPUs	Control Center (OAM) VMs	16	100	12+	
Blade with 24 CPUs	Policy Director VMs	32	100	12+	
Blade with 24 CPUs	Cluster Manager	12	-	2+	

Important On VMware, virtual NUMA topology is enabled by default when the number of virtual CPUs is greater than 8. You can ignore the following warning message displayed on mongo console which is generated when you configure more than 8 vCPUs on VM:

```
2016-06-03T21:40:03.130-0400 [initandlisten]
** WARNING: You are running on a NUMA machine.
2016-06-03T21:40:03.130-0400 [initandlisten]
** We suggest launching mongod like this to avoid performance problems:
2016-06-03T21:40:03.130-0400 [initandlisten] **
```

Note For large scale deployments having Policy Server (qns) VMs more than 35, Session Manager (sessionmgr) VMs more than 20, Policy Director (lb) VMs more than 2, recommended RAM for OAM (pcrfclient) VMs is 64GB.

Note For large scale deployments having Policy Server (qns) VMs more than 32, Session Manager (sessionmgr) VMs more than 16, Policy Director (lb) VMs more than 2, recommended vCPU for OAM (pcrfclient) VMs is 12+.

Note If CPS is deployed in a cloud environment where over-allocation is possible, it is recommended to enable hyper-threading and double the number of vCPUs.

Note The hard disk size of all VMs are fixed at 100 GB (thin provisioned). Contact your Cisco Technical Representative if you need to reduce this setting.

The `/var/data/sessions.1` directory size of all sessionmgr VMs are 60% of actual allocated RAM size of that VM and this directory is mounted on tmpfs file system and used for session replica set. If you want to change `/var/data/sessions.1` directory size you must update (increase/decrease) the RAM size of that VM and re-deploy it.

For example, if 24 GB RAM is allocated to the Session Manager VM, 16 GB is allocated to `/var/data/sessions.1` directory on tmpfs.

If you need to update `sessions.1` directory settings consult your Cisco Technical Representative.

Considerations

- Each blade should have at least 2 CPU's reserved for the Hypervisor.
- When supported by the Hypervisor, deployments must enable CPU and memory reservation.
- For VMware environments, hardware must be ESX/ESXi compatible.

- The total number of VM CPU cores allocated should be 2 less than the total number of CPU cores per blade.
- CPU must be a high performance Intel x86 64-bit chipset.



Note BIOS settings should be set to high-performance values, rather than energy saving, hibernating, or speed stepping (contact hardware vendor for specific values).

- CPU benchmark of at least 13,000 rating per chip and 1,365 rating per thread.
- Monitor the CPU STEAL statistic. This statistic should not cross 2% for more than 1 minute^1.



Note A high CPU STEAL value indicates the application is waiting for CPU, and is usually the result of CPU over allocation or no CPU pinning. CPS performance cannot be guaranteed in an environment with high CPU STEAL.

- Scaling and higher performance can be achieved by adding more VM's, not by adding more system resources to VM's.
- For deployments which cannot scale by adding more VM's, Cisco will support the allocation of additional CPU's above the recommendation, but does not guarantee a linear performance by increasing more number of the VMs.
- Cisco will not support performance SLA's for CPS implementations with less than the recommended CPU allocation.
- Cisco will not support performance SLA's for CPS implementations with CPU over-allocation (assigning more vCPU than are available on the blade, or sharing CPU's).
- RAM latency should be lower than 15 ns.
- RAM should be error-correcting ECC memory.
- Disk storage performance needs to support less than 2ms average latency.
- Disk storage performance needs to support greater than 5000 input/output operations per second (IOPS) per CPS VM.
- Disk storage must provide redundancy and speed, such as RAID 0+1.
- Cisco does not validate its CPS solution on external storage (SAN storage, shared block storage, shared file systems).
- Hardware must support 1 Gbps ports/links for each VM network interface.
- Hardware and hardware design must be configured for better than 99.999% availability.
- For HA deployments, Cisco requires the customer designs comply with the Cisco CPS HA design guidelines, such as:
    - At least two of each CPS VM type (PD, PS, SM, CC) for each platform.

- Each CPS VM type (PD, PS, SM, CC) must not share common HW zone with the same CPS VM type.

- VMware memory (RAM) Reservation must be enabled at the maximum for each CPS VM (no over-subscription of RAM).

Deployment Examples

High Availability (HA) Deployment Example

Here are examples for high availability (HA) deployment:



Note The session replica-set for mongo port 27717 must always be built by using sessionmgr01 and sessionmgr02.

If you build session replica-set for mongo port 27717 with other session managers other than SM01 and SM02, the Policy Server (qns) process does not come up.

It is not recommended to use 2 or 4 blades layout for production.

Table 4: 2 Blade Setup with 16 CPU

Blade	VM Type	Replica-sets
1	CC 6, LB 8, QNS 8, SM 6, CM 2	SM: ADMIN, Balance, Session, SPR, Reporting
2	CC 6, LB 8, QNS 8, SM 6	SM: ADMIN, Balance, Session, SPR, Reporting

Table 5: 2 Blade Setup with 24 CPU

Blade	VM Type	Replica-sets
1	CC 6, LB 12, 2 x QNS 8, SM 8, CM 4	SM: ADMIN, Balance, Session, SPR, Reporting
2	CC 6, LB 12, 2 x QNS 8, SM 8	SM: ADMIN, Balance, Session, SPR, Reporting

Table 6: 4 Blade Setup with 16 CPU

Blade	VM Type	Replica-sets
1	CM 4, CC 8, LB 8, QNS 8	-
2	CC 8, LB 8, QNS 8	-
3	2 x QNS 8, SM 8	SM: ADMIN, Session RS1,2, Balance RS1, SPR RS1, Reporting RS1

Blade	VM Type	Replica-sets
4	2 x QNS 8, SM 8	SM: ADMIN, Session RS1,2, Balance RS1, SPR RS1, Reporting RS1

Table 7: 4 Blade Setup with 24 CPU

Blade	VM Type	Replica-sets
1	CM 4, CC 8, LB 8, QNS 10, SM 8, HSF 8	SM: ADMIN, Session (Backup), Balance (Backup), SPR
2	CC 8, LB 8, QNS 10, SM 8, HSF 8	SM: ADMIN, Session (Backup), Balance (Backup), SPR
3	3 x QNS 10, 2 x SM 8	SM: Session RS1,2, Balance RS1
4	3 x QNS 10, 2 x SM 8	SM: Session RS1,2, Balance RS1

Table 8: 8 Blade Setup with 16 CPU

Blade	VM Type	Replica-sets
1	CM 4, CC 6, LB 12, HSF 6	SM: ADMIN, Session (Backup), Balance (Backup)
2	CC 6, LB 12, HSF 6	SM: ADMIN, Session (Backup), Balance (Backup)
3	2 x QNS 12, SM 6	SM: Session RS1,2, Balance RS1
4	2 x QNS 12, SM 6	SM: Session RS2,1, Balance RS2
5	2 x QNS 12, SM 6	SM: Session RS3,4, SPR RS1
6	2 x QNS 12, SM 6	SM: Session RS4,3, SPR RS2
7	2 x QNS 12, SM 6	SM: Session RS5,6, Reporting RS1
8	2 x QNS 12, SM 6	SM: Session RS6,5, Reporting RS2

Table 9: 9 Blade Setup with 24 CPU

Blade	VM Type	Replica-sets
1	CC 12, 2 x LB 12, HSF 8	SM: ADMIN, Session (Backup), Balance (Backup)
2	CC 12, 2 x LB 12, HSF 8	SM: ADMIN, Session (Backup), Balance (Backup)

Blade	VM Type	Replica-sets
3	3 x QNS 10, 2 x SM 8	SM: Session RS1,2,7,8, Balance RS1
4	3 x QNS 10, 2 x SM 8	SM: Session RS2,1,8,7, Balance RS2
5	3 x QNS 10, 2 x SM 8	SM: Session RS3,4,9,10, SPR RS1
6	3 x QNS 10, 2 x SM 8	SM: Session RS4,3,10,9, SPR RS2
7	3 x QNS 10, 2 x SM 8	SM: Session RS5,6,11,12, Reporting RS1
8	3 x QNS 10, 2 x SM 8	SM: Session RS6,5,12,11, Reporting RS2
9	CM 4	CM: Cluster Manager

Platform WSP File Sizing Calculation

Note This section is for reference purposes only. For your deployment specific calculations, contact your Cisco Account representative.

For calculation purposes, consider there are 10 VMs in a standard deployment.

- pcrfclient (OAM)- 2
- Policy Server (qns) - 4
- Policy Director (lb) - 2
- Session Manager (sessionmgr) - 2

This table provides the statistics sizing details for different VM types in a standard deployment:

Note The size of 1 WSP file is 1.59 MB.

Table 10: Platform Statistics Sizing Details

Statistics Type	VM Details (size in MB for each VM)				Total Disk Usage (in MB)
	pcrfclient (OAM)	Policy Server (qns)	Policy Director (lb)	Session Manager	
cpu	76.32	76.32	101.76	76.32	814.08[12]
disk	76.32	76.32	73.14	73.14	750.48

Statistics Type	VM Details (size in MB for each VM)				Total Disk Usage (in MB)
	pcrfclient (OAM)	Policy Server (qns)	Policy Director (lb)	Session Manager	
memory	11.13	11.13	11.13	11.13	111.30
interface	38.16	25.44	38.16	25.44	305.28
fhcount	4.77	4.77	4.77	4.77	47.70
df	57.24	57.24	57.24	62.01	581.94
process	305.28	57.24	267.12	171.72	1717.20
set	1.59	0	0	0	3.18
collectd	4.77	0	0	0	9.54
swap	7.95	7.95	7.95	7.95	79.5
load	4.77	4.77	4.77	4.77	47.7
tcpconns	17.49	17.49	17.49	17.49	174.9
db	0	0	0	213.06	426.12
Total Size in MB	**605.79**	**338.67**	**583.53**	**667.8**	**5068.20**
Total Size in GB	**4.950117188**				

[1] CPU = 76.32 x 2 OAM + 76.32 x 4 qns + 101.76 x 2 lbs + 76.32 x 2 SM = disk Used (total)
[2] Similar calculations are applied for all the statistics.

Sample Customer Deployment

Let us consider customer has deployed 36 VMs.

- pcrfclient (OAM) - 2
- Policy Server (qns) - 12
- Policy Director (lb) - 2
- Session Manager (sessionmgr) - 12



Note The size of 1 WSP file is 1.59 MB.

Table 11: Platform Statistics Sizing Details - Customer Deployment

Statistics Type	VM Details (size in MB for each VM)				Total Disk Usage (in MB)
	pcrfclient (OAM)	Policy Server (qns)	Policy Director (lb)	Session Manager	
cpu	152	152	102	102	3052.8
disk	91	73	76	73	2121.06
memory	11	11	11	11	311.64
interface	38	115	25	25	915.84
fhcount	5	5	5	5	133.56
df	110	81	72	62	1984.32
process	496	267	57	114	3587.04
set	16	0	0	0	31.8
collectd	5	0	0	0	9.54
swap	8	8	8	8	222.6
load	5	5	5	5	133.56
tcpconns	18	18	18	18	489.72
db	0	0	0	145	1729.92
Total Size in MB	**955**	**735**	**379**	**568**	**14723.4**
Total Size in GB	**14.37832031**				

Application KPI Metrics Sizing Calculation



Note This section is for reference purposes only. For your deployment specific calculations, contact your Cisco Account representative.

For calculation purposes, Node2 to Node4 is the Diameter Endpoint. Number of nodes can be increased based on the Endpoints that are configured. The calculations are done based on the data received from the customer site. If any new interfaces such as, Sd are configured, the statistics generated will increase. This results in increase in the number of WSP files generated.

Here are the setup details:

- pcrfclient (OAM) - 2

- Policy Director (lb) - 2
- Policy Server (qns) - 4

Table 12: Values based on All Possible Conditions

Node/Each VM	Policy Director (lb)	Policy Server (qns)	Total
Node1	30	494	524
Node2	696	0	696
Node3	696	0	696
Node4	696	0	696
No. of WSP Files per VM	**2118**	**494**	**2612**
No. of VMs	2	4	6
Total No. of WSP Files on all VMs	**4236**	**1976**	**6212**

Table 13: Size of Disk Required in MB by one VM

	Policy Director (lb)	Policy Server (qns)	Total
All possible conditions	3367.62	785.46	4153.08

Table 14: Size of Disk Required in MB by All VMs

	Policy Director (lb)	Policy Server (qns)	Total
All possible conditions	6735.24	3141.84	9877.08

Sample Customer Deployment

The calculations done in this section are based on the data received from the customer site. Node2 to Node4 is the Diameter Endpoint. Number of nodes can be increased based on the Endpoints that have been configured. If any new interfaces such as, Sd are configured, the statistics generated will increase. This results in increase in number of WSP files generated.

Setup details:

- pcrfclient (OAM) - 2
- Policy Director (lb) - 2
- Policy Server (qns) - 12
- Session Manager (sessionmgr) - 12

Table 15: Values based on All Possible Conditions

Node/Each VM	Policy Director (lb)	Policy Server (qns)	Total
Node1	54	2998	3052
Node2	28267	0	28267
Node3	28267	0	28267
Node4	28267	0	28267
No. of WSP Files per VM	**84855**	**2998**	**87853**
No. of VMs	2	12	14
Total No. of WSP Files on all VMs	**169710**	**35976**	**205686**

Table 16: Number of wsp Files from Customer Site

Node/Each VM	Policy Director (lb)	Policy Server (qns)	Total
Node1	54	1627	1681
Node2	1143	0	1143
Node3	1143	0	1143
Node4	1143	0	1143
No. of WSP Files per VM	**3483**	**1627**	**5110**
No. of VMs	2	12	14
Total No. of WSP Files on all VMs	**6966**	**19524**	**26490**

Table 17: Size of Disk Required in MB by one VM

	Policy Director (lb)	Policy Server (qns)	Total
As per data gathered from customer site	5537.97	2586.93	8124.9
All possible conditions for Customer Site	134919.5	4766.82	139686.27

Table 18: Size of Disk Required in MB by All VMs as per Customer Site

	Policy Director (lb)	Policy Server (qns)	Total
As per data gathered from customer site	11075.94	31043.16	42119.1
All possible conditions for Customer Site	269838.9	57201.84	327040.74

Install and Configure VMware

Prior to installing CPS make sure you have the ESXi hosts details like, blade IP address, user name, password, datastore name, and network name.

Install VMware vSphere Hypervisor (ESXi)

VMware ESXi™ 7.0 until 7.0.3 (or) VMware ESXi™ 8.0 until 8.0.3 must be installed on all the blades that are hosting CPS. For more details see the VMware VSphere 7.0 and VMware VSphere 8.0.

You can install upgrade or migrate ESXi from a CD/DVD drive using a script that specifies the installation or upgrade options.

You can start the installation or upgrade script by entering a boot option when you start the host. You can also create an installer ISO image that includes the installation script. With an installer ISO image you can perform a scripted unattended installation when you boot the resulting installer ISO image.

Important User must use simple passwords (not containing special characters) during ESXi Installation. The CPS script uses this EXSi password to deploy the CPS VMs. Once the installation is complete, user can change the password to a more complex one.

In vSphere 7.0 until 7.0.3 (or) vSphere 8.0 until 8.0.3, the vSphere Web Client is installed as part of the vCenter Server on Windows or the vCenter Server Appliance deployment.

Prerequisites

- You must have the ESXi installer ISO in one of the following locations:
  - On CD or DVD. If you do not have the installation CD/DVD you can create one. Download and burn the ESXi Installer ISO Image to a CD or DVD.
  - On a USB flash drive.
- Verify that the server hardware clock is set to UTC. This setting is in the system BIOS.
- Consider disconnecting your network storage. This action decreases the time it takes the installer to search for available disk drives. Note that when you disconnect network storage any files on the disconnected disks are unavailable at installation.

- Do not disconnect a LUN that contains an existing ESX or ESXi installation. Do not disconnect a VMFS datastore that contains the Service Console of an existing ESX installation. These actions can affect the outcome of the installation.
- Gather the information required by the ESXi installation wizard.
- Verify that ESXi Embedded is not present on the host machine. ESXi Installable and ESXi Embedded cannot exist on the same host.

Installation

For more information related to ESXi installation, refer to https://www.vmware.com/products/vsphere-hypervisor.html.

Procedure

Step 1 Download the ESXi installable ISO file.

Step 2 Mount the ISO file to a CD and feed the CD to the server where you want to install ESXi to.

Step 3 After you boot from the CD, the installer loads. Press Enter to begin and then F11 to accept the licensing agreement. Next, choose a disk to install to (All data will be erased). After ejecting the install CD, press **F11** to start the installation.

Step 4 After the installation is completed, press **Enter** to reboot, and ESXi starts.

What to do next

Open a Web browser and enter the URL for the vSphere Web Client: https://*vcenter_server_ip_address*.

If a warning message about an untrusted SSL certificate appears, select the appropriate action based on your security policy.

Note After you complete installation, IPv6 must be enabled on each blade. For more information on enabling IPv6, refer to IPv6 Support - VMware, on page 118.

Enable SSH

CPS software installation requires SSH to be enabled for each blade server host.

After you complete the installation and configuration of CPS, you can disable SSH for security purposes.

To enable SSH, perform the given steps:

Procedure

Step 1 Login to the vSphere Web Client.

Step 2 Select the host by IP address or name in the left panel.

Step 3 Click **Configure** tab from the top menu from the right panel.

Step 4 Under **System**, click **Security Profile** from the options available.

Step 5 Click **Edit...** in the upper right corner of the **Firewall** panel.

The **Edit Security Profile** window opens.

Step 6 Check **SSH Server** and configure the required port and protocol. Click **OK**.

Note
By default, daemons will start automatically when any of their ports are opened, and stopped when all of their ports are closed.

Configure VMware ESXi Timekeeping

Both the VMware ESXi and Load Balancers time must be configured correctly. Other VMs in CPS use Load Balancers as the NTP source.

To configure VMware ESXi Timekeeping, you must coordinate with customers or gain access to their NTP servers.

Login as an administrator to every VMWare ESXi host to be used for the deployment using the VMware vSphere client.

For each host, perform the given steps:

Procedure

Step 1 Click the host (IP address or name) in the left column.

Step 2 Click **Configure** tab from the top menu from the right panel.

Step 3 Under **System**, click **Time Configuration** from the options available.

Step 4 Click **Edit...** in the upper right corner of the **Time Configuration** panel.

The **Edit Time Configuration** window opens.

Step 5 Check Use Network Time Protocol (Enable NTP Client). The following parameter can be set:

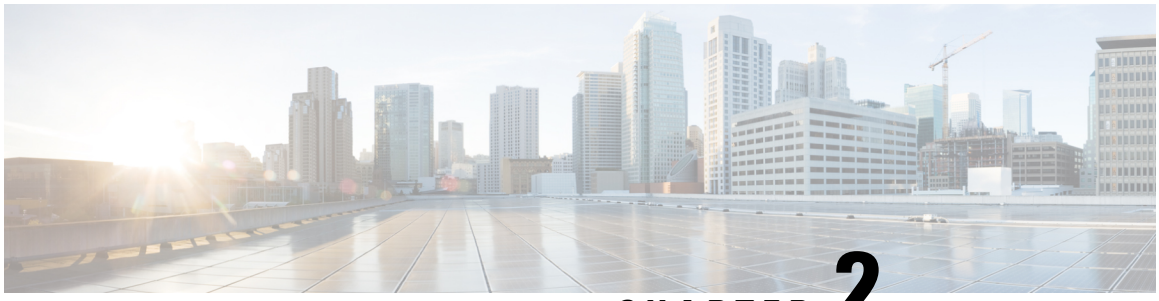
a) NTP Service Status: Options are Start, Stop and Restart. The NTP Service settings are updated when you click Start, Restart or Stop.
b) NTP Server Startup Policy: Options are Start and stop with host, Start and stop with port usage, Start and stop manually.
c) STP Servers: Add NTP Server given by or coordinated with the customer.

Step 6 After configuring the parameters according to your requirement click **OK**.
Date and Time should now show correctly in the Time Configuration window in vSphere Client. Date and Time displayed in red color indicates NTP skew that should be resolved.

Collect Virtualization Parameters for CPS Installation

Before starting the CPS deployment prepare and make sure the given items are available:

- The traffic analysis for the capacity needed for this deployment.
- Number of VMs (the type of VMs such as Policy Director (LB), OAM (PCRFCLIENT), sessionmgr, Policy Server (QNS), node).
- The size of VMs, for each type of VMs, the size of disk memory CPU etc.
- The number of blades.
- The number of networks that the deployment will be deployed to.



CHAPTER 2

CPS Installation

- Obtain the CPS Software, on page 19
- Cluster Manager VM, on page 20
- Configure System Parameters for Deployment, on page 30
- Import the Excel Information into the Cluster Manager VM, on page 78
- Customize Features in the Deployment, on page 81
- License Generation and Installation, on page 84
- SSL Certificates, on page 88
- Enable Custom Puppet to Configure Deployment, on page 90
- Installing Platform Scripts for MongoDB Health Monitoring - VMware, on page 92

Obtain the CPS Software

Obtain the CPS software from the download link provided in the CPS Release Notes for this release.

The CPS software distribution includes the following files:

- The `CPS_x.x.x.release.iso` file which serves as a temporary virtual CD driver containing the installation software.
- A compressed tar file that contains a `base.vmdk` which serves as the virtual hard drive in building the Cluster Manager virtual machine (VM).
- An Excel spreadsheet included in the `*.iso` which you manually edit to contain the IP addresses, virtual topology, and cluster settings for a High Availability (HA) deployment.

  Instructions are provided later in this document on how to obtain this Excel spreadsheet.

A VMware OVF tool is also needed to install CPS. This utility can be downloaded as described later in this guide.

Cluster Manager VM

Overview

Cluster Manager is the main virtual machine that manages the deployment, installation, upgrade, configuration and patching of the CPS cluster. The Cluster Manager stages artifacts such as Configuration, Puppet scripts, Shell script tools and CPS application software. The artifacts are applied to the CPS virtual machines (VMs) during initial CPS installation, CPS upgrades, and application of patches to the CPS.

There are four categories of artifacts:

- **Cluster Deployment Configuration**

  All the cluster deployment configuration files used for full deployment as well as individual VM deployment are stored in `/var/qps/config/deploy`. These files are created by exporting the CSV files from the CPS Deployment Template Excel spreadsheet and contains the cluster deployment configuration. For more information related to deployment template and CSV files, refer to the section Configure System Parameters for Deployment, on page 30.

  These configuration files are used by the deployment scripts (`deploy.sh` and `deploy_all.py`) during VM deployment.

- **CPS Software Configuration**

  All the CPS software configuration files which includes the configuration files in `/etc/broadhop` such as features file, qns.conf, jvm.conf and policy files (such as charging rules policy) are stored in `/var/qps/current_config/`. These configurations are applied to CPS VMs after CPS software is installed. The configuration files are copied to Cluster Manager's `/var/www/html` directory. After a VM is deployed, the puppet script in the VM downloads the configuration files and applies the configuration to the CPS software in the VM.

  The iomanager configuration file (`/etc/broadhop/iomanager/qns.conf`) is controlled by puppet. So in case you want to modify iomanager configuration file, you must modify `/etc/puppet/modules/qps/templates/etc/broadhop/iomanager/qns.conf.erb` file.



Note When you are upgrading/migrating from one release to another, you need to modify the iomanager configuration files again with the changes.

- **Puppet**

  Puppet (http://puppetlabs.com/) is the tool utilized for installing, deploying, and upgrading cluster virtual machines and configurations. Refer to Puppet Overview, on page 21 for more information.

- **Tools**

  - Various tools used for operation and maintenance in Cluster Manager.

    `/var/qps/bin` -> `/var/qps/install/current/scripts/bin` (-> is a Linux softlink)

  - Deployment Scripts: Scripts used for VM deployment.

- Build Scripts: Scripts that are used to tar the configuration, puppet scripts and software into the `/var/www/html` directory on the Cluster Manager for download by each VM during deployment.
- Control Scripts: Scripts that are used on Cluster Manager to perform tasks such as start/stop of the CPS processes running on the VM nodes.

Directory Structure

- All the artifacts for a release are stored in:

  `/var/qps/install/current` -> `/var/qps/install/`*`CurrentRelease`* (-> is a Linux softlink)
- Tools: `/var/qps/bin` -> `/var/qps/install/current/scripts/bin` (-> is a Linux softlink)
- Deployment scripts are used to deploy VMs.
- Build scripts that zips the configuration, puppet and CPS software to `/var/www/html` directory in Cluster Manager.
- Control scripts
- Configurations includes the configuration files in `/etc/broadhop` such as features file, `qns.conf`, `jvm,conf` and policy files. All the configurations in this directory are pushed to the VMs during deployment.
- Files unchanged after upgrade: All the files in `/etc/broadhop` after upgrade remain unchanged.

Puppet Overview

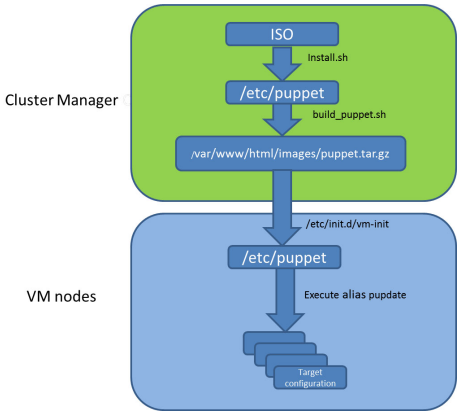
Puppet (http://puppetlabs.com/) is a tool utilized for installing, deploying, and upgrading CPS virtual machines and configurations.

Puppet operations are initiated automatically when CPS installation or upgrade scripts are run. These scripts in turn utilize numerous utility scripts to configure system modules.

For example: `reinit.sh` (used for upgrades) triggers `/etc/init.d/vm-init`.

1. When the Cluster Manager VM is deployed, puppet scripts are copied from the CPS ISO to `/etc/puppet`.
2. The `build_puppet.sh` moves them to `/var/www/html/images/puppet.tar.gz`.
3. `vm-init` downloads the `puppet.tar.gz` from cluster manager and populates them to the `/etc/puppet` directory in the VM nodes.

Figure 1: Installation Flow



Many CPS modules are managed by Puppet, including: java, ntp, zero mq, haproxy, mongo, socat, memcache, diameter, elasticsearch, monit, iomanager, unifiedapi, license manager, policybuilder, collectd, logserver, snmp, grafana.

Puppet files are stored centrally in `/etc/puppet` on the Cluster Manager.

CPS VM nodes and their software and configurations are staged in the `/var/www/html/` directory in zip files. When a VM is rebooted, the VM downloads and runs the appropriate puppet scripts to update the configuration.

Once puppet files are downloaded to each VM node, they reside in `/etc/puppet/` directory on the each VM node.

- `/etc/puppet/puppet.conf`: Basic configuration of puppet.
- `/etc/puppet/classfyNode.sh`: Determines the node type and the appropriate puppet script from `/etc/broadhop.profile`.
- `/etc/puppet/modules/qps/manifests/roles/*.pp`: These are the corresponding scripts for a node to run. For example: pcrfclient01.pp, pcrfclient02.pp, lb01.pp, qns.pp, sessionmgr.pp, etc.
- `/etc/puppet/modules/`: Contains all the puppet code.
- `env_config` -> `/var/qps/env_config`: Contains custom puppet files.

Puppet scripts can be started manually using the `pupdate` command, however this should be reserved for troubleshooting, reconfiguration, or recovery of CPS systems with assistance from a Cisco representative.

Modification or execution of puppet scripts should only be performed under the direction of a Cisco Advanced Services representative. Puppet scripts require root level permissions to be modified.

Additional information about custom deployments is provided in Enable Custom Puppet to Configure Deployment, on page 90.

For more information about Puppet, refer also to the Puppet documentation available at: https://docs.puppetlabs.com/puppet/.

Deploy the Cluster Manager VM

The Cluster Manager is a server that maintains the system (Operating System) and application artifacts such as software and CPS and Linux configuration for the CPS cluster. It also is responsible for deploying, installing/upgrading the software for the Virtual Machines in the CPS cluster.

Important User must use standard VMware Switch during VM deployment and avoid using distributed switches. If distributed switches are really needed, initial deployment should be made using standard Switch and post deployment user can change the switch type to distributed.

To deploy the cluster manager VM, perform the following steps:

Procedure

Step 1 Login to the vSphere Web Client and select the blade where you want to create a new VM to install the cluster manager VM.

Step 2 Right-click on the blade and select **New Virtual Machine**. **New Virtual Machine** window opens up.

Step 3 Select **Create a new virtual machine** and click **Next** to open **Select a name and folder**.

Step 4 Enter a name for the virtual machine (for example, CPS Cluster Manager) and select the location for the virtual machine. Click **Next**.

Step 5 Select blade IP address from **Select a compute resource** window and click **Next** to open **Select storage** window.

Step 6 From **Select storage** window, select *datastorename* and click **Next** to open **Select compatibility** window.

Step 7 From **Compatible with:** drop-down list, select **`ESXi 7.0.3/8.03 and later`** and click **Next** to open **Select a guest OS** window.

Note
Support for VMX11 is added only for fresh install. For upgrade flow (option 2/option 3), upgrade of VMX is not supported.

Step 8 From **Guest OS Family:** drop-down list, select **`Linux`** and from **Guest OS Version:** drop-down list, select **`CentOS 4/5 or later (64-bit)`**.

Step 9 Click **Next** to open **Customize hardware** window.

Step 10 In **Virtual Hardware** tab:

a) Expand **CPU** node and select **CPU** and **Cores per Socket** as given in Virtual Machine Requirements, on page 3.

b) Select **Memory** size as **`12 GB`**.

c) Expand **New SCSI controller** and from **Change Type** drop-down list, select **VMware Paravirtual**.

d) 2 NICs are required (one for eth1 as internal and second for eth2 as management). One NIC already exists as default under **New Network**.

Under **New Network**, check **Connect At Power On** is selected.

e) To add another NIC, click **ADD NEW DEVICE** and from the list select **Network Adapter**.

Under **New Network**, check **Connect At Power On** is selected.

f) Click **Next** to open **Ready to complete** window.

Step 11 Review the settings displayed on **Ready to complete** window and click **Finish**.

Step 12 Login to EXSi blade through the vSphere client software.

Step 13 Select the ESXi host (not the new VM they just created) and select **Summary** tab from the right pane.

Step 14 Under **Storage**, select the *datastorename*.

Step 15 Right-click on the *datastorename* and select **Browse Datastore…** to open **Database Browser** window..

Step 16 To upload the CPS software to the datastore, select the new directory created for your VM, and click  (Upload a file to this datastore) button and select **Upload File...**.

Step 17 Navigate to the location of the *CPS_*.tar.gz* file which you downloaded earlier. Select it and click **Open**.

Step 18 To upload CPS ISO, repeat Step 16, on page 24.

Step 19 Navigate to the location of the *CPS_*.release.iso*, select it and click **Open**.

Step 20 To upload base.vmdk, repeat Step 16, on page 24. Navigate to the location of the *Base*.vmdk_signed*/*Base*.vmdk*, select it and click **Open**

Step 21 To upload *Cluman seed ISO image*, repeat Step 16, on page 24. Navigate to the location of the *Base*.vmdk_signed*/*cluman_seed.iso*, select it and click **Open**.

Step 22 Open a secure shell (ssh) connection to the blade ESXi host.

Step 23 cd to the directory with the data store.

```
cd /vmfs/volumes/<datastore_name>/<foldername>
```

For example:

```
cd /vmfs/volumes/datastore5/CPS Cluster Manager
```

Step 24 Convert the vmdk file to ESX format:

```
vmkfstools --diskformat thin -i Base*.vmdk newbase.vmdk
```

Note
This command can take several minutes to complete.

Step 25 Cluster Manager Memory Reservation: Press Ctrl + Alt +2 to go back to **Hosts and Clusters** and select the VM created above (*CPS Cluster Manager*).

a) Right-click and select **Edit Settings..**. **Virtual Hardware** tab is displayed as default.
b) Click on **Resources** tab and select **Memory** settings.
c) From the right-side options, select **Reserve all guest memory (All Locked)**.
d) Click **OK** to save the changes.

Step 26 Press **Ctrl + Alt +2** to go back to **Hosts and Clusters** and select the VM created above (*CPS Cluster Manager*).

a) Right-click and select **Edit Settings...**. **Virtual Hardware** tab is displayed as default.
b) Click **ADD NEW DEVICE** and from the list select **Existing Hard Disk** to open **Select File** window.
c) Navigate to the location of your new VM and select *newbase.vmdk* (created in Step 24, on page 24) and click **OK**.
d) Expand **New Hard disk** and select **Virtual Device Node** as **`SCSI Controller 0`** from the drop-down list.
e) Click **SCSI Controller 0** and from **Change Type** drop-down list, select **VMware Paravirtual** and click **OK**.

Step 27 Mount Cluster Manager seed ISO on CD/DVD:

a) From **New device** drop-down list, select **`CD/DVD Drive`** and click **Add**. The newly added New CD/DVD Drive will appear at the end of the window.
b) Change the **New CD/DVD Drive** option from **Client Device** to **Datastore ISO File**. Browse to required *Cluman seed ISO image* (For example, base/cluman_seed.iso), select it and click **OK**.

c) Check **Connect At Power On** to connect the device when the virtual machine turns on.
d) Select **Virtual Device Node** as **`IDE(1:0)`** from the drop-down list and click **OK**.

Important
If the selected **Virtual Device Node** is busy, select any alternate node (IDE) from the drop-down list.

Step 28 Mount ISO on CD/DVD:

a) Expand **CD/DVD drive 1** and change the option from **Client Device** to **Datastore ISO File**. Browse to the ISO image file, select the required ISO image and click **OK**.
b) Check **Connect At Power On** to connect the device when the virtual machine turns on and click **OK**.
c) Select **Virtual Device Node** as **`IDE(1:1)`** from the drop-down list and click **OK**.

Important
If the selected **Virtual Device Node** is busy, select any alternate node (IDE) from the drop-down list.

Step 29 Select **VM Options** tab and expand the **BOOT Options**. Ensure **Firmware** is selected as **`BIOS`**.

Step 30 Power on the *CPS Cluster Manager* VM.

Note
The following message may be reported on the Cluster Manager VM console. You can disregard this message.

Probing EDD (edd=off to disable)

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

Configure Cluster Manager VM

To configure cluster manager VM, perform the following steps:

Common Steps

Procedure

Step 1 Login to the vSphere Web Client.

Step 2 To open VM console, you have two options:

Option	Description
1	You can launch the console by selecting the Cluster Manager VM, right-click on VM and select **Open Console**.
2	Select the Cluster Manager VM from the left panel and click **Summary** tab from the top menu on the right panel. For 7.0.3 / 8.03 version: Click on the small gear icon and select **Launch Web Console** to open the VM console.

Step 3 Login to the VM as the root user. The default password is **CpS!^246**.

Step 4 Configure the network settings:

a) Private LAN for future VMs (a private sub network).

For example, `/etc/sysconfig/network-scripts/ifcfg-eth0` as:

This is specific to VMware deployments:

```
DEVICE=eth0
TYPE=Ethernet
#ONBOOT=yes
NM_CONTROLLED=no
IPADDR=XX.XX.XX.XX
NETMASK=XX.XX.XX.XX
```

b) Public address (access the cisco network).

For example, `/etc/sysconfig/network-scripts/ifcfg-eth1` as:

This is specific to VMware deployments:

```
DEVICE=eth1
TYPE=Ethernet
#ONBOOT=yes
NM_CONTROLLED=no
IPADDR=XX.XX.XX.XX
NETMASK=XX.XX.XX.XX
GATEWAY=XX.XX.XX.XX
```

Step 5 Restart the network.

`service network restart` or `/etc/init.d/network restart`

Step 6 Login to the CPS Cluster Manager VM as a **root** user using SSH and public address (or via the console).

Step 7 Edit/add the eth0 private IP address of the Cluster Manager in `/etc/hosts`.

For example:

```
XX.XX.XX.XX installer
```

Note
If the actual hostname for Cluster Manager VM is other than 'installer', then modify installer/cluman entry in `/etc/hosts` accordingly.

Example:

```
XX.XX.XX.XX installer <actual-hostname>
```

Step 8 Mount the ISO from CD/DVD:

```
mkdir -p /mnt/iso
```

```
mount -o loop /dev/sr0 /mnt/iso/
```

Note
Verify whether `install.sh` command is available in `/mnt/iso`. If `install.sh` command is not available, perform the following:

a) Unmount the CPS ISO:

```
umount /mnt/iso
```

b) Mount the ISO from CD/DVD:

```
mount -o loop /dev/sr1 /mnt/iso/
```

Step 9 Proceed to the next sections to continue the installation for a High Availability (HA) deployment.

HA Installation

To proceed with a new High Availability (HA) installation:

Procedure

Step 1 Run the **install.sh** script from the ISO directory.

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

Note
The default root password created during `install.sh` is not in compliance with the PSB requirements. Hence, it is recommended to change the default root password post completion of CPS deployment using `change_passwd.sh` script. For more information, refer to *Update Default Credentials* section in the *CPS Installation Guide for VMware*.

Step 2 When prompted for the install type, enter the required type based on your CPS deployment requirements.

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

Enter `mobile` to install Diameter, mog to install `mog` module, pats to install `pats`, `arbiter` to install Arbiter.

Important
- For more information on Arbiter installation, refer to *Standalone Arbiter Deployment on VMware* section in *CPS Geographic Redundancy Guide*.
- For more information on MOG/PATS, contact your Cisco Technical Representative.

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

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

Step 4 When prompted for the type of installation, enter *1* (New Deployment).

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

Note
Refer to the *CPS Migration and Upgrade Guide* for detailed instructions on option 2 and 3.

Step 5 When prompted to change the Cluster Manager default root password, enter the new password.

```
Need to change the default root password for security reasons..
Changing password for user root.
New password: XXXXX
Retype new password:
```

Note

You can create passphrase or password with the following limitations, when you create or change passwords:

- You can provide or update a password of a minimum length of 4 characters where it must consist of all 4 classes (1 capital letter, 1 small letter, 1 numeric and 1 special character).
- You can provide or update a password of length of 5 or more where it must consist of 3-4 classes (1 capital letter, 1 small letter, 1 numeric and 1 special character).
- You can provide or update a passphrase of 127 characters.

Step 6 After finishing the installation (or upgrade) process, unmount the ISO image using the following commands. This prevents any "device is busy" errors when a subsequent upgrade/new installation is performed.

```
cd /root
umount /mnt/iso
```

Note
If you are not able to unmount the ISO using `umount` command, then use `umount -l`.

Step 7 (Optional) After unmounting the ISO, delete the ISO image to free system space.

```
rm -rf /dev/sr0/xxxx.iso
```

where, *xxxx.iso* is the name of the ISO image.

Step 8 (Optional) Change the host name of the Cluster Manager.

a) Run `hostname xxx`, where *xxx* is the new host name for the Cluster Manager.
b) Edit `/etc/hostname` to add the new host name for the Cluster Manager.

Install VMware OVF tool 4.6.3

This procedure outlines the steps to install VMware OVF tool version 4.6.3 on the Cluman VM.

Before you begin

The prerequistes are:

- Ensure you have a Cluman VM.
- Verify that you have root access to the Cluman VM.

Note This procedure is specifically for the Linux 64-bit version. The ZIP archive installation method is an alternative to the bundle installers used in older version.

Procedure

Step 1 Download the VMware OVF tool:

a) Download version 4.6.3 for VMware 7.0/8.0 from the following URL. Choose the specific file : VMware-ovftool-4.6.3-24031167-lin.x86_64.zip, which is available under "OVF Tool for Linux Zip".

Step 2 Create a directory to store the OVF tool files:

```
mkdir /usr/lib/vmware-ovftool/
```

Step 3 Copy the downloaded ZIP file to the newly created directory:

```
cp VMware-ovftool-4.6.3-24031167-lin.x86_64.zip /usr/lib/vmware-ovftool/
```

Step 4 Unzip the file to extract the content of the Zip file:

```
cd /usr/lib/vmware-ovftool/
unzip VMware-ovftool-4.6.3-24031167-lin.x86_64.zip
```

This will create an `ovftool` directory containing the necessary libraries and executables.

Step 5 Create a symbolic link to the `ovftool` executable in `/usr/bin/` to make it accessible system-wide:

```
ln -s /usr/lib/vmware-ovftool/ovftool/ovftool /usr/bin/ovftool
```

Step 6 Verify the installation. Check the OVF tool version to confirm the installation:

```
[root@localhost ~]# ovftool --version
VMware ovftool 4.6.3 (build-24031167)
```

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.

Change Password

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.

For more information, refer to *Update Default Credentials*.

Note 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.

Configure System Parameters for Deployment

Note This section applies only for High Availability CPS deployments.

The following section guides you through the steps needed to properly configure a new installation of CPS. The Deployment Template file is a spreadsheet used for populating deployment parameters.

This file is available on the Cluster Manager VM at the following location:

`/var/qps/install/current/scripts/deployer/templates/QPS_deployment_config_template.xlsm`

After entering your parameters into the spreadsheet (as described in the following sections), the information from the spreadsheet is loaded onto the Cluster Manager VM. The Cluster Manager uses the information to configure the other CPS VMs in the cluster.

Note All alphabet characters used in virtual IPv6 addresses configured in csv files must be in small case letters.

To add values to the corresponding sheets in the template file, refer to the following sections:

Definitions Configuration

The **Definitions** sheet defines default parameters used by other sheets.

Select the **Definitions** sheet.

Figure 2: Definitions

	A	B	C	D	E
1	**Diskmode**	**Datastores**	**Alias**		
2	thin	datastore1	lb01		
3	monolithicSparse	datastore2	lb02		
4	monolithicFlat	datastore3	pcrfclient01		
5	twoGbMaxExtentSparse	datastore4	pcrfclient02		
6	woGbMaxExtentFlat	datastore5	portal01		
7	seSparse	datastore6	portal02		
8	eagerZeroedThick		sessionmgr01		
9	thick		sessionmgr02		
10	sparse		sessionmgr03		
11			sessionmgr04		
12			sessionmgr05		
13			sessionmgr06		
14			sessionmgr07		
15			sessionmgr08		
16			sessionmgr09		
17			sessionmgr10		

The following parameters can be configured in this sheet:

Table 19: Definitions Configuration Sheet Parameters

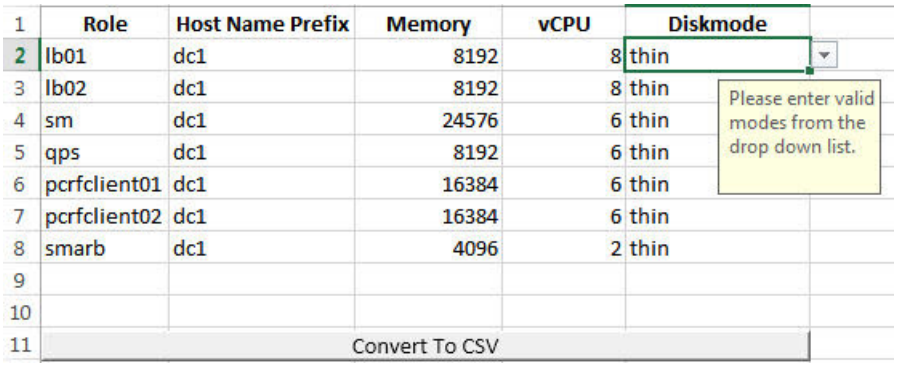
Parameter	Description
Diskmode	Do not modify this column. The Diskmode column defines the disk mode for VMware. This is used by the VMSpecification sheet.
Datastores	The Datastore column defines all the storages in the virtualization environment. It is used by the datastore column in the Hosts sheet. Add an entry here for each datastore in the virtualization environment. The datastore name must not contain spaces.
Alias	Be cautious modifying the values of the column. Add new names only if the number of session manager node names exceed 20, Policy Server (QNS) node names exceed 20. Use the naming convention: • For Policy Server (QNS) nodes: qnsxxx • For session manager: sessionmgrxx

VMSpecifications Configuration

In a CPS cluster, there are few types of nodes: Policy Director (LB), sessionmgr, Policy Server (QNS), and OAM (PCRFCLIENT). Each VM is assigned with a particular type of node. The following sheet defines the attributes for each type of node:

Select the **VMSpecification** sheet.

Figure 3: VM Specifications Configuration Sheet



The following parameters can be configured in this sheet:

Table 20: VMSpecification Configuration Parameters

Parameter	Description
Role	Do not change the value in this column. The Role column defines different types of VMs: lb01, lb02, sm, qps, pcrfclient01, pcrfclient02.
Host Name Prefix	The Host Name Prefix is prepended to the Guest Name (the host name of the VM in the Hosts sheet), which is used as the VM name in the ESX server, i.e dc1-sessionmgr01 is the VM name in vCenter and sessionmgr01 is the host name in the VM's Linux OS.
Memory	The Memory column is the size of memory needed for the type of the VMs in Megabytes (MB).
vCPU	The vCPU column is the number of CPU needed for the VM.
Diskmode	The Diskmode is how the Hypervisor should keep the disk of the VM in the storage. See VMware documentation for the meaning of different modes. Our recommendation is to keep it as thin mode unless specific needs arise in your Hypervisor environment.

Note Reserving Memory on the Virtual Machines (VMs):

To avoid performance impact, CPS reserves all the allocated memory to each CPS virtual machine. It is recommended to allocate 8 GB memory for the Hypervisor. For example, if the total memory allocated on a blade/ESXi host is 48 GB then you should only allocate 40 GB to CPS VMs and keep 8 GB for the Hypervisor.

VLANs Configuration

The VLAN Configuration sheet defines different subnets in the virtual infrastructure.

Select the **VLANs** sheet.

Figure 4: VLANS Configuration

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

Contact your Cisco Technical Representative for further information on VLANs.

The following parameters can be configured in this sheet:

Table 21: VLANs Configuration Parameters

Parameter	Description
VLAN Name	The VLAN Name column defines the name for a particular VLAN. It is recommended to use a name representing the network for certain traffic. For additional networks, add more as needed. The "Internal" VLAN Name is always needed. Names must consist only of alphanumeric characters and underscores, and must not start with a number.
Network Target Name	The Network Target Name column is the name of the networks defined in the Hypervisor (VMware), for example the network in vSphere for a blade server.
Netmask	The Netmask column is the network mask for the network. If the VLAN supports IPv6, the network mask can be IPv6 mask. If the VLAN interface supports both IPv4 and IPv6, add both netmasks in the cell, separated by space.
Gateway	The Gateway column is the gateway for the network, If the VLAN supports IPv6, the gateway can be IPv6 gateway address. If the VLAN interface supports both IPv4 and IPv6, add both gateways in the cell, separated by space. An example is provided in the table.
VIP Alias	Enter the alias name for the virtual interfaces in Policy Director (lb). The virtual addresses are used to distribute the traffic between two Policy Directors (LBs).

Parameter	Description
Pcrfclient VIP Alias	Enter the alias name for the virtual interfaces between OAM (PCRFCLIENTS) whenever you want VIP between pcrfclient01 and pcrfclient02 (for example, lbvip02 is VIP between lb01 and lb02). This virtual IP is used to support redundancy for arbiter member of replica set.
guestNic	This field is optional and it supports custom NIC/interface name other than default one i.e. eth0/1/2, which can support SR-IOV enabled interfaces. If guestNic field is empty, it takes the value eth0, eth1, eth2 in order of its appearance. For more information on bonding configuration, see guestNic, on page 34.

Table 22: Example

VLAN Name	Network Target Name	Netmask	Gateway	VIP Alias	Pcrfclient VIP Alias
Internal	VLAN_2017	255.255.255.0	NA	lbvip02	arbitervip
Management	VLAN_2025	255.255.255.0	172.20.25.1	lbvip01	-
Gx	VLAN_3041	64	2003:3041::22:1	lbvip03	-
Rx	VLAN_3043	64	2003:3043::22:1	lbvip05	-
Syp	VLAN_3042	64	2003:3042::22:1	lbvip04	-

guestNic

This field is optional and it supports custom NIC/interface name other than default one i.e. eth0/1/2, which can support SR-IOV enabled interfaces. If guestNic field is empty, it takes the value eth0, eth1, eth2 in order of its appearance.

The following table provides an example if bond interface is created for management and Gx network:

Table 23: VLANs

VLAN Name	Network Target Name	Netmask	Gateway	VIP Alias	guestNic
Internal	Sriov_network1	255.255.255.0	NA	lbvip02	eth0
Management	Sriov_network2	255.255.255.0	NA	lbvip01	bond03168
Gx	Sriov_network3	255.255.255.0	NA	lbvip03	bond01004

In the above example, lbvip01 is created on top of bond03168 and lbvip03 on bond01004.

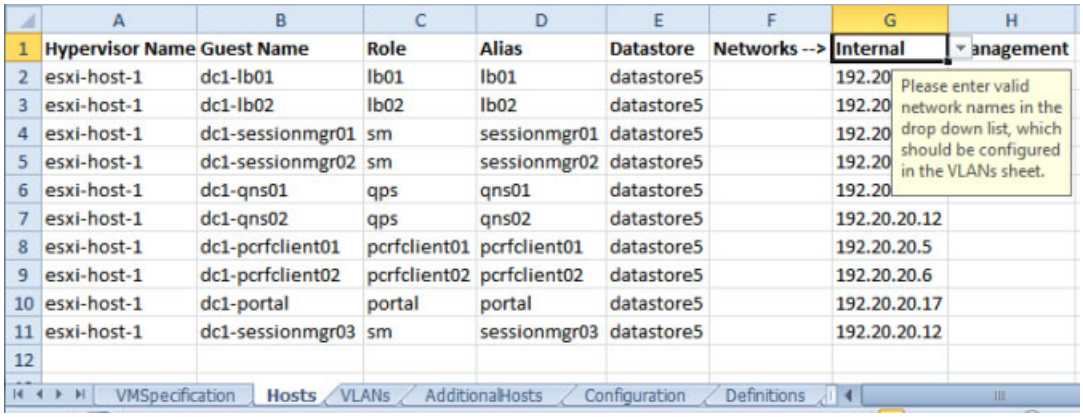
Hosts Configuration

In this sheet, all the VM/nodes are defined. The deployment uses the information here to deploy the VMs.

Note The host addresses used in the examples may be different from those in your deployment.

Select the **Hosts** sheet.

Figure 5: Hosts Configuration



The following parameters can be configured in this sheet:

Table 24: Hosts Configuration Parameters

Parameter	Description
Hypervisor Name	The Hypervisor Name column specifies the host names for the blade servers. The names should be routable by the Cluster Manager VM.
Guest Name	The Guest Name column is the host name of the VMs resolvable in the enterprise DNS environment. **Note** Host name is a text string up to 24 characters and can include alphabets, digits (0-9), minus sign (-), and period (.). The first letter of the host name can be either a letter or a digit. For more information on host names, refer to the following links: https://tools.ietf.org/html/rfc952 https://tools.ietf.org/html/rfc1123

Parameter	Description
Role	The role defines the type of VM within the CPS cluster. The Role column is a drop-down entry from a list specified in VMSpecification sheet. • lb01, lb02: Policy Director • pcrfclient01, pcrfclient02: OAM • qps: Policy Server • sm: Session Manager
Alias	The Alias is the internal host name used by CPS nodes for internal communication, such as qns01.
Datastore	The Datastore column is the datastore name used by the Hypervisor for the physical storage of the VM. The datastore is a drop-down list from column data in the Definition sheet.
Networks -->	The Networks --> column is a read only column. Do not write anything to it.
Internal/Management	The columns following the Networks --> specifies all the IP addresses for the VMs. For each VLAN Name in the VLANS sheet for the VM, a new column should be added for that network. The title of the column should come from the VLAN name in the VLANS sheet. The content should be the IP address. If the network is IPv6, add IP v6 address. If the interface has both IPv4 and IPv6 addresses, add both addresses in the cell, separated by space. The “Internal” network name is reserved and should always be present. The IP address for the internal network can only be either IPv4 or IPv6, but not both. **Note** Use the uncompressed IPv6 address. For example: 2345:f170:8306:8118:e0:208:0:100



Important Verify that all VM IP addresses and host names (Guest Name) are configured properly in the Hosts sheet. You cannot modify the IP addresses or host names manually on the VMs (excluding Cluster Manager) after deploying the VMs. Instead, you must correct the IP addresses and host names in the Hosts sheet, then import the file to the Cluster Manager and re-deploy the VMs with the updated IP address or host names.

Additional Hosts Configuration

There are many hosts in the environment that CPS needs to interact with, for example: NTP server, NMS server, etc. The AdditionalHosts sheet contains all these hosts and IP addresses. The host entries are copied to the `/etc/hosts` file of the Cluster Manager during the deployment.

Note Each line in the `/etc/hosts` file must start with an IP Address.

For additional information about `/etc/hosts`, refer to http://unixhelp.ed.ac.uk/CGI/man-cgi?hosts.

Select the **AdditionalHosts** sheet.

Figure 6: Additional Hosts

	Host	Alias	IP Address
2	ntp-primary	ntp	155.165.201.253
3	ntp-secondary	btp	155.165.132.253
4	lbvip01	lbvip01	10.105.94.232
5	lbvip02	lbvip02	192.20.20.27
6	snmp-trapdest	nms-destination	155.174.11.118
7	esxi-host-1	esxi-host-1	10.105.93.226
8	esxi-host-2	esxi-host-2	10.105.93.227
9	esxi-host-3	esxi-host-3	10.105.93.228
10	esxi-host-4	esxi-host-4	10.105.93.229
11	corporate_nms_ip	nms_manager	155.174.11.118
12			
13			

... AdditionalHosts Conf ...

The following parameters can be configured in this sheet:

Table 25: Additional Hosts Configuration Parameters

Parameter	Description
Host	The Host column is the arbitrary value that can be added by user as the name of the virtual machines added to the Hypervisor. **Attention** Make sure lbvip01, lbvip02 and sslvip01 host values are not changed from their default values. By default, the values for lbvip01, lbvip02 and sslvip01 are lbvip01, lbvip02 and sslvip01 respectively.
Alias	The Alias is the internal host name used by CPS nodes for internal communication, such as qns01.

Parameter	Description
IP Address	IP address of the host. Currently, IPv6 is supported only for policy director (lb) external interfaces. An example is provided in the Table 26: Example, on page 38. **Note** For IPv6, use the uncompressed IPv6 address. For example: 2345:f170:8306:8118:e0:208:0:100

Table 26: Example

Host	Alias	IP Address
lbvip04	lbvip04	2607:f160:8205:8018:e0:108:0:100
lbvip05	lbvip05	2607:f160:8205:8018:e0:108:0:10d

NTP Configuration

For HA, add a row for each NTP server under additionalHosts section in YAML file. The Alias for the primary has to be **ntp** and the Alias for the secondary has to be **btp**. The NTP servers are configured in the `/etc/ntp.conf` of lb01/lb02/cluster manager.

Configuration based on Diameter Endpoints Interface

If the CPS platform is acting as a Diameter Server and using HAProxy, then you can configure AdditionalHosts and VipProxyConfiguration with interface hostname in the CPS Deployment Configuration Template (Excel Worksheet) based on the following table:

Table 27: Configuration with/without VIP Proxy

Traffic on Interface	Description
Only on LBvips	Configuration can be done using `VipProxyConfiguration.csv` file or `AdditionalHosts.csv` file. • **VipProxyConfiguration.csv** If using `VipProxyconfiguration.csv` file, remove diam-int* entries from `AdditionalHosts.csv` file. Configure all your VIPs in `VipProxyConfiguration.csv` file. For more information, refer to VIP Proxy Configuration, on page 72. • **AdditonalHosts.csv** Remove `VipProxyconfiguration.csv` file. All VIPs must be added in `AdditionalHosts.csv` file. For more information, refer to Diameter Related Configuration, on page 39.

Traffic on Interface	Description
Only on Policy Director (lb) interface For example, eth1	All entries should be present in `AdditionalHosts.csv` file. Remove `VipProxyconfiguration.csv` file.
On both the interfaces. For example, eth1 and eth1:1	All entries should be present in `AdditionalHosts.csv` file. Remove `VipProxyconfiguration.csv` file.

Diameter Related Configuration

If the CPS platform is acting as a Diameter Server and using HAProxy, then configure the AdditionalHosts tab with interface hostname in the CPS Deployment Configuration Template (Excel Worksheet) using the format and naming standard as described below. For a proper diameter stack configuration, the Policy Builder configuration must match ports defined in this tab (see the mapping table below for the port mapping in the Additional Notes:, on page 41 section).

The Cluster Manager supports the following scenarios for HAProxy Diameter:

- Single Endpoint:

  All diameter traffic comes into one NIC and same port. This is defined by adding an entry to **AdditionalHosts** tab of the Excel spreadsheet. The HAProxy binds to port 3868 on the defined IP for each host. Format of the hostname is *diam-int1-{hostname}*.

Note The format of the Hostname is *diam-int1-{hostname}*, where *{hostname}* is the guest name of a Policy Director (LB) VM. There will be one *{hostname}* for each Policy Director (LB) node (lb01, lb02...). Refer to your **Hosts.csv** file to get the required *{hostname}* values. An example is provided in the above screen shot.

For example:

Table 28: Single Endpoint

Hostname	IP Address
diam-int1-lb01	*XXX.XXX.XXX.XXX*
diam-int1-lb02	*YYY.YYY.YYY.YYY*

where, *XXX.XXX.XXX.XXX* is the IP address of diam-int1-lb01 and *YYY.YYY.YYY.YYY* is the IP address of diam-int1-lb02.

- Multiple VIP Endpoints:

  Diameter traffic for different interfaces (Gx, Rx and so on) can come into different NICs either on lb01 or lb02. This is defined by adding multiple 'diam-int*x*-vip' entries to **AdditionalHosts** tab of the deployment template spreadsheet. The HAProxy binds to port 3868 on the defined VIP on each host (that is, lb01 and lb02). Format of the hostname is diam-int*x*-vip.

Note For each VIP Endpoint, you must add the respective entry in VLANs tab.

For example,

Hostname IP Address

diam-int*x*-vip *XXX.XXX.XXX.XXX*

where,

x can have value from 1 to 4.

and *XXX.XXX.XXX.XXX* is the VIP address of the respective diameter interface.

If using `VipProxyConfiguration.csv` file, no need to configure the diam-int* entries in `AdditionalHosts.csv` file. Configure all your VIPs in `VipProxyConfiguration.csv` file. For more information, refer to VIP Proxy Configuration, on page 72.

- Multiple Endpoint/Multiple Interfaces:

  Multiple Interface/Endpoints are used when different diameters are coming from different networks and ports to provide more isolation of traffic. Diameter traffic comes into multiple NICs in Load Balancer, but all other traffic comes into the same interface and shares the same port. This is defined by adding multiple entries to **AdditionalHosts** tab of the Excel spreadsheet. The HAProxy binds to port 3868 on the defined IP for each host. Format of the hostname is *diam-int[1-4]-{hostname}*.

  For example:

  Table 29: Multiple Endpoint/Multiple Interfaces

Hostname	IP Address
diam-int1-lb01	*XXX.XXX.XXX.XXX*
diam-int1-lb02	*YYY.YYY.YYY.YYY*
diam-int2-lb01	*AAA.AAA.AAA.AAA*
diam-int2-lb02	*BBB.BBB.BBB.BBB*

  where, *AAA.AAA.AAA.AAA* is the IP address of diam-int2-lb01 and *BBB.BBB.BBB.BBB* is the IP address of diam-int2-lb02.

- Multiple Endpoint/Single Interface/Multiple Ports:

  Diameter traffic comes into Load Balancer via the multiple NIC, and also through different ports such as 3868, 3869,etc. This is defined by adding multiple entries to **AdditionalHosts** tab of the Excel spreadsheet. The HAProxy binds to port 3868 through 3871 on the defined IP for each host. Format of the hostname is *diam-int1-{hostname}* for port 3868 and *diam-int1-{hostname}-[69|70|71]* for ports 3869, 3870 and 3871.

  For example:

Table 30: Multiple Endpoint/Single Interface/Multiple Ports

Hostname	IP Address
diam-int1-lb01	*XXX.XXX.XXX.XXX*
diam-int1-lb01-69	*XXX.XXX.XXX.XXX*
diam-int1-lb01-70	*XXX.XXX.XXX.XXX*
diam-int1-lb01-71	*XXX.XXX.XXX.XXX*
diam-int1-lb02	*YYY.YYY.YYY.YYY*
diam-int1-lb02-69	*YYY.YYY.YYY.YYY*
diam-int1-lb02-70	*YYY.YYY.YYY.YYY*
diam-int1-lb02-71	*YYY.YYY.YYY.YYY*

Additional Notes:

The HAProxy configuration that is generated routes the requests to local endpoints in the same Policy Director VM (LB) where the diameter endpoints are anchored. In order to utilize this, the Policy Builder settings for diameter ports must be: 3868 for haproxy server 1, 3878 for haproxy server 2, 3888 for haproxy server 3 and 3898 for haproxy server 4. For example, setting up two stacks on separate VIPs would require setting the two hosts settings: stack 1 to port 3868 and stack 2 to 3878.

```
diam-int1-lb01(3868) - base port defined in stack as 3868, 3869, 3870
diam-int2-lb01 (3868)- base port defined in stack as 3878, 3879, 3880
diam-int3-lb01(3868) - base port defined in stack as 3888, 3889, 3890
diam-int4-lb01(3868) - base port defined in stack as 3898, 3899, 3900
diam-int1-lb01-69(3869) - base port defined in stack as 3878, 3879, 3880
diam-int1-lb01-70(3870) - base port defined in stack as 3888, 3889, 3890
diam-int1-lb01-71(3871)- base port defined in stack as 3898, 3899, 3900
```

HAProxy is used to perform least connection load balancing within a VM in CPS implementation and does not load balance across a VM.

In a CPS cluster which is configured with more than 2 Policy Directors (LBs), HAproxy and the VIPs are hosted only on LB01 and LB02. The additional LBs serve only as diameter endpoints to route diameter traffic.

Add Diameter Endpoints

To add diameter endpoints manually, modify the `/var/qps/current_config/image-map` file as follows.

In CPS 10.0.0 and higher releases, the `lb01` and `lb02` entries are replaced with a single `lb` entry, as shown in the following example:

```
lb=iomanager
lb=diameter_endpoint
lb=diameter_endpoint
lb=diameter_endpoint
lb=diameter_endpoint
lb=diameter_endpoint
lb=diameter_endpoint
lb=diameter_endpoint
lb=diameter_endpoint
```

```
qns=pcrf
pcrfclient=controlcenter
pcrfclient=pb
```

In releases prior to CPS 10.0.0:

```
lb01=iomanager01
lb02=iomanager02
lb=diameter_endpoint
lb=diameter_endpoint
lb=diameter_endpoint
lb=diameter_endpoint
lb=diameter_endpoint
lb=diameter_endpoint
lb=diameter_endpoint
lb=diameter_endpoint
qns=pcrf
pcrfclient=controlcenter
pcrfclient=pb
```

General Configuration

The Configuration sheet contains values for ESXi Users and the default CPS users, and some global variables that the puppet scripts use to generate the VMs.

To change the values on this tab, contact your Cisco Technical representative.

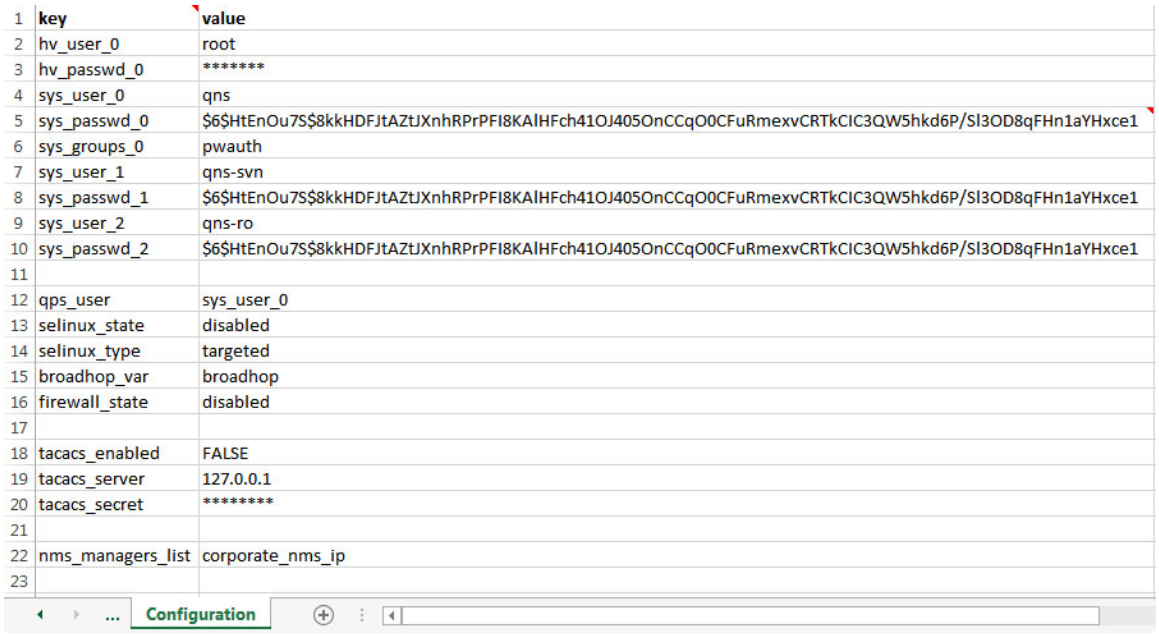
For users specified in this Configuration sheet, such as qns-admin, qns-svn, qns-ro, the password entered in the sheet is used. Any changes done manually to the system passwords after deployment would be overwritten by the password in the csv file after upgrade.



Note 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.

Figure 7: General Configuration



	key	value
2	hv_user_0	root
3	hv_passwd_0	*******
4	sys_user_0	qns
5	sys_passwd_0	6HtEnOu7S$8kkHDFJtAZtJXnhRPrPFI8KAlHFch41OJ405OnCCqO0CFuRmexvCRTkCIC3QW5hkd6P/Sl3OD8qFHn1aYHxce1
6	sys_groups_0	pwauth
7	sys_user_1	qns-svn
8	sys_passwd_1	6HtEnOu7S$8kkHDFJtAZtJXnhRPrPFI8KAlHFch41OJ405OnCCqO0CFuRmexvCRTkCIC3QW5hkd6P/Sl3OD8qFHn1aYHxce1
9	sys_user_2	qns-ro
10	sys_passwd_2	6HtEnOu7S$8kkHDFJtAZtJXnhRPrPFI8KAlHFch41OJ405OnCCqO0CFuRmexvCRTkCIC3QW5hkd6P/Sl3OD8qFHn1aYHxce1
11		
12	qps_user	sys_user_0
13	selinux_state	disabled
14	selinux_type	targeted
15	broadhop_var	broadhop
16	firewall_state	disabled
17		
18	tacacs_enabled	FALSE
19	tacacs_server	127.0.0.1
20	tacacs_secret	********
21		
22	nms_managers_list	corporate_nms_ip
23		

... Configuration

The following parameters can be configured in this sheet:

Table 31: General Configuration Parameters

Parameter	Description
hv_user_0	Hypervisor username. This is the username of a user with root access to the VMware host/blade. If installing CPS to multiple blade servers, it is assumed that the same username and password can be used for all blades. This parameter is optional[3].
hv_passwd_0	Hypervisor Password for Hypervisor User. User can also use special (non-alpha numeric) characters in the password. This parameter is optional. **Note** To pass special characters in the hv_passwd_0, they need to be replaced with its "% Hex ASCII". For example, "$" would be "%24" or "hello$world" would be "hello%24world".
vcenter_hostname	vCenter hostname. Example: qps-vcenter.cisco.com
vcenter_user	vCenter user[4]. Example: administrator@vsphere.local

Parameter	Description

vcenter_passwd	vCenter password[5]. You need to add the encrypted password.
	To encrypt the password,
	`cd /var/qps/bin/support`
	`./encrypt_pass.sh vcenter <vcenter_passwd>`
	where, <*vcenter_passwd*> is the vCenter password in plain text format.
	Note
	The `./encrypt_pass.sh vcenter <vcenter_passwd>` command must be run on every Cluster Manager and the `Configuration.csv` file should have the password generated for the respective Cluster Manager. The encrypted passwords cannot be reused on other Cluster Managers or setups.
	Note
	The encrypted password must be added in the `Configuration.csv` spreadsheet. To make the new password persisent, execute `import_deploy.sh`.

sys_user_0	The CPS System user (qns) is the main user set up on the VMs. By default, this is **qns**.
sys_passwd_0	Encrypted System Password for System User 0. Refer to System Password Encryption, on page 60 to generate an encrypted password. For High Availability (HA) environments or Geographic Redundancy (GR) environments, the password entered here in the spreadsheet is not used even if you specify one. You must set the password for the user prior to first access by connecting to the Cluster Manager after deployment and running the `change_passwd.sh` command.
sys_group	Group for the previous System User. **Note** User group can be qns-svn, qns-ro, qns-su, qns-admin and pwauth. pwauth group is valid only for `qns` username and no other username.
sys_user_1	The qns-svn system user is the default user that has access to the Policy Builder subversion repository. Default: qns-svn

Parameter	Description
sys_passwd_1	By default, the encrypted password for qns-svn is already added in `Configuration.csv` spreadsheet. If you want to change the password for qns-svn user after CPS is deployed, you can use `change_passwd.sh` script. You also need to generate an encrypted password. To generate an encrypted password, refer to System Password Encryption, on page 60. **Note** 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.
qps_user	-
selinux_state selinux_type	By default, Security Enhanced Linux (SELinux) support is disabled. **Note** Cisco recommends not to change this value.
firewall_state	Enables or disables the linux firewall on all VMs (IPtables). Valid Options: enabled/disabled Default: enabled (This field is case sensitive) **Note** An alternate parameter 'firewall_disabled' can be used with true/false options to control the IPtables functionality. **Note** In case the firewall is disabled, mongo authentication functionality for Policy Server (QNS) read-only users is also disabled. When firewall is enabled, mongo authentication functionality for read-only users is enabled by default. **Note** If the firewall is enabled/disabled, ICMP should not be blocked. If ICMP is blocked between VMs many of the dependent scripts and underlying framework fails to work. For example, blocking of ICMP can result in upgrade or migration failure, replica creation failure, and so on.
broadhop_var	Default: broadhop
tacacs_enabled	Enter **`true`** to enable TACACS+ authentication. For more information related to TACACS+, refer to TACACS+, on page 122.

Parameter	Description
tacacs_server	Enter the IP address of the TACACS+ server. **Note** If configured TACACS server is not reachable, Installation gets interrupted. To avoid interruption, make sure that the TACACS server is reachable and working before it makes part of the configuration.
tacacs_secret	Enter the password/secret of the TACACS+ server.
tacacs_on_ui	This parameter is used to enable the TACACS+ authentication for Policy Builder and Control Center. Default value is false. Possible values are true or false.
allow_user_for_cluman	This parameter is used to update the `/etc/sudoers` with CPS entries on cluman. Default value is false. Possible values are true or false.
nms_managers_list	Define the SNMP Network Management Station (NMS) address or hostname by replacing *corporate_nms_ip* with the hostname or IP address of your NMS. To add Multiple SNMP NMS destinations, replace *corporate_nms_ip* with a space separated list of hostnames or IP addresses of your NMS managers. For example: **`10.105.10.10 10.202.10.10`** or **`10.105.10.10 10.202.10.10 2003:3041::22:22`** or **`nms_main nms_bck`** To change the NMS trap receiver port, update **nms_managers_list** <*nms_manager_list:port_num*> For example, **nms_managers_list `corporate_nms_ip:6100`** **Note** Any hostnames defined should also be defined in the **Additional Hosts** tab of the deployment spreadsheet.
free_mem_per_alert	By default, a low memory alert is generated when the available memory of any CPS VM drops below 10% of the total memory. To change the default threshold, enter a new value (0.0-1.0) for the alert threshold. The system generates an alert trap whenever the available memory falls below this percentage of total memory for any given VM. Default: 0.10 (10% free).
free_mem_per_clear	Enter a value (0.0-1.0) for the clear threshold. The system generates a low memory clear trap whenever available memory for any given VM is more than 30% of total memory. Default: 0.3 (30% of the total memory).

Parameter	Description
syslog_managers _list	Entries are space separated tuples consisting of `protocol:hostname:port`. Currently, only UDP is supported. Default: 514 For example: **`udp:corporate_syslog_ip:514`** **`udp:corporate_syslog_ip2:514`**
syslog_managers _ports	A comma separated list of port values. This must match values in the syslog_managers_list.
logback_syslog_ daemon_port	Port value for the rsyslog proxy server to listen for incoming connections, used in the rsyslog configuration on the Policy Director (LB) and in the `logback.xml` on the OAM (PCRFCLIENT). Default: 6515
logback_syslog_ daemon_addr	IP address value used in the `/etc/broadhop/controlcenter/logback.xml` on the OAM (PCRFCLIENT). Default: lbvip02
cpu_usage_alert_ threshold	The following *cpu_usage* settings are related to the High CPU Usage Alert and High CPU Usage Clear traps that can be generated for CPS VMs. Refer to *CPS SNMP and Alarms Guide*, Release 9.1.0 and prior releases or *CPS SNMP, Alarms and Clearing Procedures Guide*, Release 10.0.0 and later releases for more details about these SNMP traps. Set the higher threshold value for CPU usage. System generates an Alert trap whenever the CPU usage is higher than this value.
cpu_usage_clear_ threshold	Set the lower threshold value for CPU usage. System generates a Clear trap whenever the CPU usage is lower than this value and alert trap already generated.
cpu_usage_trap_ interval_cycle	This value is used as an interval period to execute the CPU usage trap script. The interval value is calculated by multiplying 5 with the given value. For example, if set to 1 then the script is executed every 5 sec. The default value is 12, which means the script is executed every 60 seconds.
snmp_trap_community	This value is the SNMP trap community string. Default: broadhop
snmp_ro_community	This value is the SNMP read-only community string. Default: broadhop

Parameter	Description
monitor_replica_timeout	This value is used to configure timeout value. The default value is 540 sec considering four replica sets. The customer can set timeout value according to the number of replica sets in their network. To recover single session replica-set, it takes approx 120 sec and adding 20% buffer to it; we are using 540 sec for default (for four replica sets). Without any latency between sessionmgr VMs, one replica-set recovers in ~135 sec. If latency (40 -100 ms) is present between sessionmgr VMs we can add 10% buffer to 135 sec and set the timeout value for the required number of replica sets in customer's network.
snmpv3_enable	This value is used to enable/disable the SNMPv3 support on CPS. To disable the SNMPv3 support, set this value to FALSE. Default: TRUE
v3User	User name to be used for SNMPv3 request/response and trap. Default: cisco_snmpv3
engineID	This value is used for SNMPv3 request/response and on which NMS manager can receive the trap. It should be a hex value. Default: 0x0102030405060708
authProto	This value specifies the authentication protocol to be used for SNMPv3. User can use MD5/SHA as the authentication protocol. Default: SHA
authPass	This value specifies the authentication password to be used for SNMPv3 requests. It should have minimum length as 8 characters. Default: cisco_12345
privProto	This value specifies Privacy/Encryption protocol to be used in SNMPv3 request/response and SNMP trap. User can use AES/DES protocol. Default: AES
privPass	This value specifies Privacy/Encryption password to be used in SNMPv3. It is an optional field. If it is blank then value specified in authPass is used as privPass. Default: <blank>
sctp_enabled	By default, SCTP support is enabled. For more information about enabling/disabling this functionality, refer to SCTP Configuration, on page 63. Default: TRUE

Parameter	Description
corosync_ping_hosts	Moving corosync resources (like VIPs) when the connectivity is lost between lb01 or lb02 (or pcrfclient01/02) to hosts configured in this field. So if lb01 cannot connect to sessionmgr01 and sessionmgr02 then corosync resources (like VIPs) are moved from lb01 to lb02. Example: key = corosync_ping_hosts and Value = sessionmgr01 sessionmgr02
avoid_corosync_split_brain	If this field is not defined or value is 0, and when both nodes fail to connect to the configured corosync_ping_hosts, then the resources stay on the last active node. If value is 1, and both nodes fail to connect to configured corosync_ping_hosts, then the resources are not available on any nodes. **Remember** A split brain scenario (that is, VIPs are up on both nodes) can still occur when there is connectivity loss between lb01 and lb02 and not with other hosts.
rsyslog_tls	This field is used to enable or disable encryption for rsyslog. Default: TRUE
rsyslog_cert	This field is used to define the path for trusted Certificate of server.
rsyslog_ca	This field is used to define the Path of certifying authority (CA). Default: `/etc/ssl/cert/quantum.pem`
rsyslog_key	This field is used to define the path of private key.
haproxy_stats_tls	This field is used to enable or disable the encryption for HAproxy statistics (including diameter statistics). Default: TRUE
redis_authentication _enabled	This field is used to enable or disable Redis authentication. Default: TRUE (For fresh installations) To enable or disable redis authentication for upgrade and migration, refer to Redis Authentication for Upgrading/Migrating Systems, on page 63.
redis_authentication _passwd	This field is used to add an encrypted password for Redis. For more information on about generating encrypted password, refer to Redis Authentication, on page 61.
redis_server_count	This value specifies the number of redis server instances running on each policy director (lb) VM. For more information on redis functionality, refer to Configure Multiple Redis Instances, on page 124. Redis can be enabled with the number of instances as defined in *redis_server_count*. If the value for redis server count is not provided, default value of 3 for redis_server_count is considered. To disable redis explicitly, redis server count should have value 0. Default: 3 Value range: 0 to 64

Parameter	Description
remote_redis_server_count	This value can be added for Geographic Redundancy (GR) deployments only. This value specifies the number of redis server instances running on each remote policy director (lb) VM. If this value is not configured, remote redis server instances are not added for GR deployments.
snmpRouteLan	This field contains the value of a VLAN name which can be used to access the KPIs value provided by SNMP. Default: Oam

redis_for_ldap_required	This parameter is used only when dedicated LDAP instance is required.
	Default: false
	Possible Values: true, false
	If you configure LDAP instance explicitly, first redis instance on policy director (lb) VMs running on port 6379 is used for LDAP and the remaining are used for diameter.
	Note
	If you configure `redis_for_ldap_required` parameter, then the following changes are automatically added in configuration files.
	In `/etc/broadhop/qns.conf` file, an additional parameter `-DldapRedisQPrefix=ldap` is added.
	`/etc/broadhop/redisTopology.ini` file has the following content if `redis_for_ldap_required=true` and `redis_server_count=3`:
	`ldap.redis.qserver.1=lb01:6379`
	`policy.redis.qserver.2=lb01:6380`
	`policy.redis.qserver.3=lb01:6381`
	`ldap.redis.qserver.4=lb02:6379`
	`policy.redis.qserver.5=lb02:6380`
	`policy.redis.qserver.6=lb02:6381`
	If a dedicated LDAP instance is required, you many also want to consider increasing the total redis servers to accomodate the diameter traffic.
	For example, if `redis_for_ldap_required` property was not configured, and `redis_server_count=3` then after configuring `redis_for_ldap_required` as true, you want to increase total redis server count to 4 by setting `redis_server_count=4`.

Parameter	Description
database_nics	This parameter allows user to provide interface names on which firewall must be opened for replica-set on a VM. If `database_nics` is not configured, firewall is opened only for internal interface for a replica-set. If `database_nics` is configured, then firewall is opened for configured interfaces and internal interface as well (even if it is not mentioned in `database_nics`). This field has semicolon (;) separated interface names for firewall ports to be opened for a replica-set on a VM. **Note** This field is effective only when the firewall is enabled.
db_authentication_ enabled	This field is used to enable or disable MongoDB authentication. Possible Values: TRUE, FALSE **Note** You must configure `db_authentication_ enabled` parameter. This parameter cannot be left empty. To disable the authentication, the parameter value must be set as FALSE. To enable, the value should be TRUE, and admin and readonly passwords must be set. This is applicable only for new installs and not for upgrades. For more information, refer to MongoDB Authentication, on page 64.
db_authentication_ admin_passwd	This parameter is the encrypted password for admin user and is applicable only when *db_authentication_enabled* is set to **`TRUE`**. The following command is used to generate encrypted password from Cluster Manager: `/var/qps/bin/support/mongo/encrypt_passwd.sh <Password>` For more information, refer to MongoDB Authentication, on page 64.
db_authentication_ readonly_passwd	This parameter is the encrypted password for readonly user. The following command is used to generate encrypted password from Cluster Manager: `/var/qps/bin/support/mongo/encrypt_passwd.sh <Password>` For more information, refer to MongoDB Authentication, on page 64.
remote_site_ip	This parameter is used to update the remote site Cluster Manager IP address. **Note** This parameter is used only for GR and multi-cluster setups.
enable_ssh_login_security	This parameter allows user to enable or disable SSH login security. Default: disabled Possible Values: enabled, disabled
cps_admin_user_cluman	This parameter is used to configure Cluster Manager administrator user.

Parameter	Description
cps_admin_password_ cluman	This parameter is the encrypted password for administrator user.
whitelisted_hosts_for_ssh	Valid values are colon separated host names/IP addresses of the machine for which SSH access needs to be allowed. This configuration is effective only when the SSH login security is enabled. If the hostname is mentioned then it should be resolvable by CPS VM's. No validation on hostname/IP addresses is provided. You can specify both IPv4/IPv6 address. **Note** New whitelisted host list overwrites the old list. If the new whitelist host configuration is empty then all old additional whitelisted hosts (apart from standard local CPS VM's host) are deleted.
LDAP SSSD Configuration	For more information, refer to LDAP SSSD Configuration, on page 66.
enable_prometheus	This parameter is used to enable/disable Prometheus in CPS. Default: disabled Possible Values: enabled, disabled For more information, refer to *Prometheus and Grafana* chapter in *CPS Operations Guide*.
stats_granularity	This parameter is used to configure statistics granularity in seconds. Default: 10 seconds Possible Values: Positive Number For more information, refer to *Prometheus and Grafana* chapter in *CPS Operations Guide*.
restrict_access_http_port	When set to true, the http port (80) on pcrfclient and Cluster Manager VMs listen only on internal guest NIC and loopback interface. By default, this parameter is not present in `Configuration.csv` file. Possible Values: true, false
service_log_tmpfs_enabled	This parameter is used to enable or disable service log on tmpfs. Currently, this is supported only on Policy Director (LB), Policy Server (QNS) and UDC VMs. Default: false Possible Values: true, false If this parameter is not configured, then by default, the value is false.

Parameter	Description
pcrf_proc_mon_list	This parameter is used to configure additional processes on OAM (pcrfclient) VMs. Multiple processes need to be semicolon separated. By default, the following processes are monitored: • Collectd • Qns-* (all instances of qns-java process) • Logstash • Httpd • Snmpd • Carbon-cache • Carbon-cache@b • Carbon-cache@c • Carbon-aggregator • Carbon-aggregator@b • Monit
lb_proc_mon_list	This parameter is used to configure additional processes on Policy Director (LB) VMs. Multiple processes need to be semicolon separated. By default, the following processes are monitored: • Collectd • Qns-* (all instances of qns java processes) • Snmpd • Snmptrapd • Corosync • Redis-* (all instances of redis processes) • Haproxy • Haproxy-diameter • Memcached • zing-licensem • zing-licensed

Parameter	Description
qns_proc_mon_list	This parameter is used to configure additional processes on Policy Server (QNS) VMs. Multiple processes need to be semicolon separated. By default, the following processes are monitored: • Collectd • Qns-* (all instances of qns-java process) • Monit • zing-licensem • zing-licensed
sm_proc_mon_list	This parameter is used to configure additional processes on sessionmgr VMs. Multiple processes need to be semicolon separated. By default, the following processes are monitored: • Collectd • Memcached • All SM replica-set members mongodb processes
udc_proc_mon_list	This parameter is used to configure additional processes on UDC VMs. Multiple processes need to be semicolon separated. By default, the following processes are monitored: • Collectd • Qns-* (all instances of qns-java process)
lwr_proc_mon_list	This parameter is used to configure additional processes on LWR VMs. Multiple processes need to be semicolon separated. By default, the following processes are monitored: • Collectd • monit

Parameter	Description
perf_mod	1 or undefined: CPS java processes are run by Zulu on Policy Server (QNS), Policy Director (LB), and UDC VMs. **Note** Zing is only supported on Policy Server (QNS), Policy Director (LB), UDC, and LWR VMs. It is not supported on pcrfclient and session manager VMs. **Note** In CPS 21.2 release and later releases, Zing package is no longer installed on Policy Director (LB) abd UDC VMs. By default (1), CPS java process is run by Zulu on Policy Server (QNS), Policy Director (LB), and UDC VMs. If 2: CPS java processes are run by Zing on Policy Server (QNS), Policy Director (LB), and UDC VMs in the VMware. To disable Zing, refer to Disable Zing, on page 70.
gc_alarm_state	This parameter is used to enable or disable the GC alarm. Default: false Possible Values: true, false
gc_alarm_ trigger_count	This parameter is used to configure the value of continous GCs after which the GC alarm is generated from the system. Default: 3
gc_alarm_ trigger_interval	This parameter is used to indicate the interval under which the gc_alarm_trigger_count occurs to generate the GC alarm. Default: 600 (10 mins)
gc_clear_ trigger_interval	This parameter is used to indicate the interval under which the there is no GC event and GC clear notofication is generated. Default: 900 (15 mins)
oldgen_ alarm_state	This parameter is used to enable or disable the Old generation alarm. Default: false Possible Values: true, false
oldgen_alarm_ trigger_thr_per	This parameter is used to indicate the threshold in percentage for Old Generation post GC event to generate the Old Generation alarms. Default: 50
oldgen_clear_ trigger_thr_per	This parameter is used to indicate the threshold in percentage for Old Generation post GC event to generate the Old Generation clear notification. Default: 40

Parameter	Description
no_of_cont_ fullgc_for_oldgen	This parameter is used to indicate the number of continuous GC events under which the Old generation value is more than oldgen_alarm_trigger_thr_per to generate the Old generation alarm. Default: 2
alarm_resync_enabled	This parameter is used to store and forward the alarms based on NMS availability. Default: false Possible Values: true, false, TRUE, FALSE, True, False **alarm_resync_enabled** should start at the first characer in the line and there must be no additional characters after true/false. **Restriction** When NMS comes up, it takes almost 5 mins for system to start sending the stored alarms to NMS. In between if any alarm gets generated by the system, it is sent to NMS. So there is possibility that NMS may receive latest alarms first and all the older alarms later. This happens only when NMS is unreachable and comes back to rechable state. • Multiple NMSs are configured: If few NMS servers are down, then the alarm resync feature will not store the alarms to be sent to NMS. • Multiple NMSs are configured: If all NMS servers are down, then the alarm resync feature stores the alarms in Admin database and sends them to NMS when it is reachable. • Single NMS is configured: If NMS server is down, then the alarm resync feature stores the alarms in Admin database and sends them to NMS when it is reachable. • Single NMS is configured: If NMS server is up, then the alarm resync feature does not not store the alarms but just forwards the alarms to NMS.
autoheal_qns_enabled	`autoheal_qns_enabled` parameter helps `app_monitor.sh` script (application monitor script) to take the decision to restart the QNS process or not. • **FALSE:** To disable the restart QNS process in case of the MongoDB health monitor failed to reset the MongoDB client connection. • **TRUE:** To enable the restart QNS process in case of the MongoDB health monitor failed to reset the MongoDB client connection. For installing platform script, refer to Installing Platform Scripts for MongoDB Health Monitoring - VMware.

Parameter	Description
prevent_primary_ flapping_enabled	This parameter is used to prevent primary flapping from impacting the remote sites. Default: false **Restriction** • When the local site is handling traffic, during local site reboot scenario, if the latency is more between the local and remote sites, then there may be some timeout or high response time from remote site since the PRIMARY is shifted to remote site. • If the member state is not stable within the stipulated 300 seconds time, then the priority level is retained as 1 for those members until it becomes stable for minimum 300 seconds. • If `mon_db*` is enabled, make sure not to enable the `prevent_primary_flapping_enabled` flag. If both the parameters are enabled in a setup, it creates conflicts in MongoDB operations.

auto_haproxy_
balancing_list
	This parameter is used to add the list of diameter endpoints that are enabled for Policy Director (LB) HAProxy Balancing.
	For example:
	`$ cat /var/qps/config/deploy/csv/Configuration.csv | grep auto_haproxy_balancing_list`
	`auto_haproxy_balancing_list,diameter-int1-vip diameter-int2-vip,`
	To disable the HAProxy balancing, `auto_haproxy_balancing_list` is set to empty.
	For example:
	`$ cat /var/qps/config/deploy/csv/Configuration.csv | grep auto_haproxy_balancing_list`
	`auto_haproxy_balancing_list, ,`

gx_alarm_ccr_i _avg_threshold	This parameter is used to specify the threshold value for Gx CCR-I response time in *Gx Average Message processing Dropped* alarm. Default: 20 millisec For alarm information, refer to *Gx Average Message processing Dropped* in *CPS SNMP, Alarms, and Clearing Procedures Guide*.
gx_alarm_ccr_u_ avg_threshold	This parameter is used to specify the threshold value for Gx CCR-U response time in *Gx Average Message processing Dropped* alarm. Default: 20 millisec For alarm information, refer to *Gx Average Message processing Dropped* in *CPS SNMP, Alarms, and Clearing Procedures Guide*.

Parameter	Description
gx_alarm_ccr_t_ avg_threshold	This parameter is used to specify the threshold value for Gx CCR-T response time in *Gx Average Message processing Dropped* alarm. Default: 20 millisec For alarm information, refer to *Gx Average Message processing Dropped* in *CPS SNMP, Alarms, and Clearing Procedures Guide*.
ldap_alarm_retry_ threshold	This parameter is used to specify the threshold value for *Percentage of LDAP retry threshold Exceeded* alarm. Default: 10 % For alarm information, refer to *Percentage of LDAP retry threshold Exceeded* in *CPS SNMP, Alarms, and Clearing Procedures Guide*.
ldap_alarm_ccr_i_ req_threshold	This parameter is used to specify the threshold value for *LDAP Requests as percentage of CCR-I Dropped* alarm. Default: 25 % For alarm information, refer to *LDAP Requests as percentage of CCR-I Dropped* in *CPS SNMP, Alarms, and Clearing Procedures Guide*.
ldap_alarm_result_ threshold	This parameter is used to specify the threshold value for *LDAP Query Result Dropped* alarm. Default: 0 (recommended) For alarm information, refer to *LDAP Query Result Dropped* in *CPS SNMP, Alarms, and Clearing Procedures Guide*.
ldap_alarm_request_ threshold	This parameter is used to specify the threshold value for *LDAP Requests Dropped* alarm. Default: 0 For alarm information, refer to *LDAP Requests Dropped* in the *CPS SNMP, Alarms, and Clearing Procedures Guide*.
client_Alive_Interval	This parameter represents SSH idle timeout. This value is configured in seconds. For example: client_Alive_Interval, 500 Default value is 0 (zero).
MongoDB Replication Health Monitoring	For more information, refer to MongoDB Replication Health Monitoring, on page 70.

Parameter	Description

pcrfclient_memcache_ memory_size	This parameter is used to change the memcached memory size on pcrfclients.
	This parameter doesn't change memcached memory on other VMs.
	For example: pcrfclient_memcache_memory_size,8192
	Default: 2048
	Note You can change the value as per your deployment requirements. Make sure you have enough memory on pcrfclients to support the change in memcached memory size.
	If you are adding `pcrfclient_memcache_memory_size` parameter to an existing installation, execute the following:
	`/var/qps/install/current/scripts/import/import_deploy.sh`
	`/var/qps/install/current/scripts/build/build_all.sh`
	`/var/qps/install/current/scripts/upgrade/reinit.sh`

balance_mgmt_ fragmentation_threshold	This parameter is used to specify the threshold value for **balance_mgmt MongoPrimaryDB fragmentation** that exceeded the threshold value alarm. Default: 40
diameter_ fragmentation_threshold	This parameter is used to specify the threshold value for **diameter MongoPrimaryDB fragmentation** that exceeded the threshold value alarm. Default: 40
session_cache_ fragmentation_threshold	This parameter is used to specify the threshold value for **session_cache MongoPrimaryDB fragmentation** that exceeded the threshold value alarm. Default: 40
sk_cache_ fragmentation_threshold	This parameter is used to specify the threshold value for **sk_cache MongoPrimaryDB fragmentation** that exceeded the threshold value alarm. Default: 40
spr_fragmentation_ threshold	This parameter is used to specify the threshold value for **spr MongoPrimaryDB fragmentation** that exceeded the threshold value alarm Default: 40
db_fragmentation_ alarm_enable	This parameter is to enable or disable **MongoPrimaryDB fragmentation** that exceeded the threshold value alarm. Default: false
enable_mongodb_majority_failover_monit	Set the value as true or false to enable or disable the feature respectively. Default Value: False
majority_failover_monit_cycles	Set the time interval value in seconds to periodically run the MongoDB recovery Script.

Parameter	Description
majority_failover_action	Choose any of the following actions on the member which is down: • REDUCE_PRIORITY - Reduce the priority and vote the member to 0. (Default, Recommended) • REMOVE_MEMBER - Remove the member from the replica set.
majority_failover_iteration_threshold	Number of iterations the MongoDB recovery Script can wait before taking `majority_failover_action` on a member that is is down. Default Value: 3

[3] In CPS 11.0.0 and later releases, these two parameters (hv_user_0 and hv_password_0) are optional in /var/qps/config/deploy/csv/Configuration.csv file and the user is prompted for the parameters at runtime while executing `deploy_all.py` and `deploy.sh` scripts if not configured in Configuration.csv file. Now during installation on VMware, hypervisor password is not displayed on terminal by any scripts. Also, hypervisor password is not logged into any of the log files.

[4] vcenter_user and vcenter_passwd should have the administrative privilege credentials.

[5] If user misses to add vcenter_hostname, vcenter_user and vcenter_passwd in the `Configuration.csv` file, after executing `deploy_all.py` script, the user is prompted to enter the vcenter information in the command line. User has to enter the unencrypted vcenter_passwd.

Note Execute the following command to import changes done in `Configuration.csv` files into the Cluster Manager VM:

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

System Password Encryption

Use the following step to generate a password hash:

Important Password encryption method has changed. This method can be used for fresh install and new user. Existing users and passwords work without any problem. You need to update your old CSV/YAML files with new encrypted passwords.

When ISSM is performed from an older release to this release, use `generate_encrypted_password.sh` script.

1. Execute `/var/qps/install/current/scripts/bin/support/generate_encrypted_password.sh` script to get encrypted password.

```
[root@installer support]# generate_encrypted_password.sh
                        ##########################################
                        #          CISCO SYSTEMS PVT LTD         #
                        ##########################################

Hello, user! You are attempting to change your password.
Great! A few ground rules:

1. No short passwords. The longer your password is, the harder it is for someone
```

```
   to guess or figure out with brute force. Minimum password length is 8.

2. There needs to be at least: one uppercase letter, one lowercase letter, one
   digit and one special character. This increases the search space and makes
   brute force guessing more difficult.

3. You can use spaces. Feel free to use a sentence as your password. For example:
   I bring 2 gifts! is easy to remember, not cumbersome to type, meets all the
   above criteria.

4. You will need to change your password every 6 months.

Enter the password:
```

2. After script execution the encrypted password is displayed as follows.

   Here is a sample encrypted password:

```
+-------------------------------------------------------------------------------------------------------------+
| Fri May 29 11:43:47 UTC 2020
                           |
| Encrypted key
                           |
|
$6$bc732ffd2a5ad85e$dYuQfGowAsAS6E2mQyWgGtcSUY4IKss11.4AY1u852gGwZzr4Y54rBdkHG6zQytFPXXDJGwknx.IYIeDeW.jP.
     |
+-------------------------------------------------------------------------------------------------------------+
```

Redis Authentication



Important All access to Redis Server from application would require password after the server is enabled with authentication. Application reads the encrypted password from environment variable, decrypts it and uses it to connect to Redis Server.

The following sections provide information about redis password encryption and authentication for fresh or an existing installation setups:

Password Encryption

Run the following command to generate an encrypted password:

`/var/qps/bin/support/redis/encrypt_passwd.sh <XXXXX>`

where, <*XXXXX*> is the plain text password for Redis.

Run `import_deploy.sh` script.

`/var/qps/install/current/scripts/import/import_deploy.sh` creates a readonly file called `.redis` with encrypted password under home folder of the user based on the `redis_authentication_enabled` and `redis_authentication _passwd` parameter values.

Redis Authentication

For fresh installations, redis authentication must be enabled by configuring `redis_authentication_enabled` and `redis_authentication_passwd` parameters in `Configuration.csv` file.

Installation fails if `redis_authentication_enabled` field is not present in `Configuration.csv` file. If you want to disable Redis Authentication by default, then `redis_authentication_enabled` must be set to FALSE.

Note A readonly file `.redis` is not created under home folder of the user when `redis_authentication_enabled` is set to FALSE.

A readonly file `.redis` is created under home folder of the user when `redis_authentication_enabled` is set to TRUE.

Enable or Disable Redis Authentication on an Existing System

Caution Enabling or disabling Redis authentication on an existing system requires application downtime.

`/var/qps/bin/support/redis/redis_auth_upgrade.sh` command must be used to enable or disable Redis authentication on an existing system.

```
/var/qps/bin/support/redis/redis_auth_upgrade.sh
Valid arguments are not provided to the script
redis_auth_upgrade.sh <OPTION> <PASSWORD>
OPTION:
  -e / --enable                         Enable Redis Password Authentication
  -d / --disable   <password>           Disable Redis Password Authentication
  -c / --chpass    <password>           Change Redis Password
  -h / --help                           Display this help and exit
PASSWORD:
  <password>       Existing plaintext password
```

Enable Redis Authentication: Here is an example configuration:

```
/var/qps/bin/support/redis/redis_auth_upgrade.sh -e

Enabling Redis Authentication...
Reading password file...
Enabling Redis Authentication on lb01:6379
OK
Enabling Redis Authentication on lb01:6380
OK
Enabling Redis Authentication on lb01:6381
OK
Enabling Redis Authentication on lb02:6379
OK
Enabling Redis Authentication on lb02:6380
OK
Enabling Redis Authentication on lb02:6381
OK
```

Disable Redis Authentication: Here is an example configuration:

```
/var/qps/bin/support/redis/redis_auth_upgrade.sh -d cisco123
Disabling Redis Authentication...
Disabling Redis Authentication on lb01:6379
OK
Disabling Redis Authentication on lb01:6380
OK
Disabling Redis Authentication on lb01:6381
OK
```

```
Disabling Redis Authentication on lb02:6379
OK
Disabling Redis Authentication on lb02:6380
OK
Disabling Redis Authentication on lb02:6381
OK
```

Redis Authentication for Upgrading/Migrating Systems



Caution Enabling or disabling Redis authentication for upgraded or migrated systems require application downtime.

Change Redis User Password

1. Modify password in `Configuration.csv` file.
2. Update configuration file using `import_deploy.sh` and `reinit.sh` scripts.
3. Run `redis_auth_upgrade.sh` script to change the password and provide the old plain text password.

   ```
   /var/qps/bin/support/redis/redis_auth_upgrade.sh -c <old_plaintext_password>
   ```
4. Restart all the java processes.

Disable Redis Authentication

1. Modify redis authentication in `Configuration.csv` file.
2. Update configuration file using `import_deploy.sh` and `reinit.sh` scripts.
3. Run `redis_auth_upgrade.sh` script to disable authentication and provide the plain text password.

   ```
   /var/qps/bin/support/redis/redis_auth_upgrade.sh -d <plaintext_password>
   ```
4. Restart all the java processes.

Enable Redis Authentication

1. Modify redis authentication in `Configuration.csv` file.
2. Update configuration file using `import_deploy.sh` and `reinit.sh` scripts.
3. Run `redis_auth_upgrade.sh` script to enable authentication and provide the plain text password.

   ```
   /var/qps/bin/support/redis/redis_auth_upgrade.sh -e <plaintext_password>
   ```
4. Restart all the java processes.

SCTP Configuration

CPS also support Stream Control Transmission Protocol (SCTP). By default, SCTP support is enabled.

To disable or enable SCTP on an existing deployment:

Procedure

Step 1 Update the field *sctp_enabled* to FALSE or TRUE in `/var/qps/config/deploy/csv/Configuration.csv` file with the following information:

`sctp_enabled,FALSE,`

or

`sctp_enabled,TRUE,`

Step 2 Import the new configuration by executing the following command:

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

Step 3 For an existing deployed lb0X VM, after changing *sctp_enabled* (such as, TRUE to FALSE or FALSE to TRUE), re-initialize lb0*X* VM by executing the following command:

```
ssh lb0X /etc/init.d/vm-init-client
```

Note
If setting it from TRUE to FALSE, then restart the VM for the changes to take effect.

MongoDB Authentication

For upgrades/migration, `/var/qps/install/current/scripts/import/import_deploy.sh` updates `dbPassword` parameter in `/etc/broadhop/qns.conf` file based on `db_authentication_enabled` and `db_authentication_admin_passwd` fields. It also creates `<user-home-directory>/.dbadmin` and `<user-home-directory>/.dbreadonly` files, which stores the encrypted password for admin and readonly users respectively.

- *`<user-home-directory>`*`/.dbadmin` file is created for root, qns, qns-su and qns-admin users.
- *`<user-home-directory>`*`/.dbreadonly` file is created for root, qns, qns-su, qns-admin and qns-ro users.



Note Traffic errors and timeouts might be seen during Enable/Disable MongoDB Authentication.

Use Cases

- Disable authentication (Fresh install):

  ```
  db_authentication_enabled,FALSE
  ```

  Output: dbPassword field is not present in `/etc/broadhop/qns.conf` file and there is no *`<user-home-directory>`*`/.dbadmin` and *`<user-home-directory>`*`/.dbreadonly` files.

- Enable authentication (Fresh install):

  ```
  db_authentication_enabled,TRUE
  db_authentication_admin_passwd,XXXX
  ```

```
db_authentication_readonly_passwd,YYYY
remote_site_ip,X.X.X.X <--- Only required for GR and multi-cluster setups
```

where, XXXX and YYYY are encrypted passwords.

Output: dbPassword field is added in `/etc/broadhop/qns.conf` file and *<user-home-directory>*`/.dbadmin` and *<user-home-directory>*`/.dbreadonly` files are created for users with permission 400 set to (read only to that user).

- Enabling or disabling authentication on an existing system:

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

Example:

```
/var/qps/install/current/scripts/modules/mongo_auth_upgrade.py:
INFO     ===== mongo upgrade =====
INFO     Parsing Mongo Config file
INFO     Mongo authentication is enabled on this system
INFO     Following replica sets need to enable authentication: ['set01', 'set02']
Do you want to enable mongo auth on these sets? (y/n):
```

MongoDB Authentication Process

- Change MongoDB user password:
  - Modify password in `Configuration.csv` file.



Note Update encrypted password in `configuration.CSV`.

  - After modifying the password, update the configuration using `/var/qps/install/current/scripts/import/import_deploy.sh` and `/var/qps/install/current/scripts/upgrade/reinit.sh` scripts.
  - Execute change password script (`/var/qps/install/current/scripts/modules/mongo_change_password.py`) and enter the old password.

    Syntax:

    ```
    /var/qps/install/current/scripts/modules/mongo_change_password.py <old_password>
    ```



Note The `old_password` is an unencrypted password.

- Disable MongoDB authentication:
  - Modify MongoDB authentication configuration in `Configuration.csv` file.

    ```
    db_authentication_enabled,FALSE
    remote_site_ip,X.X.X.X <--- Only required for GR and multi-cluster setups
    ```

  - Update the configuration using `/var/qps/install/current/scripts/import/import_deploy.sh` and `/var/qps/install/current/scripts/upgrade/reinit.sh` scripts.

- Execute disable MongoDB authentication script:
  `/var/qps/install/current/scripts/modules/mongo_auth_upgrade.py`

- Enable MongoDB authentication:
  - Execute `/var/qps/bin/support/mongo/encrypt_passwd.sh <Password>` command to encrypt the password.
  - Modify MongoDB authentication configuration in `Configuration.csv` file.

    ```
    db_authentication_enabled,TRUE
    db_authentication_admin_passwd,XXXX
    db_authentication_readonly_passwd,YYYY
    remote_site_ip,X.X.X.X <--- Only required for GR and multi-cluster setups
    ```

    where, *XXXX* and *YYYY* are encrypted passwords.
  - Update the configuration using `/var/qps/install/current/scripts/import/import_deploy.sh` and `/var/qps/install/current/scripts/upgrade/reinit.sh` scripts.
  - Execute enable MongoDB authentication script:
    `/var/qps/install/current/scripts/modules/mongo_auth_upgrade.py`

LDAP SSSD Configuration

Note For LDAP SSSD routable IP is required. LDAP server must be accessible from CPS VMs (LDAP client).

For information on Policy Builder and Grafana configuration, refer to *LDAP SSSD* section in *CPS Operations Guide*.

Note Add the LDAP server IP address and server name in `AdditionalHost.csv` file. For more information, refer to Additional Hosts Configuration, on page 37.

HA Setup

For LDAP SSSD configuration, the following parameters can be configured in `Configuration.csv` sheet:

Note Change the parameter values as per your deployment.

Table 32: LDAP SSSD Parameters

Parameter	Description
ldap_on_all	When set to true, it installs the LDAP SSSD on all CPS VMs. When set to false, it install the LDAP SSSD only on pcrfclient/policy directors (lb) VMs. **Note** true or false must be in small case.
ldap_enabled	When set to true, applies the SSSD configuration as per input provided by user. When set to false, use the default configuration. **Note** true or false must be in small case.
ldap_server	Contains server IP:port to configure LDAP. Format: ldaps://<*serverip*>:<*port*>
ldap_search_base	This is required for SSSD configuration. The default base DN to use for performing LDAP user operations. Format: ou=users,dc=cisco,dc=com
ldap_default_bind_dn	The default bind DN to use for performing LDAP operations. Format: uid=admin,ou=system
ldap_secret	The authentication token for the default bind DN. Currently, only clear text passwords are supported. For example, secret
ldap_default_user	The default LDAP user to be configured in LDAP server. For example, admin
ldap_ou_user	The default LDAP user OU. For example, users
ldap_ou_group	The default LDAP group user OU. For example, groups
ldap_default_group	The LDAP attribute that corresponds to the group name. For example, Admin
ldap_default_group_editor	This is a user group which has the editor access to Grafana. For example, User
ldap_dc_name	This is a single entity of all domains. Format: dc=cisco,dc=com

Here is an example configuration:

```
ldap_on_all,true
ldap_enabled,true
ldap_server,"ldaps://<serverip>:10648"
ldap_search_base,"ou=users,dc=cisco,dc=com"
ldap_default_bind_dn,"uid=admin,ou=system"
ldap_secret,secret,
ldap_default_user,admin,
ldap_ou_user,users,
ldap_ou_group,groups,
ldap_default_group,Admin,
ldap_default_group_editor,User,
ldap_dc_name,"dc=cisco,dc=com"
```

Run `/var/qps/install/current/scripts/bin/support/enable_ldap clustermgr` to install the LDAP SSSD configuration on Cluster Manager.

Run `puppet apply --logdest /var/log/cluman/puppet-run.log --modulepath=/opt/cluman/puppet/modules --config /opt/cluman/puppet/puppet.conf /opt/cluman/puppet/nodes/node_repo.pp` from Cluster Manager to update the puppet.



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

Arbiter Setup

You need to create `ldapconf` file to add the required parameters to configure LDAP SSSD.

Here is an example configuration:

```
# /var/qps/config/deploy/ldapconf
ldap_on_all,true
ldap_enabled=true
ldap_server="ldaps://<serverip>:<port>"
ldap_search_base="ou=users,dc=cisco,dc=com"
ldap_default_bind_dn="uid=admin,ou=system"
ldap_secret=secret,
ldap_default_user=admin,
ldap_ou_user=users,
ldap_ou_group=groups,
ldap_default_group=Admin,
ldap_default_group_editor=User,
ldap_dc_name="dc=cisco,dc=com",
NODE_TYPE=arbiter
```

Run `/var/qps/install/current/scripts/bin/support/enable_ldap clustermgr` to install the LDAP SSSD configuration on arbiter.

Run `puppet apply --logdest /var/log/cluman/puppet-run.log --modulepath=/opt/cluman/puppet/modules --config /opt/cluman/puppet/puppet.conf /opt/cluman/puppet/nodes/node_repo.pp` from Cluster Manager to update the puppet.



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

LDAP SSSD Certificate Authentication

LDAP certificate needs to copied to `/etc/openldap/certs/`on all VMs.

After copying the certificate, run the following commands on pcrfclient01 and pcrfclient02:

Note LDAP certificate must be provided by the customer.

```
export CLASSPATH=/usr/java/default/bin
keytool -import -noprompt -trustcacerts -alias ldap_l -file /etc/openldap/certs/ldap_local.cer
 -keystore /usr/lib/jvm/zulu-8/jre/lib/security/cacerts
```

This prompts for the password and the keytool password is "changeit".

Once the certificate authentication is complete, `/var/broadhop/scripts/update-uaf.sh` script runs every hour in crontab. This updates the user information in the `/var/www/svn/users-access-file` file on pcrfclient01 and pcrfclient02.

After pcrfclient VM is rebooted/re-deployed or vm-init script is executed, check whether the class path (`CLASSPATH=/usr/java/default/bin`) has been set on pcrfclient01 and pcrflcient02 by running `echo $CLASSPATH` command.

Also check whether the certificate (`/etc/openldap/certs/ldap_local.cer`) is present or not, run ls - l command.

If the class path or certificate path is missing, run the following commands:

```
export CLASSPATH=/usr/java/default/bin
keytool -import -noprompt -trustcacerts -alias ldap_l -file /etc/openldap/certs/ldap_local.cer
 -keystore /usr/lib/jvm/zulu-8/jre/lib/security/cacerts
```

Note After installing LDAP SSSD on all VMs if you want to remove from LDAP SSSD from policy server (qns) and sessionmgr, then you need to run `reinit.sh` script twice or run `/etc/init.d/vm-init` on individual policy servers (qns) and sessionmgr VMs.

Upgrade Consideration

After upgrading, LDAP SSSD configuration is installed on default VM (pcrfclient/lb) and not on all VMs. You need to configure LDAP SSSD on all the other VMs.

Once LDAP SSSD configuration is complete, you need to authenticate the LDAP certificate. For more information, refer to LDAP SSSD Certificate Authentication, on page 69.

Note If upgrading from a lower version and do not want the LDAP SSSD package, modify the LDAP parameters as follows in `Configuration.csv`:

```
ldap_on_all=false
ldap_enable=false
```

After the modification, run `import_deploy.sh` so that LDAP SSSD is not installed by default.

Troubleshooting

- Monitor the following important log files to debug grafana and httpd service:
  - `/var/log/messages`
  - `/var/log/secure`
  - `/var/log/audit/audit.log`
  - `/var/log/sssd/*.log`
  - `/var/log/grafana/grafana.log, /var/log/httpd/*.log`
  - `/var/log/broadhop/scripts/ldap*.log`
- Restart the httpd service and grafana-server in case grafana status is `Not Running` in monit summary after configuring LDAP SSSD.
- If any error is found for HA deployments after configuring LDAP SSSD, restart the http/grafana-server.
- If LDAP SSSD user information is not automatically added in `/var/www/svn/users-access-file` on pcrfclient01/02, then check `/var/log/broadhop/scripts/ldap*.log` for error information.

Disable Zing

Configure perf_mod with the list of diameter endpoints and follow the steps to disable Zing in the VMware:

1. Add **perf_mod** with value set to 1.

   ```
   $ cat /var/qps/config/deploy/csv/Configuration.csv | grep -i perf
   perf_mod,1,
   ```

2. Execute the following command to import changes in `Configuration.csv` files into the Cluster Manager VM:

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

3. Execute the following command from Cluster Manager to rebuild puppet.

   ```
   $ /var/qps/install/current/scripts/build/build_puppet.sh
   ```

4. Execute the following command from Cluster Manager to replace the puppet on all VMs.

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

MongoDB Replication Health Monitoring

CPS supports monitoring secondary members of the replica sets and if any of them lags behind the primary member it recovers automatically. To support this functionality, a new script `auto_recovery_replica.sh` is added. The following parameters can be configured.

Table 33: MongoDB Health Monitoring Parameters

Parameter	Description
auto_replica_monitor	When set to true, it enables the script for monitoring of replica sets. When set to false, it removes the script monitoring from cron. Example: auto_replica_monitor,true, Default: false Possible Values: true, false
max_replica_lag_time	(Optional) This parameter allows you to customize the maximum number of seconds a secondary replica set is allowed to lag from its primary member. For example, if the value is set to 60 that means the configuration allows all the secondary members of the replica sets to have a maximum of 60 seconds lag. By default, the maximum allowed lag is set to 30 seconds. **Note** The configured value should always be greater than 30 seconds. If you configure value less than 30 seconds, the script forces itself to select a replication lag of 30 seconds by default.
auto_replica_cron_hour	(Optional) This parameter allows you to configure the iteration in which the script for monitoring replica set has to be triggered using cron. For example, if the value is set to 5, the cron triggers the replica recovery script every 5th hour. Default: 5 hours Possible Range: 0—23 hours
auto_replica_cron_minute	(Optional) This parameter allows you to set the minute interval for the cron job. For example, if the value is set to 30, it ensures that the cron triggers this script every 30th minute. Default: 30 minutes Possible Range: 0—59 minutes

If all the parameters described in Table 33: MongoDB Health Monitoring Parameters, on page 71 are configured, the cron configuration on Cluster Manager is displayed as follows:

```
crontab -l
---------------------
# Puppet Name: Monitor replica sets
30 */5 * * * /var/qps/install/current/scripts/bin/support/mongo/auto_recovery_replica.sh
```

Note If the `auto_replica_monitor` is set to true and other parameters are not configured, the script automatically takes the default values for the remaining parameters.

By default, auto heal script runs every 1 hour via cron and the maximum lag allowed for a secondary member is 30 seconds. The default configuration adds cron entry for auto_replica_monitor on Cluster Manager.

```
crontab -l
---------------------
# Puppet Name: Monitor replica sets
0 */1 * * * /var/qps/install/current/scripts/bin/support/mongo/auto_recovery_replica.sh
```

Once the parameters are configured, you can check the configuration using the following `facter` command and then grepping for the respective values.

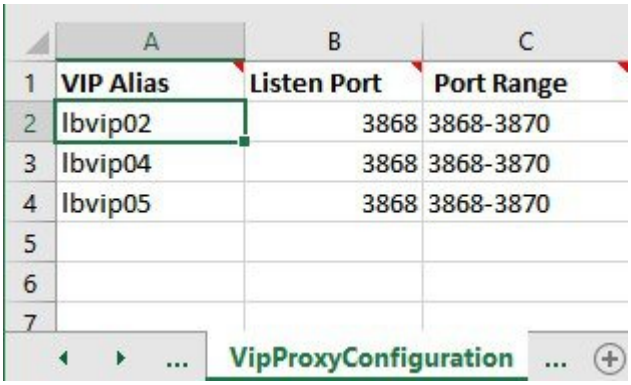
```
facter | grep auto
auto_replica_cron_hour => 1
auto_replica_cron_minute => 45
auto_replica_monitor => true
```

```
facter | grep max
max_replica_lag_time => 60
```

VIP Proxy Configuration

This file is used to specify the listen port for each VIP in HAProxy diameter configuration and the port range for the backend diameter endpoints to which the requests are load balanced. Values in this file are used to generate the HAProxy diameter configuration (`/etc/haproxy/haproxy-diameter.cfg` file) on Policy Director 01/02 VMs. Here is an example:

Figure 8: VipProxyConfiguration.csv



The following parameters can be configured in this sheet:

Table 34: VIP Proxy Configuration Parameters

Parameter	Description
VIP Alias	Name of the VIP supporting multiple diameter endpoints.

Parameter	Description
Listen Port	Front facing diameter endpoint port in HAProxy configuration.
Port Range	List of backend ports for each front end port in HAProxy configuration.

The following restriction applies to the `haproxy-diameter.cfg` file for all the installation types:

- You should not use the following list of VIP Aliases in `VipProxyConfiguration.csv` file. The VIP aliases in `AdditionalHosts.csv` invokes the legacy method of haproxy-diameter configuration. Hence, Cisco does not recommend the use of legacy VIP aliases listed below:

  diam_int1, diam_int1_vip, diam_int2, diam_int1_69, diam_int2_vip, diam_int1_69_vip, diam_int3, diam_int3_vip, diam_int1_70_vip, diam_int4, diam_int4_vip, diam_int1_71_vip

Secure Configuration

The **SecureConfig** sheet defines the Transport Layer Security (TLS) related configuration for secure services in CPS.

Select the **SecureConfig** sheet.

Figure 9: Secure Configuration Sheet

key	value
enable_tlsv1.1_pb	disabled
enable_tlsv1.1_cc	disabled
enable_tlsv1.1_uapi	disabled
enable_tlsv1.1_grafana	disabled
min_tls_pb	1.1
max_tls_pb	1.2
min_tls_cc	1.1
max_tls_cc	1.2
min_tls_uapi	1.1
max_tls_uapi	1.2
min_tls_grafana	1.1
max_tls_grafana	1.2
default_tls_grafana	1.2
default_tls_pb	1.2
default_tls_cc	1.2
default_tls_uapi	1.2

The following parameters can be configured in this sheet:

Table 35: Secure Configuration Sheet Parameters

Parameter	Description	Possible Values	Default Value
enable_tlsv1.1_pb	Enables TLSv1.1 for the Policy Builder interface.	Enabled Disabled	Disabled
enable_tlsv1.1_cc	Enables TLSv1.1 for the Control Center interface.	Enabled Disabled	Disabled
enable_tlsv1.1_uapi	Enables TLSv1.1 for the Unified API interface.	Enabled Disabled	Disabled
enable_tlsv1.1_grafana	Enables TLSv1.1 for the Grafana interface.	Enabled Disabled	Disabled
min_tls_pb	Defines the minimum TLS version supported by the Policy Builder interface.	1.1 1.2	1.1
max_tls_pb	Defines the maximum TLS version supported by the Policy Builder interface.	1.1 1.2	1.2
min_tls_cc	Defines the minimum TLS version supported by the Control Center interface.	1.1 1.2	1.1
max_tls_cc	Defines the maximum TLS version supported by the Control Center interface.	1.1 1.2	1.2
min_tls_uapi	Defines the minimum TLS version supported by the Unified API interface.	1.1 1.2	1.1
max_tls_uapi	Defines the maximum TLS version supported by the Unified API interface.	1.1 1.2	1.2
min_tls_grafana	Defines the minimum TLS version supported by the Grafana interface.	1.1 1.2	1.1
max_tls_grafana	Defines the maximum TLS version supported by the Grafana interface.	1.1 1.2	1.2
default_tls_grafana	Defines the default TLS version to use for Grafana.	1.1 1.2	1.2
default_tls_pb	Defines the Default TLS version to use for Policy Builder.	1.1 1.2	1.2

Parameter	Description	Possible Values	Default Value
default_tls_cc	Defines the default TLS version to use for Control Center.	1.1 1.2	1.2
default_tls_uapi	Defines the default TLS version to use for Unified API.	1.1 1.2	1.2

Note

- From CPS 19.1.0 release, TLSv1.1 is deprecated. By default, TLSv1.2 is supported.

  If you want to use TLSv1.1, it needs to enabled in `Secureconfig.csv` file.
- All the configuration changes are applied on the HAProxy server during **vm-init** on all Load Balancer VMs.
- For configuration parameters that are not defined in the `SecureConfig.csv` file, its logical default value is considered.
- If you enter a wrong value for any parameter, that value is discarded and the default value for that parameter is used. The Puppet log file displays a warning message.

DSCP Configuration

You can configure DSCP bits using DSCP class or DSCP value on the following for IPv4 and/or IPv6:

- Out-interface
- Protocol
- Destination IP
- Destination Port

DSCPConfig.csv format is: Role,IP Family,Out Interface,Protocol,Destination IP,Destination Port,SourcePort,DSCP Class,DSCP Value

Table 36: DSCP Configuration

Parameter	Description
Role	This parameter is used to specify the VM type. Valid values are: lb, pcrfclient, qns, sessionmgr, udc.
IP Family	This parameter is used to specify ipv4 or ipv6 address. If no parameter is configured, then the value ipv4 and ipv6 are used.
Out Interface	This parameter is used to specify the interface name i.e., eth0/eth1. If no parameter is configured, then DSCP marking is applied to any interface.

Parameter	Description
Protocol	This parameter is used to specify tcp/udp and so on. If no parameter is configured, then DSCP marking is applied to any protocol.
Destination IP	This parameter is used to specify destination IP.
Destination Port	This parameter is used to specify destination port.
Source Port	This parameter is used to specify source port.
DSCP Class	This parameter is used to specify DSCP class. Supported values are: af11, af12, af13, af21, af22, af23, af31, af32, af33, af41, af42, af43, cs1, cs2, cs3, cs4, cs5, cs6, cs7, ef
DSCP Value	This parameter is used to specify DSCP value.

`DSCPConfig.csv` file location: `/var/qps/config/deploy/csv/DSCPConfig.csv`

Example:

VM Role,IP Family,Out Interface,Protocol,Destination IP,Destination Port,Source Port,DSCP Class,DSCP Value

lb,,eth1,tcp,,27717,af11eth0,udp,,5405,,af21,,

Iptables output:

```
pkts bytes target prot opt in out  source    destination
2545K 403M DSCP   udp  --  *  eth0 0.0.0.0/0 0.0.0.0/0 multiport dports 5405 /* 100 IPv4
DSCP rules outInterface=eth0 protocol=udp destPort=5405 class=af21 */ DSCP set 0x12
```

Ip6tables output:

```
pkts bytes target     prot opt in out   source   destination
0    0     DSCP       udp      *  eth0  ::/0      ::/0 multiport dports 5405 /* 100 IPv6
DSCP rules outInterface=eth0 protocol=udp destPort=5405 class=af21 */ DSCP set 0x12
```

Critical File Monitoring Configuration

You can configure the critical file names to be monitored for write, execute or any other attribute changes.



Important Critical Files configuration is specific to Cluster Manager. If you are using Geographic Redundancy configuration, then you need to do the configuration across all the Cluster Managers.

CriticalFiles.csv format is: File To Be Monitored,Action To Be Monitored

Table 37: Critical Files Configuration

Parameter	Description
File To Be Monitored	File name with absolute path of the file that needs to be monitored.
Action To Be Monitored	Action for file that needs to be monitored. Supported options are: • w –write • x - execute • a – attribute changes

Important File monitoring for read operation is not supported.

`CriticalFiles.csv` file location: `/var/qps/config/deploy/csv/CriticalFiles.csv`

Rules configured in `CriticalFiles.csv` are added in `#BEGIN_CPS_AUDIT_RULES` and `#END_CPS_AUDIT_RULES` block in `/etc/audit/rules.d/audit.rules` file on Cluster Manager VM.

Sample output of AUDIT block in audit.rules:

```
#BEGIN_CPS_AUDIT_RULES
-w /etc/hosts -p wxa -k watch_critical_files
-w /etc/broadhop.profile -p wxa -k watch_critical_files
#END_CPS_AUDIT_RULES
```

Note Do not modify the rules in `#BEGIN_CPS_AUDIT_RULES` and `#END_CPS_AUDIT_RULES` block manually. Any modification done in this block is overwritten every time you execute `/var/qps/install/current/scripts/bin/support/update_audit_conf.py` script.

You can add the custom rules in `/etc/audit/rules.d/audit.rules` file outside of the `#BEGIN_CPS_AUDIT_RULES` and `#END_CPS_AUDIT_RULE` block but notification (SNMP trap) is not sent for the rules.

SNMP alarm with version v2c or v3 is generated based on SNMP confirmation done in `Configuration.csv` file. There is no clear alarm.

Audit daemon logs all the audit events occurred in `/var/log/audit/audit.log` file with no delay.

`/var/qps/install/current/scripts/bin/support/snmp-traps/vm-traps/gen-crit-file-mod-traps.py` script monitors audit.log file for any file modification event since last execution of script and send traps for all the events occurred during this time.

`gen-crit-file-mod-traps.py` scripts last execution time is stored in `/var/tmp/lastGenCritFileModExeTime`. If the file does not contain any entry for last execution or the file is not present, then trap for events occurred during last 60 seconds is sent.

You can execute the following command to validate particular audit logs:

```
ausearch -i -k watch_critical_files
```

Sample Output:

```
type=PROCTITLE msg=audit(08/26/2018 18:53:56.834:250) : proctitle=vim /etc/hosts
type=PATH msg=audit(08/26/2018 18:53:56.834:250) : item=1 name=/etc/hosts inode=5245468
dev=08:02 mode=file,644 ouid=root ogid=root rdev=00:00 objtype=CREATE
type=PATH msg=audit(08/26/2018 18:53:56.834:250) : item=0 name=/etc/ inode=5242881 dev=08:02
mode=dir,755 ouid=root ogid=root rdev=00:00 objtype=PARENT
type=CWD msg=audit(08/26/2018 18:53:56.834:250) : cwd=/root/modified_iso
type=SYSCALL msg=audit(08/26/2018 18:53:56.834:250) : arch=x86_64 syscall=open success=yes
exit=3 a0=0x1c74390 a1=O_WRONLY|O_CREAT|O_TRUNC a2=0644 a3=0x0 items=2 ppid=18335 pid=13946
auid=root uid=root gid=root euid=root suid=root fsuid=root egid=root sgid=root fsgid=root
tty=pts0 ses=9 comm=vim exe=/usr/bin/vim key=watch_critical_files
```

Finish and Save

After entering your deployment information, save the Deployment Template file in Excel format.

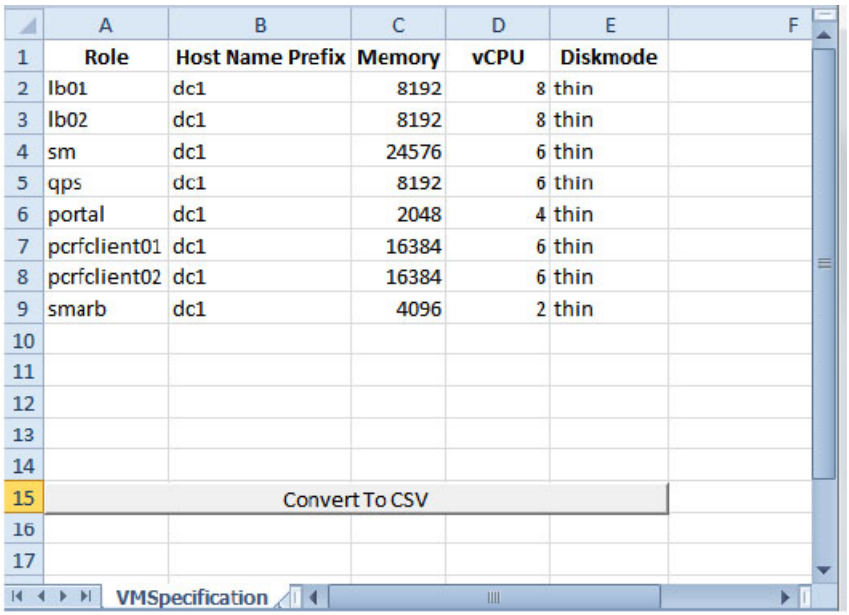
Import the Excel Information into the Cluster Manager VM

The settings in the excel sheet must be converted to a csv file and imported into CPS.

Save the csv Files

Click the **Convert to CSV** button on the VMSpecification sheet.

Figure 10: Convert To CSV



The **Convert to CSV** button exports each individual sheet into a separate CSV file in a new folder (csv_files) where the source file is located. Each csv file is named as the sheet name. Make sure the Host names, Alias, datastore, network names are all correct and created in VMware. Any mismatch of the attribute can cause the deployment to fail and restart the deployment process.

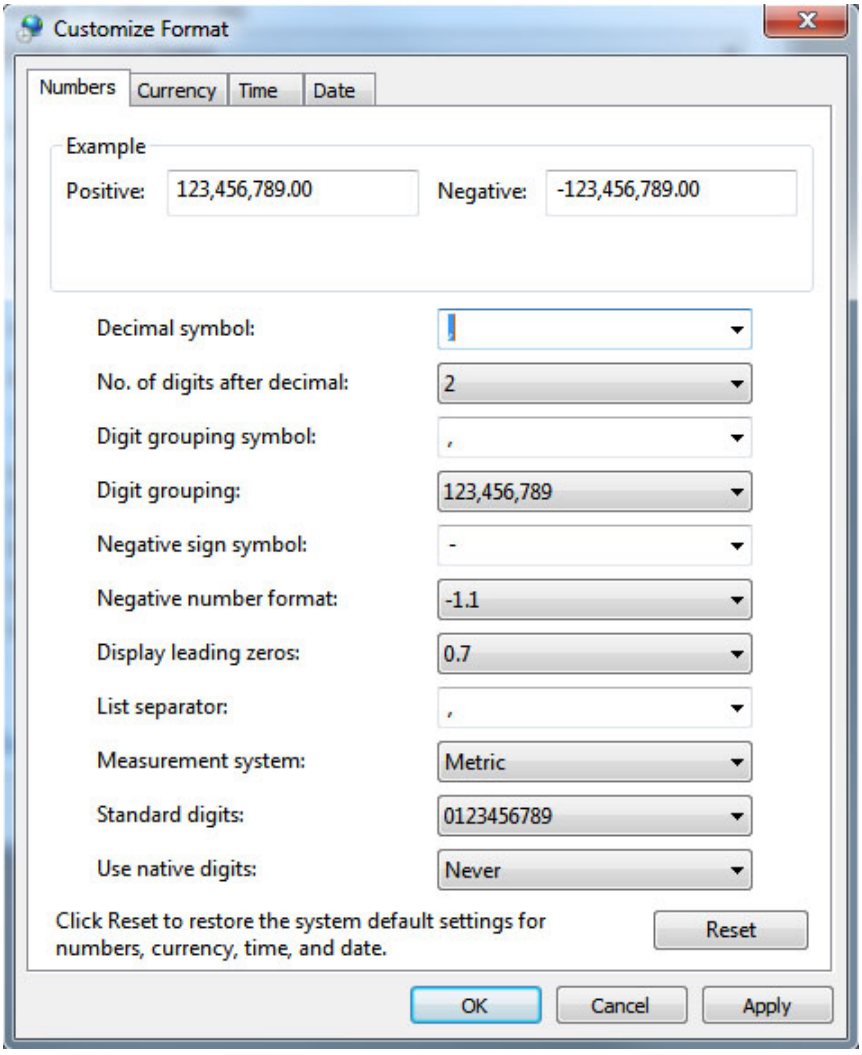
Attention It is strongly recommended to go through this list with Cisco AS and Virtualization system administrator, network administrator to make sure all the settings are correct.

The following list of csv files are generated:

- VMSpecification.csv
- Hosts.csv
- VLANs.csv
- AdditionalHosts.csv
- Configuration.csv
- Definitions.csv
- VipProxyConfiguration.csv
- SecureConfig.csv
- DSCPConfig.csv
- CriticalFiles.csv

Verify that the generated csv files are separated with commas. If needed, modify the regional settings. For reference, see the following image.

Figure 11: Regional Settings



Copy the csv Files into Cluster Manager VM

Use a tool such as Secure Copy (scp) to copy all the csv files to the Cluster Manager VM to the following directory:

```
/var/qps/config/deploy/csv/
```

Import the csv Files into the Cluster Manager VM

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

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

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

Validate Imported Data

Execute the following command to validate the imported data:

```
cd /var/qps/install/current/scripts/deployer/support/
```

```
python jvalidate.py
```

This script validates the parameters against the ESX servers to make sure ESX server can support the configuration and deploy the VMs.

Note If you are deploying the VMs using the `--nossh` feature, jvalidate will not work as jvalidate needs SSH login credentials to get the memory, CPU and network details from the ESXi. Since `--nonssh` feature doesn't have the SSH login credentials, the script execution fails for getting the details. You can login to the vCenter where the ESXi's are mapped and check the required details to plan the VM design.

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.

Continue with Customize Features in the Deployment, on page 81.

Update System Parameters

Refer to section Update the VM Configuration without Re-deploying VMs, on page 120 if you need to update any of the parameters you defined in the spreadsheet after deploying the CPS VMs.

Customize Features in the Deployment

Certain deployments require additional features to be installed. To add or remove features, perform the following steps on Cluster Manager VM:

Procedure

Step 1 Determine which features are needed with the assistance of your Cisco Technical Representative.

Step 2 If this is HA environment, edit the corresponding features files in Cluster Manager VM:

Modify the features file for the corresponding server types. Here are some examples:

```
/var/qps/current_config/etc/broadhop/controlcenter/features
```

```
# The server and infrastructure features do not need to be specified.
# IO Manager Features
com.broadhop.controlcenter.feature
com.broadhop.server.runtime.product
com.broadhop.infrastructure.feature
```

```
com.broadhop.snmp.feature
com.broadhop.faultmanagement.service.feature
```

/var/qps/current_config/etc/broadhop/diameter_endpoint/features

```
com.broadhop.server.runtime.product
com.broadhop.snmp.feature
com.broadhop.diameter2.service.feature
```

/var/qps/current_config/etc/broadhop/iomanager/features

```
# IO Manager Features
com.broadhop.iomanager.feature
com.broadhop.server.runtime.product
com.broadhop.snmp.feature
iomanager02
```

Note
In releases prior to CPS 10.0.0, there are two separate `iomanager` directories, `iomanager01` and `iomanager02`. For these older releases, changes to the iomanager features files must be populated in both directories:

/var/qps/current_config/etc/broadhop/iomanager01/features

/var/qps/current_config/etc/broadhop/iomanager02/features

/var/qps/current_config/etc/broadhop/pb/features

```
com.broadhop.client.product
com.broadhop.client.feature.ws
com.broadhop.client.feature.isg
com.broadhop.client.feature.radius
com.broadhop.client.feature.balance
com.broadhop.client.feature.spr
com.broadhop.client.feature.unifiedapi
#com.broadhop.client.feature.pop3auth
com.broadhop.client.feature.vouchers
com.broadhop.client.feature.isg.prepaid
com.broadhop.client.feature.notifications
com.broadhop.client.feature.diameter2
com.broadhop.client.feature.ldap
com.broadhop.client.feature.relianceutil
#com.broadhop.client.feature.policyintel
com.broadhop.client.feature.custrefdata
#com.broadhop.client.feature.congestionrefdata
#com.broadhop.client.feature.audit
com.broadhop.balance.crdbalance.feature
```

/var/qps/current_config/etc/broadhop/pcrf/features

```
# The server and infrastructure features do not need to be specified.
# PCRF Features
com.broadhop.server.runtime.product
com.broadhop.policy.feature
com.broadhop.externaldatacache.memcache.feature
com.broadhop.snmp.feature
com.broadhop.ws.service.feature
com.broadhop.unifiedapi.ws.service.feature
com.broadhop.spr.dao.mongo.feature
com.broadhop.spr.feature
com.broadhop.unifiedapi.interface.feature
com.broadhop.balance.service.feature
com.broadhop.vouchers.service.feature
com.broadhop.ui.controlcenter.feature
com.broadhop.diameter2.local.feature
com.broadhop.custrefdata.service.feature
```

```
com.broadhop.policyintel.service.feature
com.broadhop.balance.crdbalance.feature
```

If VMs are already deployed, after modifying the feature files, execute the following commands:

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

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

What to do next

To enable the feature **Disable Root SSH Login**, check whether there exists a user with uid 1000 on Cluster Manager.

Use the following command to check there exists a user with uid 1000:

```
cat /etc/passwd | grep x:1000
```

If a user with uid 1000 exists on the Cluster Manager, change the uid on the Cluster Manager by executing the following command:

```
usermod -u <new-uid> <user-name-with-uid-as-1000>
```

This is done because the feature **Disable Root SSH Login** creates a user with uid 1000.

LDAP Feature Installation

Enable LDAP on HA Deployment

To enable the LDAP feature on an High Availability (HA) deployment:

Procedure

Step 1 Edit the features files in Cluster Manager VM:

In the `/var/qps/current_config/etc/broadhop/pb/features` file, add the following line:

```
com.broadhop.client.feature.ldap
```

In the `/var/qps/current_config/etc/broadhop/pcrf/features` file, add the following line:

```
com.broadhop.ldap.interface.feature
```

In the `/var/qps/current_config/etc/broadhop/iomanager0X/features` file, add the following line:

```
com.broadhop.ldap.service.feature
```

Step 2 After modifying the feature files, execute the following commands from Cluster Manager:

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

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

Subscriber Lookup Feature Installation

Enable Subscriber Lookup on HA Deployment

To enable the Subscriber Lookup feature on an High Availability (HA) deployment:

Procedure

Step 1 Edit the features files in Cluster Manager VM:

In the `/var/qps/current_config/etc/broadhop/pb/features` file, add the following line:

```
com.broadhop.client.feature.ldapserver
```

In the `/var/qps/current_config/etc/broadhop/pcrf/features` file, add the following line:

```
com.broadhop.ldapserver.local.feature
```

In the `/var/qps/current_config/etc/broadhop/iomanager0X/features` file, add the following line:

```
com.broadhop.ldapserver.service.feature
```

Step 2 After modifying the feature files, execute the following commands from Cluster Manager:

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

License Generation and Installation

License Generation

For HA or GR systems, contact your Cisco Technical support representative to generate a license. You must provide the MAC addresses and hostnames for your pcrfclient01 and pcrfclient02 VMs.

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

Procedure

Step 1 To generate a unique MAC address, execute the following command on the Cluster Manager once for pcrfclient01 and again for pcrfclient02:

```
python /var/qps/install/current/scripts/deployer/support/genmac.py
```

The MAC address generated by this script is applied to pcrfclient01/02.

Important
For the pcrfclient01/pcrfclient02 VMs, the eth0 MAC address reported in the VMware Virtual Machine properties may not match what is listed in the VM's when executing the `ifconfig -a | grep HW` command output. This mismatch can be ignored. Use the MAC address displayed by `ifconfig -a | grep HW` command.

Step 2 To get the hostname, refer to the **Hosts.csv** file, and use the Guest Name that corresponds to pcrfclient01 and pcrfclient02 roles.

Step 3 Submit this information to your Cisco Technical support representative. After you receive the license, continue with License Installation, on page 85.

License Installation

The following section describes:

- How to install the license files prior to deploying all CPS VMs, as described in the Deploy the VMs, on page 95.
- The steps you perform to preserve the license files during CPS upgrade to the current release.

To install the licenses:

Procedure

Step 1 Copy the license files you received to the Cluster Manager VM.

Step 2 Create pcrfclient01 and pcrfclient02 directories in the Cluster Manager VM in `/etc/broadhop/license/`.

```
mkdir -p /etc/broadhop/license/pcrfclient01
```

```
mkdir -p /etc/broadhop/license/pcrfclient02
```

Step 3 Copy the pcrfclient01 license to the `/etc/broadhop/license/pcrfclient01` directory, and the pcrfclient02 license to the `/etc/broadhop/license/pcrfclient02` directory on the Cluster Manager VM:

`cp` <*filename1*> `/etc/broadhop/license/pcrfclient01`

`cp` <*filename2*> `/etc/broadhop/license/pcrfclient02`

where,

<*filename1*> is the license filename generated for pcrfclient01.

<*filename2*> is the license filename generated for pcrfclient02.

Step 4 If you are performing an upgrade of the system from an earlier version to the current release:

- Copy the existing pcrfclient02 license file from the pcrfclient02 VM (found in `/etc/broadhop/license`) to the `/etc/broadhop/license/pcrfclient02` directory on the Cluster Manager VM.
- During an upgrade, the license on pcrfclient01 is automatically retrieved and re-installed. Do not manually copy or move this license to the `/etc/broadhop/prcfclient01` directory on the Cluster Manager VM.

Note
As a best practice, make a backup of your existing pcrfclient01 license under `/etc/broadhop/license` on the Cluster Manager VM.

Step 5 Create a `features.properties` file in the `/etc/broadhop/license` directory on the Cluster Manager with the following content from the license file. For example:

```
LicenseFeature=POLICY-ALL,POLICY-VALUE
```

Note
The content of this file is based on the contents of the license file and your deployment.

Step 6 Execute the following command to rebuild the `/etc/broadhop/license` directory in the Cluster Manager VM.

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

This script makes a zip file with the new license file and copies it to the `/var/www/html/images` directory. Later the file is pushed to the target VMs when the `reinit.sh` script is executed.

Step 7 If pcrfclient01/pcrflient02 is already deployed, the license must be pushed to the pcrfclient01/02 VMs. For this, execute the following commands:

```
ssh pcrfclient01

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

and

```
ssh pcrfclient02

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

Note
If pcrfclient01 and pcrfclient02 VMs have not yet been deployed, the license will be automatically pushed to pcrfclient01/02 when all VMs are deployed later in section Deploy the VMs .

Step 8 If pcrfclient01/pcrflient02 is already deployed and are being updated, you must restart the LMGRD process by executing the following commands:

```
killall -9 lmgrd

service lmgrd start
```

Validate Installed License

Use the **`lmutil lmstat`** command on pcrfclient01/02 to check the status of license and list all the licenses available (Change *XXXX* to valid license file name).

Command Syntax:

```
/opt/broadhop/lmgr/x64_lsb/lmutil lmstat -a -c /etc/broadhop/license/XXXX.lic
```



Note Users of Feature-name shown is 0 in the below example (i.e. Total of 0 licenses in use). This is due to limited support for **lmgrd** from CPS side.

Example:

```
/opt/broadhop/lmgr/x64_lsb/lmutil lmstat  -a -c /etc/broadhop/license/XXXX.lic
lmutil - Copyright (c) 1989-2013 Flexera Software LLC. All Rights Reserved.
Flexible License Manager status on Fri 7/24/2015 16:11
License server status: 27000@pcrfclient01
License file(s) on pcrfclient01: /etc/broadhop/license/XXXX.lic:
pcrfclient01: license server UP (MASTER) v11.11
Vendor daemon status (on pcrfclient01):
cisco: UP v11.11
Feature usage info:
Users of SPR:  (Total of 0 licenses issued;  Total of 0 licenses in use)
Users of SP_CORE:  (Total of 0 licenses issued;  Total of 0 licenses in use)
Users of POLICY_REPORT:  (Total of 2000 licenses issued;  Total of 0 licenses in use)
Users of QUOTA:  (Total of 0 licenses issued;  Total of 0 licenses in use)
Users of Diameter_UD:  (Total of 0 licenses issued;  Total of 0 licenses in use)
Users of Diameter_SH:  (Total of 0 licenses issued;  Total of 0 licenses in use)
Users of SCE_PRPC:  (Total of 2000 licenses issued;  Total of 0 licenses in use)
Users of SCE_GY:  (Total of 2000 licenses issued;  Total of 0 licenses in use)
Users of DIAMETER_SD:  (Total of 2000 licenses issued;  Total of 0 licenses in use)
```

Upgrade License

User needs to upgrade license if the current licenses have expired or if you need to increase the session capacity of the system.

Procedure

Step 1 Contact your Cisco Technical representative to generate a license. You must provide the MAC addresses and hostnames for your pcrfclient01 and pcrfclient02 VMs.

Step 2 Copy the license files you received to the Cluster Manager VM.

Step 3 Delete the existing license files from the Cluster Manager VM.

`rm -fr /etc/broadhop/license/pcrfclient01/`<*filename1*>

`rm -fr /etc/broadhop/license/pcrfclient02/`<*filename2*>

where,

<*filename1*> is the existing license filename of pcrfclient01.

<*filename2*> is the existing license filename of pcrfclient02.

Note
As a best practice, create a backup of your existing licenses.

Step 4 Copy the pcrfclient01 license to the `/etc/broadhop/license/pcrfclient01` directory, and the pcrfclient02 license to the `/etc/broadhop/license/pcrfclient02` directory on the Cluster Manager VM:

`cp` <*filename1*> `/etc/broadhop/license/pcrfclient01`

`cp` <*filename2*> `/etc/broadhop/license/pcrfclient02`

where,

<*filename1*> is the license filename generated for pcrfclient01.

<*filename2*> is the license filename generated for pcrfclient02.

Step 5 Add the necessary content as per your license files in `features.properties` in `/etc/broadhop/license` directory on the Cluster Manager. For example:

```
LicenseFeature=POLICY-ALL,POLICY-VALUE
```

Step 6 Execute the following command to rebuild the `/etc/broadhop/license` directory in the Cluster Manager VM.

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

This script makes a zip file with the new license file and copies it to the `/var/www/html/images` directory.

Later the file is pushed to the target VMs when the `vm-init.sh` script is executed.

Step 7 To push new license to the pcrfclient01/02 VMs, restart the LMGRD process, and restart the Policy Server (QNS), execute the following commands.

```
ssh pcrfclient01 "/etc/init.d/vm-init && killall -9 lmgrd; service lmgrd start && systemctl restart qns-1"
```

```
ssh pcrfclient02 "/etc/init.d/vm-init && killall -9 lmgrd; service lmgrd start && systemctl restart qns-1"
```

Step 8 Validate installed license, refer to Validate Installed License, on page 86.

SSL Certificates

Default SSL cipher supported:

```
ciphers ECDH+AESGCM:DH+AESGCM:ECDH+AES256:DH+AES256:ECDH+AES128:DH+AES:RSA+AESGCM:RSA+AES:
```

For more information, refer to https://www.openssl.org/docs/man1.0.2/apps/x509.html.

Create SSL Certificates

Certain deployments have customized certificates (for example, *.der and *.cer files) installed in their systems. To create Self-Signed certificates (SSL) that can be used in CPS, perform the following steps on Cluster Manager VM:

Procedure

Step 1 Convert the user provided files to pem files. Consider the user has provided *.der and *.cer files.

For example, if the user has provided the following files:

- x.der: Server certificate
- y.cer: ROOT CA certificate file
- z.cer: Intermediate file

```
openssl x509 -inform der -in x.der -out x.pem
```

```
openssl x509 -inform der -in y.cer -out y.pem
```

```
openssl x509 -inform der -in z.cer -out z.pem
```

Note
For details on how to generate certificates using openssl, refer to: https://www.openssl.org/docs/man1.0.2/apps/x509.html.

Step 2 Generate your chain crt file in the following order: server > intermediate > root.

```
cat x.pem z.pem y.pem > server.crt
```

Step 3 Remove passphrase from KEY file (You will be asked to supply the passphrase of the KEY file).

```
openssl rsa -in server.key -out server.nopass.key
```

Step 4 Combine the key without pass and certificate chain to create pem file.

```
cat server.nopass.key server.crt > server.pem
```

Step 5 Copy the server.pem, server.crt and server.nopass.key to `/var/qps/install/current/puppet/modules/qps/templates/certs/`.

```
cp server.crt /var/qps/install/current/puppet/modules/qps/templates/certs/quantum.crt
cp server.nopass.key /var/qps/install/current/puppet/modules/qps/templates/certs/quantum.key
cp server.pem /var/qps/install/current/puppet/modules/qps/templates/certs/quantum.pem
```

Step 6 Execute the following commands from Cluster Manager:

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

Replace SSL Certificates

To replace the default Self-Signed certificates (SSL) during installation process, replace the crt, key and pem (contains both the crt/key) in `/var/qps/install/current/puppet/modules/qps/templates/certs` directory on the Cluster Manager VM with the new certificates.



Important The custom certificates are replaced with the default CPS certificates after the migration or upgrade. In this case, you need to apply the custom certificates again on Cluster Manager once the upgrade or migration is complete.

Consider the user has the following new SSL certificates and wants to replace the default SSL certificates in the system:

- SSL_new.crt
- SSL_new.key
- SSL_new.pem

The new certificates can be stored anywhere on the Cluster Manager. In the following steps the new certificates are stored in `/root`. To replace the old keys/certs/pem with the new ones, perform the following steps:

Procedure

Step 1 Execute the following commands from Cluster Manager to replace the old certificates with the new certificates:

```
mv /root/SSL_new.crt /var/qps/install/current/puppet/modules/qps/templates/certs/quantum.crt
mv /root/SSL_new.key /var/qps/install/current/puppet/modules/qps/templates/certs/quantum.key
mv /root/SSL_new.pem /var/qps/install/current/puppet/modules/qps/templates/certs/quantum.pem
```

Important
Retain the permissions of the old files.

Step 2 Execute the following command from Cluster Manager to rebuild puppet:

```
build_puppet.sh
```

Step 3 Execute the following command from each VM to replace the certs/keys:

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

OR

Execute the following command from Cluster Manager to replace the puppet on all VMs.

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

Enable Custom Puppet to Configure Deployment

Some customers may need to customize the configuration for their deployment. When customizing the CPS configuration, it is important to make the customization in a way that does not impact the normal behavior for VM deployment and redeployment, upgrades/migration, and rollbacks.

For this reason, customizations should be placed in the `/etc/puppet/env_config` directory. Files within this directory are given special treatment for VM deployment, upgrade, migrations, and rollback operations.



Note If system configurations are manually changed in the VM itself after the VM has been deployed, these configurations will be overridden if that VM is redeployed.

The following section describes the steps necessary to make changes to the puppet installer.

Customizations of the CPS deployment are dependent on the requirements of the change. Examples of customizations include:

- deploying a specific facility on a node (VM)
- overriding a default configuration.

To explain the process, let us consider that we modify all VMs built from an installer, so we use the Policy Server (QNS) node definition.

For the above mentioned example, add custom routes via the examples42-network Puppet module. (For more information on the module, refer to https://forge.puppetlabs.com/example42/network).

Attention In CPS 20.2.0, puppet is upgraded from 3.6.2-3 to 5.5.19 version. Puppet code has been modified to adapt to this change. Previous release puppet code is not compatible with the current puppet version (5.5.19). Customer specific puppet code must be adapted to current release puppet version (5.5.19) before applying it to CPS 20.2.0.

Procedure

Step 1 Make sure that the proper paths are available:

```
mkdir -p /etc/puppet/env_config/nodes
```

Step 2 Install the necessary Puppet module. For example:

```
puppet module install \
--modulepath=/etc/puppet/env_config/modules:/etc/puppet/modules \
example42-network
Notice: Preparing to install into /etc/puppet/env_config/modules ...
Notice: Downloading from https://forge.puppetlabs.com ...
Notice: Installing -- do not interrupt ...
/etc/puppet/env_config/modules
example42-network (v3.1.13)
```

Note
For more information on installing and updating Puppet modules, refer to https://docs.puppetlabs.com/puppet/latest/reference/modules_installing.html.

Step 3 Copy the existing node definition into the env_config nodes:

```
cp /etc/puppet/modules/qps/nodes/qps.yaml \

/etc/puppet/env_config/nodes
```

Step 4 Add a reference to your custom Puppet manifest:

```
echo ' custom::static_routes:' >> \

/etc/puppet/env_config/nodes/qps.yaml
```

Step 5 Create your new manifest for static routes:

```
cat
 >/etc/puppet/env_config/modules/custom/manifests/static_routes.pp <<EOF class custom::static_routes
 {
   network::route {'eth0':
     ipaddress => ['192.168.1.0',],
     netmask   => ['255.255.255.0',],
     gateway   => ['10.105.94.1',],
   }
}
EOF
```

Step 6 Validate the syntax of your newly created puppet script(s):

```
puppet parser validate
/etc/puppet/env_config/modules/custom/manifests/static_routes.pp
```

Step 7 Rebuild your Environment Configuration:

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

Step 8 Reinitialize your environment:

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

At this point your new manifest is applied across the deployment. For more details, refer to the installer image in the `/etc/puppet/env_config/README`.

What to do next

It is recommended that version control is used to track changes to these Puppet customizations.

For example, to use 'git', perform the following steps:

1. Initialize the directory as a repository:

   ```
   # git init
   ```

   Initialized empty Git repository in `/var/qps/env_config/.git/`.

2. Add everything:

   ```
   # git add .
   ```

3. Commit your initial check-in:

   ```
   # git commit -m 'initial commit of env_config'
   ```

4. If you are making more changes and customizations, make sure you create new revisions for those:

   ```
   # git add .
   # git commit -m 'updated static routes'
   ```

Installing Platform Scripts for MongoDB Health Monitoring - VMware

The following steps are performed to install platform scripts for MongoDB health monitoring for write operations on VMware setup.

Procedure

Step 1 Log in to the Cluster Manager or installer as a root user.

Step 2 Update `/var/qps/config/deploy/csv/Configuration.csv` file by adding following entry and save the file:

```
autoheal_qns_enabled,TRUE
```

Step 3 Execute the following scripts to make sure the changes are applied on all the required VMs.

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

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

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

Step 4 Execute the following command to validate if the parameter is applied.

```
for hn in `hosts.sh`; do echo $hn ; ssh $hn "grep autoheal_qns_enabled
/etc/facter/facts.d/qps_facts.txt"; echo; done
```

Sample Output when parameter is configured:

```
[root@installer ~]# for hn in `hosts.sh`; do echo $hn ; ssh $hn "grep autoheal_qns_enabled
/etc/facter/facts.d/qps_facts.txt"; echo; done
lb01
autoheal_qns_enabled=TRUE

lb02
autoheal_qns_enabled=TRUE

qns01
autoheal_qns_enabled=TRUE

qns02
autoheal_qns_enabled=TRUE

pcrfclient01
autoheal_qns_enabled=TRUE

pcrfclient02
autoheal_qns_enabled=TRUE
```

Step 5 Execute the following steps on each Policy Server (QNS) VMs.

a) Log in as a root user.

b) Edit `crontab` using the following command.

```
crontab –e
```

The **vi editor** page opens.

Note
Type the command on the terminal and do not copy and paste values on the terminal.

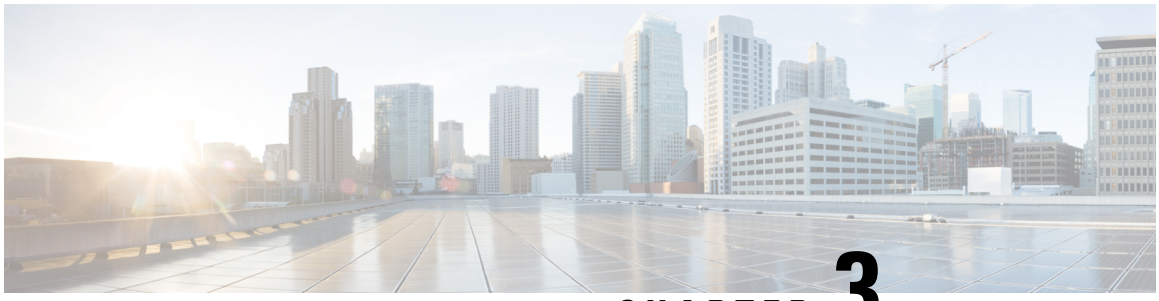
c) Add the following line in the opened **vi editor**.

```
* * * * * /var/qps/bin/support/app_mon/app_monitor.sh
```

d) Save the file and exit the editor.

Note

If any change or upgrade is performed, make sure the cronjob entry is present. If the entry is not present repeat the above steps to configure cronjob.



CHAPTER 3

Deploy CPS VMs

- Deploy the VMs, on page 95
- Update Default Credentials, on page 98
- Initialize SVN Synchronization, on page 99
- External Port Matrix, on page 100
- Memory Reservation on VMs, on page 100
- Session Manager Configuration for Data Replication, on page 100
- Validate VM Deployment, on page 111

Deploy the VMs

If there are large number of VMs in your CPS deployment it is recommended to perform a Manual Deployment for one VM (for test purposes). After the success of the first VM, then all VMs can be deployed using Automatic Deployment process.

Note During the VM deployment, do not perform any vCenter operations on the blades and VMs installed on them.

Build VM Images

Before deploying the VMs, build the VM images by executing the following command from the Cluster Manager VM:

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

Sample Output

```
Building /etc/broadhop...
Copying to /var/qps/images/etc.tar.gz...
...
Copying wispr.war to /var/qps/images/wispr.war
Output images to /var/qps/images/
[root@hostname]#
```

Manual Deployment

This section describes the steps to deploy each VM in the CPS deployment individually. To deploy all of the VMs in parallel using a single command refer to Automatic Deployment of All CPS VMs in Parallel, on page 96. To deploy a selective list of VMs in parallel using a single command refer to Automatic Deployment of Selective CPS VMs in Parallel, on page 97.

Note Before proceeding, refer to License Generation and Installation, on page 84 to confirm you have installed the license correctly.

For each host that is defined in the Hosts tab of the CPS Deployment Template spreadsheet execute the following:

Note The following command uses the short alias name (qns01 qns02 etc.) as defined in the Hosts tab of the CPS Deployment Template. It will not work if you enter the full hostname.

```
/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 qns01` <=== passed

`./deploy.sh NDC2BSND2QNS01` <=== failed

Important Newly deployed VM/VMs need to be shutdown cleanly and started with your preferred method to reserve memory:

1. To shutdown individual VM:

   ```
   cd /var/qps/install/current/scripts/deployer
   ./deploy.sh <vm alias> --shutdownvm
   ```

2. Start the VM:

   ```
   ./deploy.sh <vm alias> --poweronvm
   ```

Automatic Deployment of All CPS VMs in Parallel

This section describes the steps to deploy all VMs in parallel in the CPS deployment.

Note Before proceeding, refer to *License Generation and Installation* to confirm you have installed the license correctly.

Execute the following command:

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

The order in which VMs are deployed is managed internally.



Note The amount of time needed to complete the entire deployment process depends on the number of VMs being deployed as well as the hardware on which it is being deployed.

The following is a sample list of VM hosts deployed. The list varies according to the type of CPS deployment as well as the information you entered in the CPS Deployment Template.

- pcrfclient01
- pcrfclient02
- sessionmgr01
- sessionmgr02
- lb01
- lb02
- qns01
- qns02
- qns03
- qns04



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, on page 96.

Automatic Deployment of Selective CPS VMs in Parallel

This section describes the steps to deploy a selective list of VMs in parallel in the CPS deployment.



Note Before proceeding, refer to *License Generation and Installation* to confirm you have installed the license correctly.

Execute the following command:

```
python /var/qps/install/current/scripts/deployer/support/deploy_all.py --vms <filename-of-vms>
```

where, *<filename-of-vms>* is the name of the file containing the list of VMs such as:

```
pcrfclient01
lb01
```

```
qns01
```

Note The amount of time needed to complete the entire deployment process depends on the number of VMs being deployed as well as the hardware on which it is being deployed.

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
```

Important Newly deployed VM/VMs need to be shutdown cleanly and started with your preferred method to reserve memory:

Shut down and start Selective CPS VMs in Parallel.

1. Use your preferred editor and create `/tmp/vm-list` file and add VMs which you want to shut down and start.
2. To shutdown VMs from the given list.

```
cd /var/qps/install/current/scripts/deployer/support
python deploy_all.py --vms /tmp/vm-list --poweroffvm
```

Note Make sure that all the VMs in the list are powered OFF by using the above command.

3. To start all the VMs in list.

```
python deploy_all.py --vms /tmp/vm-list --poweronvm
```

Update Default Credentials

The passwords for the users in an HA or GR deployment are not set by default. Before you can access the deployed VMs or CPS web interfaces, you must set these passwords.

Procedure

Step 1 Log into the Cluster Manager VM as the `root` user. The default credentials are **`root/CpS!^246`**.

Step 2 Execute the `change_passwd.sh` script to set the password.

Note
`change_passwd.sh` script can also be used to change the root user password on all VMs including Cluster Manager VM.

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

Note
The `change_passwd.sh` script changes the password on all the VMs temporarily. You also need to generate an encrypted password. 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.

Step 3 When prompted, enter qns.

```
Enter username whose password needs to be changed: qns
```

Step 4 If password not set, then below user message will be displayed:

Currently password is not set, please change the password

If password exists, then user will be prompted for the current password

Enter current password:

Note
You can create passphrase or password with the following limitations, when you create or change passwords:

- You can provide one uppercase letter, one lowercase letter, one digit and one special character
- You can provide minimum 8 characters length password and up to maximum password length 256.
- Default password expiry is set to 6 months.

Step 5 When prompted, enter and reconfirm the desired password for the **qns** user.

```
Enter new password:
Re-enter new password:
Changing password on $host...
Connection to $host closed.
Password for qns changed successfully on $host
```

Note
If script prompts for `[installer] Login password for 'root'`:, enter default password (**`CpS!^246`**).

Step 6 Repeat Step 2, on page 98 to Step 5, on page 99 to set or change the passwords for **`root`** and **`qns-svn`** users.

For more information about this and other CPS administrative commands, refer to the *CPS Operations Guide*.

Initialize SVN Synchronization

After the VMs are deployed, execute the following script from the pcrfclient01 VM:

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

This command synchronizes the master/slave Policy Builder subversion repositories.

External Port Matrix

The following table lists the services and ports that CPS makes available to external users and applications. It is recommended that connectivity to these ports be granted from the appropriate networks that require access to the below services.

Table 38: External Port Matrix

Service	Common Port (For HA Environment)	Deprecated Port (For HA Environment)
Control Center	443	443
Policy Builder	443	7443
Grafana	443	9443
Unified API	443	8443
Custom Reference Data REST API	443	8443
HAProxy Status	5540	5540

For a full list of ports used for various services in CPS, refer to the *CPS Architecture Guide*, which is available by request from your Cisco Representative.

Memory Reservation on VMs

To avoid performance impact you must reserve all allocated memory to each CPS virtual machine. For more information, refer to Reserving Memory on the Virtual Machines (VMs), on page 121.

Session Manager Configuration for Data Replication

Before you perform service configuration, configure the session managers in the cluster. The database must be up and running for the CPS software.

Note Perform the steps mentioned in the following sections from the Cluster Manager.

Guidelines for Choosing MongoDB Ports

The standard definition for supported replica-set is defined in the `mongoConfig.cfg` file.

Use the `/etc/broadhop/ha_mongoconfig_template` file to create the `/etc/broadhop/mongoConfig.cfg` and modify it to your requirements.

Note If you are using VIP for arbiter, it is always recommended to keep VIP and all mongod processes on pcrfclient02 (by default).

Consider the following guidelines for choosing MongoDB ports for replica-sets:

- Port must not be in use by any other application. To check whether the port is in use, login to VM on which replica-set is to be created and execute the following command:

  `netstat -lnp | grep <port_no>`

  If no process is using same port, then port can be chosen for replica-set for binding.

- Port number used should be greater than 1024 and not in ephemeral port range i.e, not in between following range :

  net.ipv4.ip_local_port_range = 32768 to 61000

- While configuring MongoDB ports in a geographically redundant environment, there should be a difference of 100 ports between two respective sites. For example, consider there are two sites: Site1 and Site2. For Site1, if the port number used is 27717, then you can configure 27817 as the port number for Site2. This is helpful to identify a MongoDB member's site. By looking at first three digits, you can decide where the MongoDB member belongs to. However, this is just a guideline. You must avoid having MongoDB ports of two different sites to close to each other (for example, 27717 on Site-1 and 27718 on Site2).

  Reason: The `build_set.sh` script fails when you create shards on the site (for example, Site1). This is because the script calculates the highest port number in the `mongoConfig` on the site where you are creating shards. This creates a clash between the replica-sets on both sites because the port number which it allocates might overlap with the port number of `mongoConfig` on other site (for example, Site2). This is the reason why there should be some gap in the port numbers allocated between both the sites.

Supported Databases

The replica-set script is used to create replica-sets for the following databases. For more information about the script, see Script Usage, on page 103.

- session
- spr
- balance
- report
- audit
- admin

Prerequisites

- It is recommended to use the specific option for creating a single replica-set rather than `--all` option as it is easy to recreate it again if it fails to create.

- If recreating a replica-set on a production system, make sure to back up the database (Refer *CPS Backup and Restore Guide*).
- Auto Intelligent DB Operations (AIDO) server is running on Cluster Manager or third-party site Arbiter.
  - It is not active on third-party site Arbiter node, i.e., using `monit summary` you can see aido_server is running but in `/var/log/aido_server.log` you can see the following message:

    ```
    AIDO server is not needed on arbiter/site
    ```

  - It pushes latest or updated `mongoConfig.cfg` file to all database members every 60 seconds interval.
  - It checks if any database member is UP and ready to join a replica-set. If Yes, then checks whether replica-set exist or not. If replica-set exists, then join as a member in the existing replica-set. If replica-set does not exist, then create new replica sets
  - Monit process name is `aido_server`.
  - AIDO server status can be checked by using `/etc/init.d/aido_server status` and `systemctl status aido_server`
  - Log rotate file is available at: `/etc/logrotate.d/aido_server`, size limit is 10 M and 5 rotation
- AIDO client is running on sessionmgr, pcrfclient and third-party site Arbiter.
  - `mongoConfig.cfg` file is received from AIDO servers (in GR, multiple AIDO servers are available).

    `mongoConfig.cfg` file is available at: `/var/aido`

    File name format is:
    `/var/aido/mongoConfig.cfg.<<cluman-host-name>>-<<--cluman-eth0-IP-->>`

    AIDO server pushes `mongoConfig.cfg` file to all database members i.e., AIDO clients.
  - AIDO client status can be checked by using `/etc/init.d/aido_client status` and `systemctl status aido_client`
  - Log rotate file is availabe at: `/etc/logrotate.d/aido_client`, size limit is 10 M and 5 rotation

Note You have to refer to `/etc/broadhop/ha_mongoconfig_template` file and use this file to create `/etc/broadhop/mongoConfig.cfg` file based on your requirements.

All the replica set members and required information like Host Name and port number arbiter host name and port number should be defined in `/etc/broadhop/mongoConfig.cfg` file.

Note Make sure all the replica set ports defined in the `mongoConfig.cfg` file are outside the range 32768 to 61000. For more information about the port range, refer to http://www.ncftp.com/ncftpd/doc/misc/ephemeral_ports.html.

The following example shows replica-set set04:

Table 39: Replica-set Example

[SPR-SET1]	[Beginning Set Name-Set No]
SETNAME=rep_set04	Set name i.e. rep_set04
ARBITER1=pcrfclient0127720	Arbiter VM host with port number
ARBITER_DATA_PATH=/var/data/sessions.4	Arbiter data directory
MEMBER1=sessionmgr0127720	Primary Site Member1
MEMBER2=sessionmgr0227720	Primary Site Member2
DATA_PATH=/var/data/sessions.4	Data Directory Path for members
[SPR-SET1-END]	[Closing Set Name-Set No]

Run the `/var/qps/install/current/scripts/build/build_etc.sh` script from the Cluster Manager to finalize `mongoConfig.cfg` file after AIDO automatically takes care of updating it.

`build_set.sh` script copies `/etc/broadhop/mongoConfig.cfg` file to `/var/www/html/images/mongoConfig.cfg` file.

Script Usage

`build_set.sh` script is used to verify replica-set creation.

Option to view help: `/var/qps/bin/support/mongo/build_set.sh --help`

```
build_set.sh --help


              Replica-set Configuration
--------------------------------------------------------------------------------

Usage: build_set.sh <--option1> <--option2> [--setname SETNAME] [--help]
option1: Database name
option2: Build operations (create, add or remove members)
option3: Use --setname SETNAME to build or alter a specific replica-set
         replica-set setnames are defined in the /etc/broadhop/mongoConfig.cfg file

The script applies to Database: session, spr, balance, report, portal, admin, audit and
bindings db replica-sets
                  Config Server: session_configs, spr_configs and bindings_configs db
replica-sets

--all                     : Alias for all databases in the configuration
--create                  : Create a replica-set if force option is given, else it just
validate
--create-asc              : Create a replica-set with set priority in the ascending format
 if
                            force option is given, else it just validate
--create-des             : Create a replica-set with set priority in the descending format
 if
                            force option is given, else it just validate
--add-members             : Add members to a replica-set if force option is given, else it
```

```
 just validate
                            This applies to members which have been removed from the
replica-set using the
                            --remove-members and --remove-failed-members operations
--remove-members          : Remove specific members from a replica-set
                            For example, a non-active member
--remove-failed-members   : Remove failed/not reachable members from a replica-set
                           On occasion, replica-set members are not reachable due to network
 isues
--remove-replica-set      : Remove a replica-set
--create-scripts         : Create init.d script for the replica-set members if force option
 is given
--setname                 : The name of a replica-set as configured in
/etc/broadhop/mongoConfig.cfg
--force                   : This option can be used with create & add-members

Examples:
  General operation

    build_set.sh --all --create
    build_set.sh --session --create
    build_set.sh --session --create-asc
    build_set.sh --session --create-des
    build_set.sh --session --add-members
    build_set.sh --session --remove-members
    build_set.sh --session --remove-failed-members
    build_set.sh --session --remove-replica-set
    build_set.sh --session --create-scripts
    build_set.sh --help

  To perform build operations on a sepecific replica-set:

    build_set.sh --spr --create --setname set04
    build_set.sh --spr --create-asc --setname set04
    build_set.sh --spr --create-des --setname set04
    build_set.sh --spr --add-members --setname set04
    build_set.sh --spr --remove-failed-members --setname set04
    build_set.sh --spr --remove-replica-set --setname set04
    build_set.sh --spr --create-scripts --setname set04
```

If you want to use build_set.sh to create replica-set then use option --force.

Note When you execute execute `build_set.sh <databasename> --remove_replica_set <setname>`, it creates a `/var/tmp/stopped-XXXX` (*XXXX* is the port number of the replica-set member) file on the respective sessionmgr and arbiter VM. If you want to recreate the same replica-set again on the same port then you have to manually remove the `/var/tmp/stopped-XXXX` file from respective sessionmgr or arbiter VM. AIDO monitors the `/var/tmp/stopped-XXXX` file on the VM and don't do any action on the replica-set member if file is present.

Guidelines for Adding Replica-sets

You must create the database replica-set members on the same VM and the same port on both sites.

For example: For session manager database, among four replica-set members (except arbiter), if `sessionmgr01:27717` and `sessionmgr02:27717` are two members of replica-set from SITE1, then choose `sessionmgr01:27717` and `sessionmgr02:27717` of SITE2 as other two replica-set members as shown in following example:

```
[SESSION-SET]
        SETNAME=set01
        OPLOG_SIZE=5120
        ARBITER1=SITE-ARB-sessionmgr05:27717
        ARBITER_DATA_PATH=/var/data/sessions.1/set1
        PRIMARY-MEMBERS
        MEMBER1=SITE1-sessionmgr01:27717
        MEMBER2=SITE1-sessionmgr02:27717
        SECONDARY-MEMBERS
        MEMBER1=SITE2-sessionmgr01:27717
        MEMBER2=SITE2-sessionmgr02:27717
        DATA_PATH=/var/data/sessions.1/set1
        [SESSION-SET-END]
```

Defining a Replica-set

Procedure

Step 1 Update the `mongoConfig.cfg` file with the new replica-set.

Step 2 Execute the following command from the Cluster Manager to finalize `mongoConfig.cfg` file after AIDO automatically takes care of updating it:

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

Step 3 To verify replica-set has been created, run the `build_set.sh` command for the different replica-sets. The following table describes the commands for each type of replica set:

Table 40: Replica-set Commands

Replica-set	Command
Session Replica-set	`/var/qps/bin/support/mongo/build_set.sh --session`
SPR Replica-set	SPR (USum) supports MongoDB hashed sharding. `/var/qps/bin/support/mongo/build_set.sh --spr`
Balance Replica-set	`/var/qps/bin/support/mongo/build_set.sh --balance`
Report Replica-set	`/var/qps/bin/support/mongo/build_set.sh --report`
Audit Replica-set	`/var/qps/bin/support/mongo/build_set.sh --audit`
Admin Replica-set	The ADMIN database holds information related to licensing, diameter end-points and sharding for runtime. use. `/var/qps/bin/support/mongo/build_set.sh --admin`

Replica-set	Command
WT_CACHESIZEGB	This parameter configures **wiredtiger** cache in GB on Session Manager VMs. The configured **WT_CACHESIZEGB** reflects in mongo processes as **--wiredTigerCacheSizeGB** parameter. This is an optional parameter. Default value: 2 GB **Note** Starting from CPS 22.1.1 release, MongoDB Storage Engine is changed from **MMAPv1** to **WiredTiger**. WiredTigerstorage engine change in MongoDB server requires an additional CPU resources of ~15% and additional memory (RAM) resources of ~40% in the Session Manager VMs. WiredTiger consumes up to ~40% extra memory from the total memory(RAM) than MMapV1. For example, if the sessionmgr VM (150GB) with MMapV1 uses 60GB, then WiredTiger requires 120GB (MMapV1 usage 60GB + 40% of total memory). As per mongo documentation, the wiredtigercachegb can be configured as [50% of (RAM - 1 GB)] in the VM. If **n** mongo processes are running in the VM, the wiredtigercachegb can be configured as [50% of (RAM - 1 GB)]/n per mongo process. For example, in the setup: • Sessionmgr VMs configured RAM: 157GB • The number of mongo processes will be running on VM: 6 • Each process cache size can be configured : [50% of (157GB-1GB)]/6 ==> 78/6 = 13GB(can rounded to 12 GB) **Note** OS can consume 40-50GB of buffer/cache memory towards system/kernel operations. The given values must be configured in mongoConfig.cfg: • WT_CACHESIZEGB=12 • WT_CACHEARBSIZEGB=1

Replica-set	Command
WT_CACHEARBSIZEGB	This parameter configures **wiredtiger** cache in GB on arbiter VMs. The configured **WT_CACHEARBSIZEGB** will be reflected in mongo processes **--wiredTigerCacheSizeGB** parameter. This is an optional parameter. Default value: 1 GB

Instead of the specific command described in table, you can also use the following command:

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

Note
The installation logs are generated in the appropriate directories (`/var/log/broadhop/scripts/`) for debugging or troubleshooting purposes.

Example of Replica set Creation

Here are some examples for replica-sets:

Procedure

Step 1 Log in to Cluster Manager.

Step 2 Refer to `/etc/broadhop/ha_mongoconfig_template` file and use this file to create `/etc/broadhop/mongoConfig.cfg` file based on your requirements.

```
vi /etc/broadhop/mongoConfig.cfg
```

```
[SESSION-SET1]
SETNAME=set01
OPLOG_SIZE=1024
WT_CACHESIZEGB=2
WT_CACHEARBSIZEGB=1
ARBITER=pcrfclient01:27717
ARBITER_DATA_PATH=/var/data/sessions.1
MEMBER1=sessionmgr01:27717
MEMBER2=sessionmgr02:27717
DATA_PATH=/var/data/sessions.1
[SESSION-SET1-END]

[BALANCE-SET1]
SETNAME=set02
OPLOG_SIZE=1024
WT_CACHESIZEGB=2
WT_CACHEARBSIZEGB=1
ARBITER=pcrfclient01:27718
ARBITER_DATA_PATH=/var/data/sessions.2
MEMBER1=sessionmgr01:27718
MEMBER2=sessionmgr02:27718
DATA_PATH=/var/data/sessions.2
[BALANCE-SET1-END]
```

```
[REPORTING-SET1]
SETNAME=set03
OPLOG_SIZE=1024
WT_CACHESIZEGB=2
WT_CACHEARBSIZEGB=1
ARBITER=pcrfclient01:27719
ARBITER_DATA_PATH=/var/data/sessions.3
MEMBER1=sessionmgr01:27719
MEMBER2=sessionmgr02:27719
DATA_PATH=/var/data/sessions.3
[REPORTING-SET1-END]

[SPR-SET1]
SETNAME=set04
OPLOG_SIZE=1024
WT_CACHESIZEGB=2
WT_CACHEARBSIZEGB=1
ARBITER=pcrfclient01:27720
ARBITER_DATA_PATH=/var/data/sessions.4
MEMBER1=sessionmgr01:27720
MEMBER2=sessionmgr02:27720
DATA_PATH=/var/data/sessions.4
[SPR-SET1-END]
```

Step 3 After defining the admin database details, rebuild `etc.tar.gz`.

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

What to do next

After replica sets are created, you need to configure the priorities for the replica set members using `set_priority.sh` command. For more information on `set_priority.sh`, refer to *CPS Operations Guide*.

Guidelines to Configure More than Seven Replica-set Members

If it is required to configure more than seven members (including arbiters), then data members must be defined as non-voting-members in `/etc/broadhop/mongoConfig.cfg` file.

Non-voting members allow you to add additional data members for read distribution beyond the maximum seven voting members.

To configure a member as non-voting, votes and priority value must be configured to 0.

This configuration is done by `build_set.sh` and `set_priority.sh` scripts. So, it is expected to have priority as 0 for non-voting-member.

For more information, see https://docs.mongodb.com/manual/tutorial/configure-a-non-voting-replica-set-member/ (select appropriate mongo version).

Configure Non-Voting Members

If there are total eight data members and one arbiter (i.e. total nine members), six must be defined as MEMBER*n* and all other remaining data members must be defined as NON-VOTING-MEMBER*n* in `/etc/broadhop/mongoConfig.cfg` file.

where, *n* in MEMBER*n* and NON-VOTING-MEMBER*n* represents number 1, 2, 3 and so on.

```
[SPR-SET1]
SETNAME=set04
OPLOG_SIZE=3072
ARBITER=site3-arbiter:27720
ARBITER_DATA_PATH=/var/data/sessions.4
PRIMARY-MEMBERS
MEMBER1=site1-sessionmgr01:27720
MEMBER2=site1-sessionmgr02:27720
MEMBER3=site1-sessionmgr03:27720
NON-VOTING-MEMBER4=site1-sessionmgr04:27720
SECONDARY-MEMBERS
MEMBER1=site2-sessionmgr01:27720
MEMBER2=site2-sessionmgr02:27720
MEMBER3=site2-sessionmgr03:27720
NON-VOTING-MEMBER4=site2-sessionmgr04:27720
DATA_PATH=/var/data/sessions.4
[SPR-SET1-END]
```

Note You can have only maximum seven voting members including arbiter which can be defined as MEMBER*n* and ARBITER*n* and all other member must be defined as NON-VOTING-MEMBER*n*.

Session Cache Scaling

The session cache can be scaled by adding an additional sessionmgr VM (additional session replica-set). You must create separate administration database and the hostname and port should be defined in Policy Builder (cluster) as defined in the following sections:

- Service Restart, on page 109
- Create Session Shards, on page 109

Service Restart

After mongo configuration is done successfully (The `build_set.sh` script gives the status of the mongo configuration after the configuration has been finished) from Cluster Manager, run `/var/qps/bin/control/restartall.sh` script.

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

After we modify `mongoconfig.cfg` file, we can run the `synconfig.sh` script to rebuild `etc.tar.gz` image and trigger each VM to pull and extract it.

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

Create Session Shards

Procedure

Step 1 From pcrfclient01 or pcrfclient02 VM, execute the following command:

```
session_cache_ops.sh --add-shard
```

The following screen prompts are displayed:

```
Session Sharding
------------------------------------------------------
Select type of session shard Default [ ]
Hot Standby [ ]
Sessionmgr pairs :
Session shards per pair :
```

Step 2 Select either **Default** or **Hot Standby** by placing the cursor in the appropriate field and pressing **y**.

Step 3 In Sessionmgr pairs, enter the name of the sessionmgr VM pairs separated by a colon (:) with port number.

Example: `sessionmgr01:sessionmgr02:27717`

If sharding is needed for multiple sessionmgr VMs, enter the sessionmgr VM name with port separated by a colon (:), with each pair separated by a colon (:).

Example: `sessionmgr01:sessionmgr02:27717,sessionmgr03:sessionmgr04:27717`

Step 4 In Session shards per pair, enter the number of shards be added.

Example: `Session shards per pair: 4`

Step 5 Login to ADMIN DB primary mongo sessionmgr VM using port number 27721 and execute the following commands to verify the shards:

```
# mongo sessionmgr01:27721
set05:PRIMARY> use sharding
switched to db sharding
set05:PRIMARY> db.shards.find()
```

Example:

```
# mongo sessionmgr01:27721
MongoDB shell version: 2.6.3
connecting to: sessionmgr01:27721/test
set05:PRIMARY> use sharding
switched to db sharding
set05:PRIMARY> db.shards.find()
{ "_id" : 1, "seed_1" : "sessionmgr01", "seed_2" : "sessionmgr02", "port" : 27717, "db" :
"session_cache", "online" : true, "count" : NumberLong(0), "lockTime" :
ISODate("2015-12-16T09:35:15.348Z"), "isLocked" : false, "lockedBy" : null }
{ "_id" : 2, "seed_1" : "sessionmgr01", "seed_2" : "sessionmgr02", "port" : 27717, "db" :
"session_cache_2", "online" : true, "count" : NumberLong(0), "backup_db" : false, "lockTime" :
ISODate("2015-12-16T09:35:06.457Z"), "isLocked" : false, "lockedBy" : null }
{ "_id" : 3, "seed_1" : "sessionmgr01", "seed_2" : "sessionmgr02", "port" : 27717, "db" :
"session_cache_3", "online" : true, "count" : NumberLong(0), "backup_db" : false, "lockTime" :
ISODate("2015-12-16T09:34:51.457Z"), "isLocked" : false, "lockedBy" : null }
{ "_id" : 4, "seed_1" : "sessionmgr01", "seed_2" : "sessionmgr02", "port" : 27717, "db" :
"session_cache_4", "online" : true, "count" : NumberLong(0), "backup_db" : false, "lockTime" :
ISODate("2015-12-16T09:35:21.457Z"), "isLocked" : false, "lockedBy" : null }
set05:PRIMARY>
```

Verify CPS Sanity

From Cluster Manager, run `/var/qps/bin/diag/diagnostics.sh` script.

Note Currently, running `diagnostics.sh --ha_proxy` with qns-admin or qns-su user is not supported, It's only supported with root user.

Validate VM Deployment

Virtual Interface Validation

To verify that the lbvip01 and lbvip02 are successfully configured in lb01 and lb02, perform the following steps:

Procedure

Step 1 SSH to lb01. The default credentials are qns/cisco123.

Step 2 Check whether the virtual interface of the Policy Director (LB) is UP. Use `ifconfig` command to show the virtual interfaces are UP. If extra diameter interface were configured, verify the corresponding VIPs are up for the diameter interfaces.

Basic Networking

From Cluster Manager, verify that you are able to ping all the hosts in the `/etc/hosts` file.

Diagnostics and Status Check

The following commands can be used to verify whether the installation was successful or not:

- `diagnostics.sh`
- `about.sh`
- `list_installed_features.sh`
- `statusall.sh`

Note For more information on other CPS administrative commands, refer to *CPS Operations Guide*.

diagnostics.sh

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

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

This command will prompt for reboot choice. Please select **Y** for the same and proceed.

Syntax

```
/var/qps/bin/diag/diagnostics.sh -h
Usage: /var/qps/bin/diag/diagnostics.sh [options]
This script runs checks (i.e. diagnostics) against the various access, monitoring, and
configuration points of a running CPS system.
In HA/GR environments, the script always does a ping check for all VMs prior to any other
checks and adds any that fail the ping test to the IGNORED_HOSTS variable. This helps reduce
 the possibility for script function errors.
NOTE: See /var/qps/bin/diag/diagnostics.ini to disable certain checks for the HA/GR env
persistently. The use of a flag will override the diagnostics.ini value.
Examples:
    /var/qps/bin/diag/diagnostics.sh -q
    /var/qps/bin/diag/diagnostics.sh --basic_ports --clock_skew -v
--ignored_hosts='portal01,portal02'

Options:
    --basic_ports : Run basic port checks

        For HA/GR: 80, 11211, 7070, 8080, 8081, 8090, 8182, 9091, 9092, and Mongo DB ports
 based on /etc/broadhop/mongoConfig.cfg
    --clock_skew : Check clock skew between lb01 and all vms (Multi-Node Environment only)
    --diskspace : Check diskspace
    --get_replica_status : Get the status of the replica-sets present in environment.
(Multi-Node Environment only)
    --get_shard_health : Get the status of the sharded database information present in
environment. (Multi-Node Environment only)
    --get_sharded_replica_status : Get the status of the shards present in environment.
(Multi-Node Environment only)
   --ha_proxy : Connect to HAProxy to check operation and performance statistics, and ports
 (Multi-Node Environment only)
        http://lbvip01:5540/haproxy?stats
        http://lbvip01:5540//haproxy-diam?stats
    --help -h : Help - displays this help

    --ignored_hosts : Ignore the comma separated list of hosts. For example
--ignored_hosts='portal01,portal02'
        Default is 'portal01,portal02,portallb01,portallb02' (Multi-Node Environment only)
    --ping_check : Check ping status for all VM
    --qns_diagnostics : Retrieve diagnostics from CPS java processes
    --qns_login : Check qns user passwordless login
    --quiet -q : Quiet output - display only failed diagnostics
    --radius : Run radius specific checks
    --redis  : Run redis specific checks
   --svn : Check svn sync status between pcrfclient01 & pcrfclient02 (Multi-Node Environment
 only)
    --tacacs : Check Tacacs server reachability
    --swapspace : Check swap space
   --verbose -v : Verbose output - display *all* diagnostics (by default, some are grouped
 for readability)
    --virtual_ips : Ensure Virtual IP Addresses are operational (Multi-Node Environment
only)
    --vm_allocation : Ensure VM Memory and CPUs have been allocated according to
recommendations
```

Executable on VMs

- Cluster Manager and OAM (PCRFCLIENT) nodes

Example

```
[root@pcrfclient01 ~]# diagnostics.sh
QNS Diagnostics
```

```
Checking basic ports (80, 7070, 27017, 27717-27720, 27749, 8080, 9091)...[PASS]
Checking qns passwordless logins on all boxes...[PASS]
Validating hostnames...[PASS]
Checking disk space for all VMs...[PASS]
Checking swap space for all VMs...[PASS]
Checking for clock skew...[PASS]
Retrieving QNS diagnostics from qns01:9045...[PASS]
Retrieving QNS diagnostics from qns02:9045...[PASS]
Checking HAProxy status...[PASS]
Checking VM CPU and memory allocation for all VMs...[PASS]
Checking Virtual IPs are up...[PASS]
[root@pcrfclient01 ~]#
```

about.sh

This command displays:

- Core version
- Patch installed
- ISO version
- Feature version
- URLs to the various interfaces
- APIs for the deployment

This command can be executed from Cluster Manager or OAM (PCRFCLIENT).

Syntax

```
/var/qps/bin/diag/about.sh [-h]
```

Executable on VMs

- Cluster Manager
- OAM (PCRFCLIENT)

list_installed_features.sh

This command displays the features and versions of the features that are installed on each VM in the environment.

Syntax

```
/var/qps/bin/diag/list_installed_features.sh
```

Executable on VMs

- All

statusall.sh

This command displays whether the monit service and CPS services are stopped or running on all VMs. This script can be executed from Cluster Manager or OAM (PCRFCLIENT).

Syntax

```
/var/qps/bin/control/statusall.sh
```

Executable on VMs

- Cluster Manager
- pcrfclient01/02



Note Refer to *CPS Operations Guide* for more details about the output of this command.

Web Application Validation

To verify that the CPS web interfaces are running navigate to the following URLs where <*lbvip01*> is the virtual IP address you defined for the lb01 VM.



Note Run the `about.sh` command from the Cluster Manager to display the actual addresses as configured in your deployment.

- **Policy Builder:** https://<*lbvip01*>:7443/pb

  Default credentials: qns-svn/cisco123

- **Control Center:** https://<*lbvip01*>:443

  Default credentials: qns/cisco123

- **Grafana:** https://<*lbvip01*>:9443/grafana

  Default credentials: —



Note You must create at least one Grafana user to access the web interface. Refer to the *Prometheus and Grafana* chapter of the *CPS Operations Guide* for steps to configure User Authentication for Grafana.

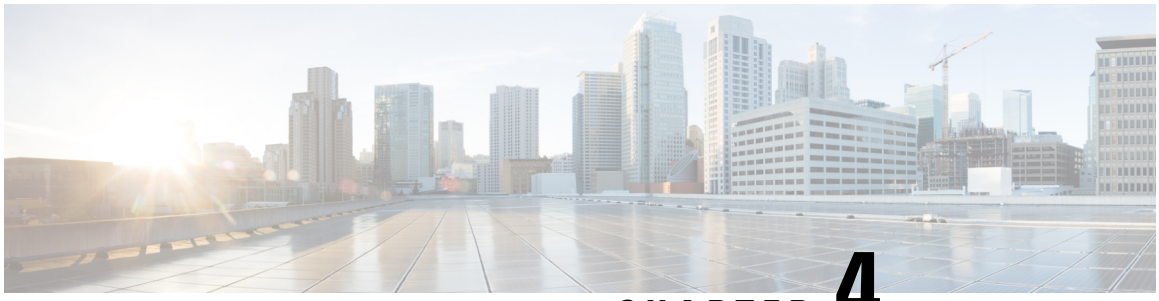
- **Unified API:** http://<*lbvip01*>:8443/ua/soap
- **CRD REST API:** http://<*lbvip01*>:8443/custrefdata

For more information related to CPS interfaces, refer to *CPS Operations Guide*.

Supported Browsers

CPS supports the most recent versions of the following browsers:

- Firefox
- Chrome
- Safari
- Microsoft IE version 9 and above



CHAPTER 4

Post Installation Processes

- Post Installation Configurations, on page 117
- Modify Configuration Files, on page 127
- Scaling Existing Installation, on page 127
- Configure Balance Shards, on page 129
- Secondary Key Ring Configuration, on page 131
- Configuring SK DB, on page 134

Post Installation Configurations

Configure Control Center Access

After the installation is complete you need to configure the Control Center access. This is designed to give the customer a customized Control Center username. For more information on Control Center Access, refer to *CPS Operations Guide*.

Configure NTP on Cluster Manager



Note By default, NTP is configured on Installer/Cluster Manager VM. No additional configuration steps are required.

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
```

IPv6 Support - VMware

Enable IPv6 Support

For VMware hypervisor, IPv6 needs to be enabled first.

Procedure

Step 1 Select the blade from the left panel where you want to enable IPv6 support.

Step 2 Click **Configure** tab from the top menu from the right panel.

Step 3 Under **Networking**, click **Advanced** from the options available.

Step 4 Click **Edit...** in the upper right corner of the **Advanced** panel.

The **Edit Advanced Network Settings** window opens.

Step 5 From **IPv6 support** drop-down list, select **`Enabled`** to enable IPv6 support.

By performing above steps, IPv6 will be enabled on the blade. Rebooting the blade is required for this setting to take effect.

Note
All CPS nodes support IPv4 and IPv6 addresses.

Set Up IPv4 and IPv6 Addresses for VMs

Any hosts in the CPS cluster can be configured to have IPv4 or IPv6 addresses. Currently, IPv6 is supported only for policy director (lb) external interfaces.

For more information on how to configure IPv6 addresses for VMs, refer to the section Hosts Configuration, on page 35.

Converting IPv4 to IPv6 on Policy Director External Interfaces

To convert an existing CPS deployment from IPv4 to IPv6 (external IP addresses on lb* VM), perform the following steps:

Procedure

Step 1 Log in to Cluster Manager.

Step 2 Backup the relevant files using the following commands:

```
mkdir /var/backup_ipv4
cp -rf /var/qps/config/deploy/csv /var/backup_ipv4
cp -rf /etc/puppet/modules/qps/templates/var/broadhop/init_pacemaker_res.sh /var/backup_ipv4
```

Step 3 Update the CSV files as per your IPv6 requirement.

The following sample configuration files for Hosts.csv, AdditionalHosts.csv, and Vlan.csv that use IPv6 address are shown:

- Hosts.csv:

```
cat /var/qps/config/deploy/csv/Hosts.csv
Hypervisor Name,Guest Name,Role,Alias,Datastore,Networks -->,Internal,Management
10.10.10.1,lb01,lb01,lb01,datastore8,,192.1.168.10,2003:3041:0000:0000:0000:0000:0022:0020
10.10.10.2,lb02,lb02,lb02,datastore9,,192.1.168.11,2003:3041:0000:0000:0000:0000:0022:0021
10.10.10.1,pcrfclient01,pcrfclient01,pcrfclient01,datastore8,,192.1.168.12,
10.10.10.2,pcrfclient02,pcrfclient02,pcrfclient02,datastore9,,192.1.168.13,
10.10.10.1,qns01,qps,qns01,datastore8,,192.1.168.14,
10.10.10.2,qns02,qps,qns02,datastore9,,192.1.168.15,
10.10.10.1,qns03,qps,qns03,datastore8,,192.1.168.16,
10.10.10.2,qns04,qps,qns04,datastore9,,192.1.168.17,
10.10.10.1,sessionmgr01,sm,sessionmgr01,datastore8,,192.1.168.18,
10.10.10.2,sessionmgr02,sm,sessionmgr02,datastore9,,192.1.168.19,
```

- AdditionalHosts.csv:

```
cat /var/qps/config/deploy/csv/AdditionalHosts.csv
Host,Alias,IP Address
ntp-primary,ntp,10.14.58.1
ntp-secondary,btp,10.14.58.2
lbvip01,lbvip01,2003:3041::22:22
lbvip02,lbvip02,192.1.168.20
arbitervip,arbitervip,192.1.168.250
sslvip01,sslvip01,10.12.12.18
qns-site-server-2,pcrf,10.12.12.24
snmp-trapdest,nms-destination,10.12.12.5
10.10.207.8,,10.10.207.8
10.10.207.9,,10.10.207.9
```

- Vlans.csv:

```
cat /var/qps/config/deploy/csv/VLANs.csv
VLAN Name,Network Target Name,Netmask,Gateway,VIP Alias
Internal,vlan467,255.255.255.0,NA,lbvip02
Management,vlan467,64,2003:3041::22:1,lbvip01
External,qps-vlan,255.255.255.0,NA,rtp-swag-vm204
```

Step 4 Execute the following commands to update the changes through puppet and redeploy the Policy Director (lb) VMs:

```
/var/qps/install/current/scripts/import/import_deploy.sh
cd /var/qps/install/current/scripts/deployer
deploy.sh lb01
deploy.sh lb02
```

Note
Configure the appropriate firewall rules required for IPv6 or disable the same.

Step 5 After modifying the files, execute the following commands:

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

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

Synchronize Time Between Nodes

To synchronize time between VM nodes, perform the following steps:

Procedure

Step 1 Login to Cluster Manager VM.

Step 2 Execute the following command to synchronize the time between nodes:

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

Note
If this is a Geographic Redundancy (GR) installation with multiple sites, refer to *CPS Geographic Redundancy Guide*.

To check the current clock skew of the system, execute the following command:

```
diagnostics.sh --clock_skew -v
```

The output numbers are in seconds. Refer to the following sample output:

```
CPS Diagnostics Multi-Node Environment
--------------------------
Checking for clock skew...
Clock skew not detected between qns01 and lb01. Skew: 1...[PASS]
Clock skew not detected between qns02 and lb01. Skew: 0...[PASS]
Clock skew not detected between lb01 and lb01. Skew: 0...[PASS]
Clock skew not detected between lb02 and lb01. Skew: 0...[PASS]
Clock skew not detected between sessionmgr01 and lb01. Skew: 0...[PASS]
Clock skew not detected between sessionmgr02 and lb01. Skew: 0...[PASS]
Clock skew not detected between pcrfclient01 and lb01. Skew: 0...[PASS]
Clock skew not detected between pcrfclient02 and lb01. Skew: 0...[PASS]
```

Update the VM Configuration without Re-deploying VMs

Sometimes, certain configurations in the excel sheet need to be modified and updated to the deployed VMs. To update the configurations in the excel sheet, perform the following steps:

Procedure

Step 1 Make the changes to the excel.

Step 2 Save them as CSV files.

Step 3 Upload the csv files to the Cluster Manager VM in `/var/qps/config/deploy/csv/`.

Step 4 Execute the following commands after uploading the csv files to Cluster Manager VM:

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

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

Reserving Memory on the Virtual Machines (VMs)

To avoid performance impact, you must reserve all allocated memory to each CPS virtual machine.

It is recommended to allocate 8 GB memory for the Hypervisor. For example, suppose the total memory allocated on a blade/ESXi host is 48 GB then we should only allocate 40 GB to CPS VMs and keep 8 GB for the Hypervisor.

Note This is required only if your ESXi host is added to vCenter. If not then the deployment takes care of the reservation.

Power OFF the virtual machine before configuring the memory settings.

Procedure

Step 1 Log in to your ESXi host with the vSphere Client.

Step 2 In the vSphere Client, right-click a virtual machine from the inventory and select **Edit Settings...**.

Step 3 In the **Virtual Machine Properties** window, select **Resources** tab and select **Memory**.

Step 4 In the **Resource Allocation** pane, set the memory reservation to allocated memory.

Step 5 Click **OK** to commit the changes.

Step 6 Power ON the Virtual Machine.

Configure Custom Route

In lb01 and lb02, if needed, custom route should be configured to route diameter traffic to the PGWs.

Add a file called route-ethxx in the `./etc/sysconfig/network-scripts`.

For example, 172.20.244.5/32 via 172.16.38.18

Destination subnet via GW of the subnet.

TACACS+

TACACS+ Configuration Parameters

Basic instructions for enabling TACACS+ AAA in the system can be found in the section Configure System Parameters for Deployment, on page 30. There are a number of advanced configuration options which allow administrators to tune their deployments for their specific needs. The following table list TACACS+ configuration parameters that can be added in the Configuration sheet:

Table 41: TACACS+ Configuration Parameters

Parameter	Description	Value Range
tacacs_enabled*	A boolean value indicating whether TACACS+ AAA must be enabled or not.	Values: 1, 0, true, false For example: tacacs_enabled,1
tacacs_server*	An ordered comma-separated list of *<ip>[:port]* pairs indicating which servers need to be queried for TACACS+ AAA.	Values: NA For example: tacacs_server“10.0.2.154:49 ,172.18.63.187:49” **Note** If multiple servers are defined, they must be separated by a comma and enclosed in double quotes, as shown in the example above. Port number with the IP address is optional.
tacacs_secret*	The 'secret' key string used for encrypting the TACACS+ protocol communications.	Values: NA For example: tacacs_secret,CPE1704TKS
tacacs_debug	An integer value indicating the debug level to run the software in. Currently, this is effectively boolean.	Value: 0 1 For example: tacacs_debug,1 Default: 0
tacacs_service	A string value indicating which service to be used when authorizing and auditing against the TACACS+ servers.	Value: NA For example: tacacs_servicepcrflinuxlogin Default: pcrflinuxlogin if no value is specified

Parameter	Description	Value Range
tacacs_protocol	A string value indicating which protocol to be used when authorizing and auditing against the TACACS+ servers.	Value: NA For example: tacacs_protocol,ssh Default: ssh
tacacs_timeout	An integer that represents how long the software needs to wait, in seconds, for the TACACS+ server to respond to the queries.	Value: in seconds For example: tacacs_timeout,2 Default: 5 seconds

The * mark indicates that the parameter is mandatory. * mark is not a part of the parameter.

Arbiter Configuration for TACACS+

Procedure

Step 1 Create the following `yaml` file on Cluster Manager: `/etc/facter/facts.d/tacacs.yaml`.

tacacs_enabled: true

tacacs_server: ip address

tacacs_secret: password

Step 2 Execute `puppet apply` command to apply the appropriate configurations changes to the system:

```
/usr/bin/puppet apply --modulepath "/etc/puppet/modules:/etc/puppet/env_config/modules" --pluginsync
/etc/puppet/manifests/init.pp --config /etc/puppet/puppet.conf --logdest /var/log/puppet.log
```

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

TACACS+ Enabler

The `enable_tacacs+` utility can be used to configure the Cluster Manager VM for TACACS+-based authentication. The utility achieves this by first validating if TACACS+ has been configured properly using the Configuration sheet of CPS Deployment Template (Excel spreadsheet). Assuming the required values are provided, the utility then selectively applies several Puppet manifests to enable TACACS+ authentication on the target VM.

To use the utility:

Procedure

Step 1 (Optional) Copy the utility to the `/var/qps/bin/support` directory.

```
cp tacacs_enabler/enable_tacacs+ /var/qps/bin/support/
```

Note
This step places the utility into a directory which should be in the PATH on the target VM. While not required, this simplifies execution of the script for the later steps.

Step 2 (Optional) Execute the `enable_tacacs+ clustermgr --check` script in 'check' mode to validate the configuration values:

```
Detected VM node type: clustermgr
        Generating facts based on current deployment configuration

        Validating TACACS+ configuration settings:
        * Found required setting for 'tacacs_secret'
        * Found optional setting for 'tacacs_debug'
        * Found optional setting for 'tacacs_timeout'
        * Found required setting for 'tacacs_server'
        * Found required setting for 'tacacs_enabled'

        Configuration appears to be complete. You should be able to enable TACACS+
        on this 'clustermgr' node by executing this command without the '--check'
        command-line option.
```

Step 3 Execute the script without the '--check' command-line option to apply the configuration:

```
enable_tacacs+ clustermgr
```

```
Detected VM node type: clustermgr
        Generating facts based on current deployment configuration

        Validating TACACS+ configuration settings:
        * Found required setting for 'tacacs_secret'
        * Found optional setting for 'tacacs_debug'
        * Found optional setting for 'tacacs_timeout'
        * Found required setting for 'tacacs_server'
        * Found required setting for 'tacacs_enabled'

        Executing Puppet to apply configuration:
        ... Puppet output ...
        Notice: Finished catalog run in 34.57 seconds
```

Step 4 Validate that TACACS+ authenticated users are available on the target VM:

`id -a <`*TACACS+ user*`>`

Configure Multiple Redis Instances



Note All the commands mentioned in the following section should be executed on Cluster Manager.

Before you begin

Redis instance must be enabled and running.

Procedure

Step 1 To configure multiple redis instances, update *redis_server_count* parameter in `Configuration.csv` spreadsheet in `QPS_deployment_config_template.xlsm` deployment template file.

Step 2 After updating the `Configuration.csv`, execute the following command to import the new configuration file into Cluster Manager VM.

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

Step 3 Edit `redisTopology.ini` file in `/etc/broadhop/` directory and add all redis endpoints:

Note
By default, three redis instances are enabled.

For example, for three redis instances, the `redisTopology.ini` file will look like:

```
policy.redis.qserver.1=lb01:6379
policy.redis.qserver.2=lb02:6379
policy.redis.qserver.3=lb01:6380
policy.redis.qserver.4=lb02:6380
policy.redis.qserver.5=lb01:6381
policy.redis.qserver.6=lb02:6381
```

Note

- For every added redis instance, you need to add two lines in `redisTopology.ini` file.
- Redis instances are monitored by monit on lb VMs.
- Guidelines on number of redis instances running on each policy director (lb):

  If Diameter TPS < 15K, then the default number of redis instances (3) will be running on each policy director (lb).

  If Diameter TPS < 28K, then the number of redis instances running on each policy director (lb) should be 4.

  If Diameter TPS > 28K, then the number of redis instances running on each policy director (lb) should be 5.

Step 4 After modifying the configuration file, to make the changes permanent, user needs to rebuild `etc.tar.gz`.

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

Step 5 Reinitialize the environment:

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

Step 6 Restart the qns service.

```
/var/qps/bin/control/restartall.sh
```

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

Configure Redis Instances for Keystore

Currently, keystore is being used internally for features such as RAN NAS Retry, Holding Rx STR, and so on.

- Keystore is a temporary cache used by application to store temporary information. It stores information in the form of key-value pair.
- Keystore internally uses redis cache for storing the key-value pair.



Note By default, keystore uses redis running on lb01:6379 and lb02:6379 if redis instances is not configured for keystore.

```
-Dredis.keystore.connection.string=lb01:lb02:6379:6379
```

Before you begin

Redis instance must be installed and running on VMs.

Procedure

If you want to add more redis instances for keystore, execute the following OSGi command:

```
telnet qns01 9091
```

Note
qns01 must be up and running.

```
setKeystoreConnectionString <start lb>:<end lb>:<start port>:<end port>
```

Range of lbs can be defined using <*start lb*>:<*end lb*>.

Range of redis ports can be defined using <*start port*>:<*end port*>.

For example, to use redis instance running on 6379, 6380 on lb01 to lb04, configure the parameter as follows:

```
telnet qns01 9091
osgi> setKeystoreConnectionString lb01:lb04:6379:6380
Keystore string updated successfully
```

Note
The parameter `-Dredis.keystore.connection.string` has been deprecated from CPS 12.0.0 release and is only used to maintain backward compatibility.

The current keystore instances which are being used in application can be checked by the running the following command:

```
telnet qns01 9091
```

Note
qns01 must be up and running.

```
listKeystoreShards
```

For example:

```
telnet qns01 9091
osgi> listKeystoreShards

Shard Id    Keystore Instances

1           lb01:6379
2           lb01:6380
3           lb02:6379
4           lb02:6380

Keystore Shards Status: BALANCED
```

Modify Configuration Files

Customers might need to change configurations in the `/etc/broadhop` directory on VMs. It is recommended not to change the configurations in the VMs. All changes must be done in the Cluster Manager VM. To modify configuration files, perform the following steps:

Procedure

Step 1 In Cluster Manager, modify the files in `/etc/broadhop/`.

Step 2 (Optional) For HA system, we need to configure TCP no delay by modifying the set *Denable_tcp_nodelay* flag in `/etc/broadhop/qns.conf` file.

-Denable_tcp_nodelay=true

Step 3 After modifying the configuration file, to make the changes permanent for future use (when any VM is redeployed or restarted... etc.), user needs to rebuild `etc.tar.gz`.

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

Step 4 In Cluster Manager, execute the following command to synchronize the changes to the VM nodes.

```
SSHUSER_PREFERROOT=true copytoall.sh /etc/broadhop/qns.conf /etc/broadhop/qns.conf
```

Step 5 Restart the CPS service if necessary on cluster manager.

```
/var/qps/bin/control/restartall.sh
```

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

Scaling Existing Installation

There might be situations when customer would want to expand existing installation e.g. add more Policy Server (QNS) or session manager virtual machines.

To add more VMs, perform the following steps:

Procedure

Step 1 Refer to the template file that were used to create earlier installation. These files are present under `/var/qps/config/deploy/csv`. For more information, refer to Import the csv Files into the Cluster Manager VM, on page 80.

Step 2 Assuming that we are not adding any new VLANs in the scaling up of setup, modify the csv files in the following order to include additional changes:

- Update the Hosts.csv to include new hosts on appropriate ESXi hosts, with corresponding guest name, role, alias etc. For more information, refer to Hosts Configuration, on page 35.
- For scaling up, if we want to add more Policy Server (QNS), say qns03 and qns04, those entries should get reflected appropriately in above Hosts file.

Note
No changes are needed in rest of the template files.

Step 3 Validate the configurations using `jvalidate.py` script. For more information, refer to Validate Imported Data, on page 81.

```
cd /var/qps/install/current/scripts/deployer/support/

python jvalidate.py
```

Step 4 Once Steps 2 and 3 are completed, import the modified csv file by executing the following command:

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

This would convert updated csv into the JSON file, and also create new `/etc/hosts` file on the Cluster Manager with entries for new virtual machines.

Step 5 For each new hosts (VM) that is defined in the Hosts sheet, we need to run `deploy.sh` script. For more information, refer to Manual Deployment .

Note
Make sure that we do not execute `deploy_all.sh` script as it would wipe out the existing deployment and recreate new VMs.

Step 6 Manually copy the new `/etc/hosts` file from cluster manager to all (new and existing) virtual machines.

Note
Currently, there is no procedure to synchronize `/etc/hosts` to all the hosts.

What to do next

Important If existing four qns are not able to handle CC (Control Center) and API traffic, we need to make changes in `/etc/haproxy/haproxy.conf` file for additional Policy Server (QNS). If we do not add entries for additional Policy Server (QNS), (e.g. qns05 and qns06), then the CC and API traffic would be handled by existing four Policy Server (QNS) VMs i.e., qns01 to qns04. Also, no changes are required to be done in `/etc/broadhop/servers` for new VMs.

For Gx, no entries are required in `haproxy.conf` file.

Adding Member to Existing Replica Set

During above process you can add new session managers for databases and expand the existing replica-set. The procedure for the same is covered in Defining a Replica-set, on page 105.

Configure Balance Shards

Balance database can be sharded logically to improve the performance of balance database. Internally it will create multiple balance dbs and distribute the data among each shards.

Prerequisites

This feature is available in 7.5.0 and higher releases. By default, there is one shard that gets created for balance.

- Adding or removing shards to the Balance database must be performed during a maintenance window.
- Back up the Balance database before adding or removing shards. Refer to the *CPS Backup and Restore Guide* for instructions.

Shard Configuration

Shard collection can be increased/decreased based on performance needs.

Add Shards to Balance Database

The following example increases the number of shards from 1 to 6.

Procedure

Step 1 Log into the Control Center and note down the balance information for a few subscribers. This information is used to confirm the balance of these users after the shards have been added.

Step 2 Log into the Cluster Manger VM.

Step 3 Edit `/etc/broadhop/pcrf/qns.conf` to add the following parameter:

```
-Dcom.cisco.balance.dbs=6
```

Step 4 Run `copytoall.sh` to synchronize the configuration changes to all VMs in the CPS cluster.

Step 5 Run `restartall.sh` to restart all Policy Server (QNS) processes.

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

Step 6 After restart, connect to qns01 OSGi console and execute `rebalanceBalanceShard <newShardCount>` command.

where, <*newShardCount*> is the new shard count equal to the value configured for `com.cisco.balance.dbs` in `/etc/broadhop/pcrf/qns.conf` file.

Example: `rebalanceBalanceShard 6`

Caution
This command may take time to complete. Monitor the rebalance shard and wait until the command finishes. Do not restart Policy Server (QNS) while rebalance is in progress.

This creates six logical shards.

Caution
To terminate the OSGi session, use the disconnect command. Do not use the exit command, as this command restarts the process.

Step 7 Verify by connecting to the Balance database. A total of six entries for balance_mgmt must be listed, (balance_mgmt – balance_mgmt_5).

Step 8 Log into Control Center again and verify that the subscribers from 1 have the same balance.

Remove Shards from Balance Database

The following example decreases the number of shards from 6 to 1.

Procedure

Step 1 Log into Control Center and note down the balance information for a few subscribers.

Step 2 Go to Policy Server (QNS) OSGi console and run `rebalanceBalanceShard <newShardCount>` command.

where, <*newShardCount*> is the new shard count equal to the value configured for `com.cisco.balance.dbs` in `/etc/broadhop/pcrf/qns.conf` file.

Caution
This command may take time to complete. Monitor the rebalance shard and wait until the command finishes. Do not restart Policy Server (QNS) while rebalance is in progress.

This reduces the shards from six to one.

Caution
To terminate the OSGi session, use the `disconnect` command. Do not use the `exit` command, as this command restarts the process.

Step 3 Log into the Cluster Manager VM.

Step 4 Edit the `/etc/broadhop/pcrf/qns.conf` file and add or modify the following parameter:

```
-Dcom.cisco.balance.dbs=1
```

Step 5 Run `copytoall.sh` to synchronize the configuration changes to all VMs in the CPS cluster.

Step 6 Run `restartall.sh` to restart all Policy Server (QNS) processes.

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

Step 7 Verify by connecting to the Balance database and see the count. Only one entry for balance_mgmt should now be listed.

Step 8 Log into Control Center again and verify that the subscribers from 1 have the same balance.

Secondary Key Ring Configuration

CPS provides a high availability solution for secondary key to primary key mappings. Rings are group of memcached servers processes running on different sessionmgr VMs (session cache) which stores the mapping of primary and secondary keys. This is used for secondary key lookup to optimize performance for Rx calls. Examples of secondary key lookups include framed IP Rx session ID IMSI MSISDN.

Architecturally the solution is divided into the following components

- Secondary Key Ring — A secondary key ring is a set of nodes that have a complete set of secondary key to primary key mappings. The secondary keys are partitioned using consistent hashing across the nodes within the ring to ensure an even distribution of the keys.
- Ring Set — Each node on a secondary key ring is called a ring set. A ring set can have 1 to many physical servers. Each server has an exact copy of the data stored for that node. Each additional server within a ring set increases the high availability capability of the system.

Using these component pieces the system supports parallel searching of key mappings across the physical servers to find a specific entry. If a physical server is shutdown or becomes unavailable the system automatically rebuilds the rings and remap the secondary keys to the primary keys when the server comes back online.

The system does not support the following scenario:

- Detecting if a ring is need of a rebuild due to issuing a `flush_all` command.

Why it is Required

- Secondary key (Rx) to primary key (Gx) lookups are cached into a set of *n* servers and failure of a server results in a loss of 1/*n* of the primary to secondary key mappings. Because of this failure number of additional queries continues to increase also the keys are not removed on session removal which ages out.
- Rings are used to handle this situation which allows the server endpoints to grow or shrink. Each key is written to multiple memcached servers within a ring.
- Keys are removed on session removal to keep the cache keys from expiring.
- Queries are parallely executed when search is done against multiple rings to allow for a ring and multiple servers within a ring.

Key Ring Commands

The following commands are provided to support this new functionality.



Note Before implementing any of these commands contact the Cisco AS team to discuss the optimal architecture for your CPS deployment.

All commands must be issued from Policy Server (QNS).

Telnet to any Policy Server (QNS) machine on port 9091 to enter the OSGi console.

Creating a New Ring

To create a new secondary key (sk) ring with a given id:

`createSkRing` *ringid*



Note The *ringid* must be numeric and the ring must initially be empty with no ring sets defined.

Example:

```
createSkRing 2
```

Adding a New Endpoint

This command assigns a set of servers to act as node on the cache ring. Each server will have an exact copy of the data. If a node exists in the ring with that id then it is replaced and the ring is automatically rebuilt.

`setSkRingSet` *ringid setid cacheserver1port[cacheserver2portcacherverNport]*

Example:

```
setSkRingSet 1 1 sessionmgr01:11211 sessionmgr02:11211
```

Removing an Endpoint

This command removes a ring set from a ring. This triggers an automatic rebuild of the ring.

`removeSkRingSet` *ringid setid*

Example:

```
removeSkRingSet 1 2
```

Removing a Ring

This command removes a ring.



Note You cannot remove the last ring from the system.

`removeSkRing` *ringid*

Example:

```
removeSkRing 2
```

Triggering a Ring Rebuild

To trigger a rebuild of a secondary key ring with a given id:

`rebuildSkRing` *ringid*

where, *ringid* is a numeric value.

Example:

rebuildSkRing 1

To track the progress of a ring rebuild refer to the following statistic:

```
skcache_ring[ring id]_entry_rebalance
```

Single Cluster Configuration

Log into pcrfcleint01 or 02 to create/update rings from Policy Server (QNS) OSGi console. Assuming, there are three session cache replica sets, by default, Ring-1 Set-1 gets configured automatically and remaining rings need to be configured manually.

```
osgi> setSkRingSet 1 2 sessionmgr03:11211,sessionmgr04:11211

Ring updated

osgi> setSkRingSet 1 3 sessionmgr05:11211,sessionmgr06:11211
```

Multi-Cluster Configuration

Log into pcrfcleint01 or 02 to create/update rings from Policy Server (QNS) OSGi console. Assuming there are three session cache replica sets by default Ring-1 Set-1 get configured automatically and remaining rings need to be configured manually.

- Configure cluster-1 (Ring-1):

  ```
  osgi> setSkRingSet 1 2 sessionmgr03:11211,sessionmgr04:11211
  Ring updated
  osgi> setSkRingSet 1 3 sessionmgr05:11211,sessionmgr06:11211
  ```

- Configure cluster-2 (Ring-2):

  ```
  telnet qns01 9091
  osgi> createSkRing 2
  Successfully added ring with ID: 2
  osgi> setSkRingSet 2 1 sessionmgr01:11211,sessionmgr02:11211
  osgi> setSkRingSet 2 2 sessionmgr03:11211,sessionmgr04:11211
  osgi> setSkRingSet 2 3 sessionmgr05:11211,sessionmgr06:11211
  ```

  Log into admin database and verify. You should be able to see such entries in `cache_config` collection.

GR Configuration with Session Replication Across Sites

Login to pcrfclient01/02 to create/update rings from Policy Server (QNS) OSGi console. Assuming there are two session cache replica-sets. By default, Ring-1 Set-1 get configured automatically and remaining rings need to be configured manually.

Configure cluster-1 (Ring-1)

```
osgi> setSkRingSet 1 1 <hostname_primary_site_sessionmgr01>:11211,
<hostname_primary_site_sessionmgr02>:11211
Ring updated
osgi> setSkRingSet 1 2 <hostname_primary_site_sessionmgr03>:11211,
<hostname_primary_site_sessionmgr04>:11211
Ring updated
```

Configure cluster-2 (Ring-2)

```
telnet qns01 9091
osgi> createSkRing 2
Successfully added ring with ID: 2
osgi> setSkRingSet 2 1 <hostname_secondary_site_sessionmgr01>:11211,
<hostname_secondary_site_sessionmgr02>:11211
osgi> setSkRingSet 2 2 <hostname_secondary_site_sessionmgr03>:11211,
<hostname_secondary_site_sessionmgr04>:11211
```

An example configuration is given below:

- Configure cluster-1 (Ring-1):

  ```
  osgi> setSkRingSet 1 1 L2-CA-PRI-sessionmgr01:11211, L2-CA-PRI-sessionmgr02:11211
  Ring updated
  osgi> setSkRingSet 1 2 L2-CA-PRI-sessionmgr03:11211, L2-CA-PRI-sessionmgr04:11211
  Ring updated
  ```

- Configure cluster-2 (Ring-2):

  ```
  telnet qns01 9091
  osgi> createSkRing 2
  Successfully added ring with ID: 2
  osgi> setSkRingSet 2 1 L2-CA-SEC-sessionmgr01:11211, L2-CA-SEC-sessionmgr02:11211
  osgi> setSkRingSet 2 2 L2-CA-SEC-sessionmgr03:11211, L2-CA-SEC-sessionmgr04:11211
  ```

Configuring SK DB

In both upgrade and fresh installations, SK DB is disabled by default. You need to manually configure SK DBs and enable SK DB.

To ease configuration change on upgrades, new qns.conf parameter is added to do auto upgrade which can be enabled to copy the existing session DBs as SK DBs. This parameter is disabled by default and needs to be enabled if you want to use this auto upgrade feature. It is recommended to do a manual upgrade as described in the following section.

Upgrading SK DB Manually

Perform the following steps to manually upgrade SK DB:

Procedure

Step 1 Create new SK DB replica sets

- In `/etc/broadhop/mongoConfig.cfg` file, there is no new option for SK DB. Use the same naming convention as for SESSION.
- There is no new option available for SK DB in `build_set.sh`. Use the `--session` option for SK DB.

Step 2 Add all new SK DB shards using OSGi command as follows:

```
addskshard sessionmgr01, sessionmgr02 27717 2
```

Note
Similar to session db, add the default SK db shards.

Step 3 Execute rebalance.

```
rebalancesk
```

or

```
rebalanceskbg
```

Step 4 Execute the following command to check the rebalance status of SK DB shards.

```
rebalanceskstatus
```

Step 5 Execute migrate.

```
migratesk
```

```
migrateskbg
```

Step 6 Execute the following command to check the rebalance status of SK DB shards.

```
rebalanceskstatus
```

Step 7 Configure “skAsync” threads equals to “rules” thread under threading configuration in Policy Builder and publish.

The SK DB insert or delete is done in parallel and is handled by separate thread pool executor. This requires adding the new threading configuration under Threading Configuration. Set Thread pool name as “skAsync” and set Threads for “skAsync” as same as threads for “rules”.

Step 8 After rebalance, set SK DB as FALLBACK to ensure new keys get stored both in SK DB and memcache.

```
setskorder MEMCACHE SK_DB
```

Step 9 Rebuild SK DBs to populate all secondary keys from session db in background.

```
rebuildskdb
```

Step 10 Check rebuild status

```
rebuildskdbstatus
```

Step 11 After rebuild, set SK DB as PRIMARY and DISABLE memcache.

```
setskorder SK_DB
```

Step 12 Disable full scan from Policy Builder and publish.

For more information, see *CPS Mobile Configuration Guide*.

Step 13 Verify SK Evaluation Time In Minutes is set to 60 under Cluster configuration in Policy Builder. If it is not set to 60, manually set to 60 and publish again.

Step 14 Add the following two parameters in `/etc/broadhop/qns.conf` file and restart the system:

```
-Dmongo.connections.per.host.secondary_key=12
-Dmongo.connections.per.host.session=12
```

Important
Recommended Values:

- HA and GR (with SKDB)

```
-Dmongo.connections.per.host.secondary_key=25

-Dmongo.connections.per.host.session=25

-Dsk.db.asyncMaxWaitInMs=200

-Ddb.full.scan.tps.non.diameter=0
```

  - Threading Configuration

    For 1 SKDB shard, 75 threads are recommended

    for 2 SKDB shards, 120 threads are recommended

- GR specific (with SKDB)

```
-Dsk.db.skipRemotePrimary=TRUE
-Dsk.db.skipRemote=TRUE
-Dsk.db.skipPrimary=true
```

Upgrading SK DB with Auto – Upgrade

Perform the following steps to upgrade SK DB using auto-upgrade:

Procedure

Step 1 Add the following parameter in `/etc/broadhop/qns.conf` file, before upgrading SK database.

```
-Dsk.db.replicateSessionSharding=true
```

Step 2 Perform upgrade.

Step 3 Perform the following post upgrade procedures in Policy Builder.

- Disable full scan from PB.
- Configure “skAsync” threads equals to “rules” thread under threading configuration in Policy Builder.

The SK DB insert or delete is done in parallel and is handled by separate thread pool executor. This requires adding the new threading configuration under Threading Configuration. Set Thread pool name as “skAsync” and set Threads for “skAsync” as same as threads for “rules”.

- Verify "SK Evaluation Time In Minutes" is set to 60 under Cluster configuration in Policy Builder. If it is not set to 60, manually set to 60.
- Publish changes.

When the upgrade parameter is present at upgrade time:

- All session shards from existing “shards” is copied into “sk_shards” and rebuild task is triggered in background automatically to populate secondary keys in SK DBs from session dbs. The configuration is updated to set memcached as PRIMARY and SK DB as FALLBACK.
- After rebuild SK DBs is completed, configuration is updated again to set SK DB as PRIMARY and memcached is disabled

Note

- The auto-upgrade is optional and it is recommended to do manual upgrade in maintenance window
- While selecting auto-upgrade option, you need to consider that the SK DB needs extra space on session managers and the existing session db sizing might be impacted. Therefore it is recommended to manually create SK DB shards.

Step 4 Add the following two parameters in `/etc/broadhop/qns.conf` file and restart the system:

```
-Dmongo.connections.per.host.secondary_key=12
-Dmongo.connections.per.host.session=12
```

Important

Recommended Values:

- HA and GR (with SKDB)

```
-Dmongo.connections.per.host.secondary_key=25

-Dmongo.connections.per.host.session=25

-Dsk.db.asyncMaxWaitInMs=200

-Ddb.full.scan.tps.non.diameter=0
```

  - Threading Configuration

    For 1 SKDB shard, 75 threads are recommended

    for 2 SKDB shards, 120 threads are recommended

- GR specific (with SKDB)

```
-Dsk.db.skipRemotePrimary=TRUE
-Dsk.db.skipRemote=TRUE
-Dsk.db.skipPrimary=true
```