CPS Release Change Reference, Release 19.4.0 (2)

First Published: 2019-07-25

Last Modified: 2019-09-09

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: www.cisco.com go trademarks. 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)

© 2019 Cisco Systems, Inc. All rights reserved.

CONTENTS

PREFACE

Preface vii

About This Guide vii

Audience vii

Additional Support vii

Conventions (all documentation) viii

Communications, Services, and Additional Information ix

Important Notes ix

CHAPTER 1

19.4.0 Features and Changes 1

19.4.0 Features and Changes 1

CHAPTER 2

ANDSF 5

ANDSF 5

CHAPTER 3

ATS 7

Parallel Feature File Execution 7

Stargen Master Slave REST APIs 8

Support for StarOS CLI Grammars 8

Updating Jenkins Credentials 9

CHAPTER 4

Geographic Redundancy 11

Geographic Redundancy 11

CHAPTER 5

LWR 13

LWR 13

CHAPTER 6 **Mobile** 15

Alarm Upon Detection of Missing/Corrupt Shards 15

MongoDB Health Monitoring for Write Operations 16

Prevention of Application from Using Missing/Corrupted Shards 19

RCS Feature Supressing Gx 20

Session Manager NTP Synchronization Improvement 24

Support for WPS Service on PCRF 24

CHAPTER 7 **MOG** 27

MOG 27

CHAPTER 8 **Operations** 29

API Additions or Changes 29

MIB Additions or Changes 29

Log Additions or Changes 29

Policy Builder and Control Center Activity Logs 29

SNMP Alarm Additions or Changes 31

Statistics/KPI Additions or Changes 31

CHAPTER 9 **Performance Improvement** 35

Performance Improvement 35

CHAPTER 10 **Platform** 37

Send SNMP Traps for Process Restarts 37

SNMP Alarm Resynchronization 38

SR-IOV Support for CPS 39

CHAPTER 11 **Policy Reporting** 41

Policy Reporting 41

CHAPTER 12 **Product Security** 43

CentOS Security Enhancements/Kernel Upgrade 43

CHAPTER 13 Security Enhancements 47
Security Enhancements 47
PSB Requirements for GUI 47
CHAPTER 14 UDC 49
UDC 49
Support for WPS Service on PCRF 49
CHAPTER 15 UI Enhancements 51
Policy Builder Enhancements 51
CHAPTER 16 vDRA 53
Application based Client Sharding 53
Binding Database Record Capacity Overload Handling 54
DRA Distributor 56
DRA Distributor Usability Improvements 56
In-Service Migration from MongoDB Sharding to Application Sharding 58
Linux Utilities for Troubleshooting 59
Load Balance Diameter Messages across vDRA Relay Connections 60
MongoDB Service without Authentication Detection 61
Redis Security Vulnerabilities on vDRA 62
Routing Method Order Enhancements 63
Support to Create Mongo Shard/Replica-set 64
Support to Create Prometheus Statistics Partition 65
Upgrade MongoDB Version 3.6 66
vDRA Resiliency Improvements and Database Overload Protection 67
Zone-Aware Sharding 67

Preface

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

About This Guide

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.

Audience

This guide is best used by these readers:

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

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

Additional Support

For further documentation and support:

- Contact your Cisco Systems, Inc. technical representative.
- Call the Cisco Systems, Inc. technical support number.
- Write to Cisco Systems, Inc. at support@cisco.com.

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

Conventions (all documentation)

This document uses the following conventions.

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

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

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

Warning IMPORTANT SAFETY INSTRUCTIONS.

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

SAVE THESE INSTRUCTIONS

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

Communications, Services, and Additional Information

- To receive timely, relevant information from Cisco, sign up at Cisco Profile Manager.
- To get the business impact 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.

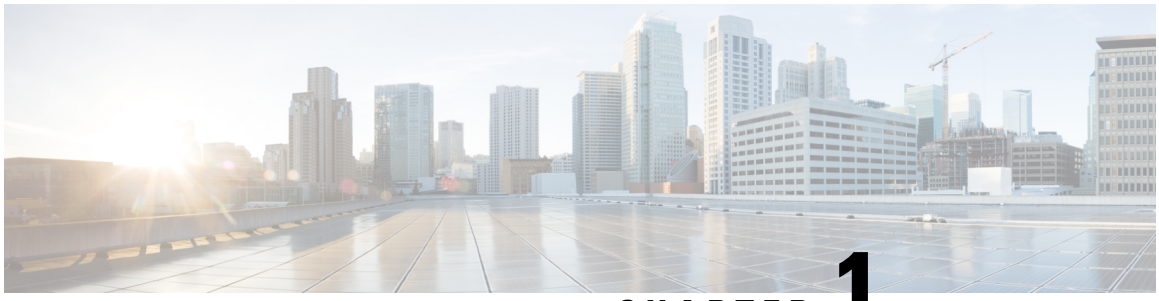

CHAPTER 1

19.4.0 Features and Changes

- 19.4.0 Features and Changes, on page 1

19.4.0 Features and Changes

Features/Behavior Changes	Applicable Product(s)/ Functional Area	Release Introduced/ Modified
Alarm Upon Detection of Missing/Corrupt Shards, on page 15	CPS	19.4.0
Application based Client Sharding, on page 53	vDRA	19.4.0
Binding Database Record Capacity Overload Handling, on page 54	vDRA	19.4.0
CentOS Security Enhancements/Kernel Upgrade, on page 43	CPS	19.4.0
DRA Distributor, on page 56	vDRA	19.4.0
DRA Distributor Usability Improvements, on page 56	vDRA	19.4.0
In-Service Migration from MongoDB Sharding to Application Sharding, on page 58	vDRA	19.4.0
Linux Utilities for Troubleshooting, on page 59	vDRA	19.4.0
Load Balance Diameter Messages across vDRA Relay Connections, on page 60	vDRA	19.4.0
MongoDB Health Monitoring for Write Operations, on page 16	CPS	19.4.0

Features/Behavior Changes	Applicable Product(s)/ Functional Area	Release Introduced/ Modified
MongoDB Service without Authentication Detection, on page 61	vDRA	19.4.0
Parallel Feature File Execution, on page 7	ATS	19.4.0
Policy Builder and Control Center Activity Logs, on page 29	CPS	19.4.0
Policy Builder Enhancements, on page 51	CPS	19.4.0
Prevention of Application from Using Missing/Corrupted Shards, on page 19	CPS	19.4.0
PSB Requirements for GUI, on page 47	CPS	19.4.0
RCS Feature Supressing Gx, on page 20	CPS	19.4.0
Redis Security Vulnerabilities on vDRA, on page 62	vDRA	19.4.0
Routing Method Order Enhancements, on page 63	vDRA	19.4.0
Send SNMP Traps for Process Restarts, on page 37	CPS	19.4.0
Session Manager NTP Synchronization Improvement, on page 24	CPS	19.4.0
SNMP Alarm Additions or Changes, on page 31	CPS	19.4.0
SNMP Alarm Resynchronization, on page 38	CPS	19.4.0
SR-IOV Support for CPS, on page 39	CPS	19.4.0
Stargen Master Slave REST APIs, on page 8	ATS	19.4.0
Statistics/KPI Additions or Changes, on page 31	CPS/vDRA	19.4.0
Support for StarOS CLI Grammars, on page 8	ATS	19.4.0

Features/Behavior Changes	Applicable Product(s)/ Functional Area	Release Introduced/ Modified
Support for WPS Service on PCRF, on page 24	UDC	19.4.0
Support to Create Mongo Shard/Replica-set, on page 64	vDRA	19.4.0
Support to Create Prometheus Statistics Partition, on page 65	vDRA	19.4.0
Updating Jenkins Credentials, on page 9	ATS	19.4.0
Upgrade MongoDB Version 3.6, on page 66	vDRA	19.4.0
vDRA Resiliency Improvements and Database Overload Protection, on page 67	vDRA	19.4.0
Zone-Aware Sharding, on page 67	vDRA	19.4.0

19.4.0 Features and Changes

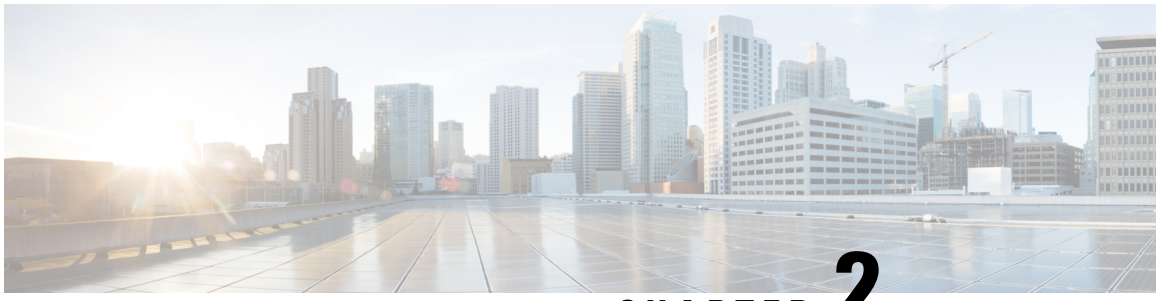

CHAPTER 2

ANDSF

- ANDSF, on page 5

ANDSF

No new features or changes were introduced in this release.

ANDSF

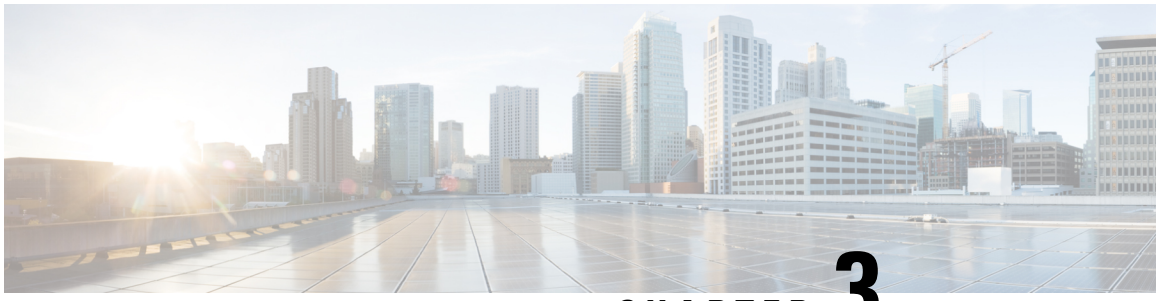

CHAPTER 3

ATS

- Parallel Feature File Execution, on page 7
- Stargen Master Slave REST APIs, on page 8
- Support for StarOS CLI Grammars, on page 8
- Updating Jenkins Credentials, on page 9

Parallel Feature File Execution

Feature Summary and Revision History

Table 1: Summary Data

Applicable Product(s) or Functional Area	CPS
Applicable Platform(s)	Not Applicable
Default Setting	Disabled - Configuration Required
Related Changes in This Release	Not Applicable
Related Documentation	Contact your Cisco Account representative

Table 2: Revision History

Revision Details	Release
First introduced	19.4.0

Feature Description

ATS now provides support for parallel execution of reference feature file from another feature file.

Currently, the parallel execution is supported for child feature file only.

For more information, contact your Cisco Account representative.

Stargen Master Slave REST APIs

Feature Summary and Revision History

Table 3: Summary Data

Applicable Product(s) or Functional Area	CPS
Applicable Platform(s)	Not Applicable
Default Setting	Enabled – Configuration Required
Related Changes in This Release	Not Applicable
Related Documentation	Contact your Cisco Account representative

Table 4: Revision History

Revision Details	Release
First introduced	19.4.0

Feature Description

You can now use new APIs supported by stargen-master tool. The following APIs are supported:

- Tool-specific Starttool API
- Tool-specific Stoptool API
- Tool-specific Installtool API

Note Understanding of testing scenario with stargen master-slave support is required.

For this feature to work, site-toolagent version 19.3.0-RELEASE and onwards must be installed.

For more information, contact your Cisco Account representative.

Support for StarOS CLI Grammars

Feature Summary and Revision History

Table 5: Summary Data

Applicable Product(s) or Functional Area	CPS
Applicable Platform(s)	Not Applicable

Default Setting	Enabled - Always-on
Related Changes in This Release	Not Applicable
Related Documentation	Contact your Cisco Account representative

Table 6: Revision History

Revision Details	Release
First introduced	19.4.0

Feature Description

ATS now provides support for StarOS CLI execution and output parsing.

For more information, contact your Cisco Account representative.

Limitations

- Currently, child elements are parsed based on spaces. There is no auto intelligent to detect specific child when spaces are given incorrectly in the command output.
- Currently, multiColumnStartPattern can be applied only once per command.

Updating Jenkins Credentials

Feature Summary and Revision History

Table 7: Summary Data

Applicable Product(s) or Functional Area	CPS
Applicable Platform(s)	Not Applicable
Default Setting	Enabled - Always-on
Related Changes in This Release	Not Applicable
Related Documentation	Contact your Cisco Account representative

Table 8: Revision History

Revision Details	Release
First introduced	19.4.0

Feature Description

ATS now provides support to update Jenkins details for accessing and triggering jobs in pats-executor pod from PATS Central.

For this feature to work, you should have latest PATS Central distribution.

To access the secured Jenkins URL of the patsexecutor pod, enter the patsexecutor ingress IP address in the browser. The Jenkins login page opens up. Provide the login credential and then click **log in**.

Note Currently, this feature is only available in PATS Central Jenkins of pats-executor pod.

For more information, contact your Cisco Account representative.

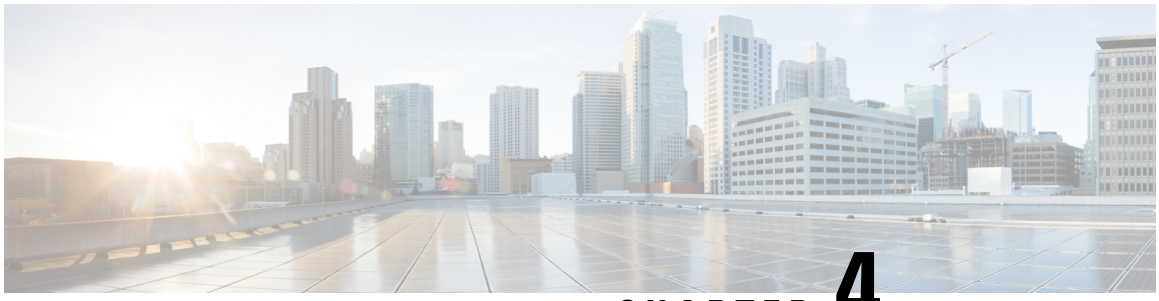

CHAPTER 4

Geographic Redundancy

• Geographic Redundancy, on page 11

Geographic Redundancy

No new features or changes were introduced in this release.

Geographic Redundancy

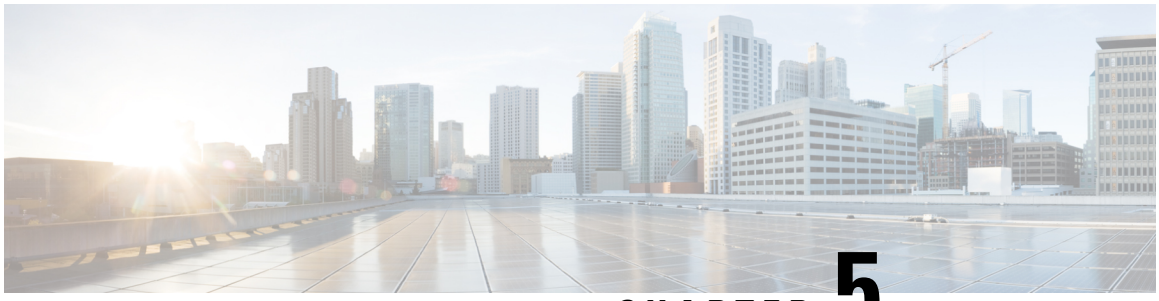

CHAPTER 5

LWR

• LWR, on page 13

LWR

No new features or changes were introduced in this release.

LWR

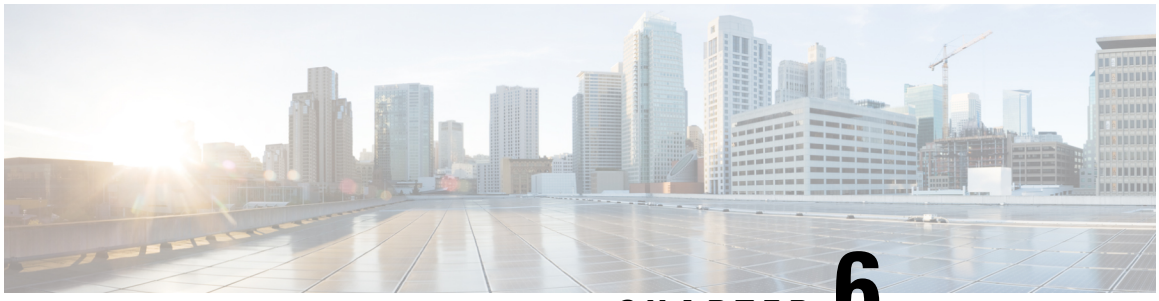

CHAPTER 6

Mobile

- Alarm Upon Detection of Missing/Corrupt Shards, on page 15
- MongoDB Health Monitoring for Write Operations, on page 16
- Prevention of Application from Using Missing/Corrupted Shards, on page 19
- RCS Feature Supressing Gx, on page 20
- Session Manager NTP Synchronization Improvement, on page 24
- Support for WPS Service on PCRF, on page 24

Alarm Upon Detection of Missing/Corrupt Shards

Feature Summary and Revision History

Table 9: Summary Data

Applicable Product(s) or Functional Area	CPS
Applicable Platform(s)	Not Applicable
Default Setting	Disabled - Configuration Required
Related Changes in This Release	Not Applicable
Related Documentation	CPS SNMP, Alarms, and Clearing Procedures Guide CPS Troubleshooting Guide

Table 10: Revision History

Revision Details	Release
First introduced	19.4.0

Feature Description

CPS has been enhanced to generate alarms and let the service provider operators know:

- When the sharding database in the ADMIN replica-set is missing session shard entries in the HA/GR environment

- When a shard is created and an entry exists in the sharding database of the ADMIN replica-set, but Session Manager VM is not reachable
- When indexes are missing on collections existing on SPR/Session database

The following new alarms have been added:

- SESSION_SHARD_UNREACHABLE
- ADMIN_DB_MISSING_SHARD_ENTRIES
- MISSING_SESSION_INDEXES
- MISSING_SPR_INDEXES

A new `qns.conf` parameter `skipUnreachableShards` has been added. By default, the value is set to false. The value must be set to true so that the above mentioned alarms are generated. Contact your Cisco Account representative for information on `qns.conf` file parameters.

For more information, see the following sections:

- *Application Notifications* table in the *CPS SNMP, Alarms, and Clearing Procedures Guide*
- *Clearing Procedures* chapter in the *CPS SNMP, Alarms, and Clearing Procedures Guide*
- *Testing Traps Generated by CPS* in the *CPS Troubleshooting Guide*

MongoDB Health Monitoring for Write Operations

Feature Summary and Revision History

Table 11: Summary Data

Applicable Product(s) or Functional Area	CPS
Applicable Platform(s)	Not Applicable
Default Setting	Disabled - Configuration Required
Related Changes in This Release	Not Applicable
Related Documentation	CPS SNMP, Alarms, and Clearing Procedures Guide CPS Troubleshooting Guide CPS Installation Guide for VMware CPS Installation Guide for OpenStack

Table 12: Revision History

Revision Details	Release
First introduced	19.4.0

Feature Description

CPS is enhanced to support MongoDB health monitor for write operations where application detects the false reporting of no PRIMARY connection available for the replica-sets and takes action accordingly.

Note The scope is limited to ADMIN, SESSION, SPR (excluding remote database), and Balance (excluding remote database) replica-sets.

The process of monitoring MongoDB write operations is divided into two sub-parts: Application side and Platform side.

- On the application side, MongoDB write operation and Reset MongoDB client connection threads are initialized to perform the write operations on replica-sets. This is done to verify the write operation failure threshold reached or not. If threshold has reached, MongoDB client connections are reset. Also, generates the counters for all the configured replica-sets.
- The platform scripts fetches the counters generated by the application and checks whether the failure counter for write operation is incrementing, and counter for reset MongoDB client connection fails after maximum retry of reset attempts. If parameter `autoheal_qns_enabled` value is set to **TRUE**, all the configured replica-sets sends an alarm and restarts the application (QNS process on QNS VMs) where the alarm is reported. If parameter value is set to **FALSE**, then the alarm is raised for all the configured replica-sets.

For more information on `autoheal_qns_enabled` configuration, see the following sections:

- *Installing Platform Scripts for MongoDB Health Monitoring - VMware* in the *CPS Installation Guide for VMware*
- *Installing Platform Scripts for MongoDB Health Monitoring - OpenStack* in the *CPS Installation Guide for OpenStack*

The following parameters should be configured in `qns.conf` file.

Table 13: qns.conf Parameters

Parameter	Description
mongo.health.monitor.enabled	This parameter is used to enable MongoDB health monitor for write operations. Default value is false.
mongo.health.monitor.scheduler.period	This parameter used to configure how frequently to execute the MongoDB write operation health monitor thread. Default value is 3000 milliseconds.
mongo.reset.scheduler.period	This parameter is used to configure how frequently to execute the thread for checking the reset MongoDB client connections in case of write operations are breached the failure threshold. Default value is 4000 milliseconds.

Parameter	Description
mongo.monitor.write.update.threshold	This parameter is used to configure number of MongoDB write operations to be performed on the database every time the thread executes. Default value is 50.
mongo.monitor.percentage.threshold	This parameter is used to calculate the success threshold percentage rate for MongoDB write operations. If success result is less than configured threshold, then only MongoDB client rest operation is triggered. Default value is 80.
wait.time.after.reset	This parameter is used to configure the time that thread for resetting MongoDB client connection should wait/sleep for after connection reset. Default value is 10000 milliseconds.
mongo.reset.retry.counter	This parameter is used to configure maximum retry attempts for resetting MongoDB client connections in case of previous reset failure. Default value is 2.

The following new statistics are added:

- mongodb_heartbeat_update_success_<replicaSetName>
- mongodb_heartbeat_update_error_<replicaSetName>
- mongo_client_reset_<replicaSetName>
- mongo_client_reset_fail_<replicasetSetName>

For more information on statistics, see Statistics/KPI Additions or Changes, on page 31.

Also, a new alarm, `Database Operations` is added.

For more information, see the following sections:

- *Application Notifications* table in the *CPS SNMP, Alarms, and Clearing Procedures Guide*
- *Clearing Procedures* chapter in the *CPS SNMP, Alarms, and Clearing Procedures Guide*
- *Testing Traps Generated by CPS* in the *CPS Troubleshooting Guide*
- *Installing Platform Scripts for MongoDB Health Monitoring - VMware* in the *CPS Installation Guide for VMware*
- *Installing Platform Scripts for MongoDB Health Monitoring - OpenStack* in the *CPS Installation Guide for OpenStack*

Performance Impact

- When the MongoDB health monitoring is enabled, application initiates two thread for each Session, ADMIN, Balance (excluding remote database replica-sets), SPR (excluding remote database replica-sets) replica sets. The number of threads depends on the number of replica-sets.
- In the scaled setups, as there are more number of replica-sets, you will observe that there is an increase in CPU utilization by 4 - 5%. The increase in the CPU utilization varies if the environment has large number of replica-sets.

Prevention of Application from Using Missing/Corrupted Shards

Feature Summary and Revision History

Table 14: Summary Data

Applicable Product(s) or Functional Area	CPS
Applicable Platform(s)	Not Applicable
Default Setting	Disabled - Configuration Required
Related Changes in This Release	Not Applicable
Related Documentation	Not Applicable

Table 15: Revision History

Revision Details	Release
First introduced	19.4.0

Feature Description

CPS can now detect missing or unrecognizable secondary session shards and direct the applications not to use them for its CRUD (create, read, update and delete) operations.

This excludes arbiter members and hot standby/backup members as arbiter nodes do not store any data as well as participate in application related CRUD operations. The hot standby/backup VMs are not used in normal traffic operations unless there is a failover/election in progress.

The following scenarios describe when a session shard of a secondary member becomes unreachable:

- When the sessionmgr VM is shutdown. All mongo process running on that VM including the one hosting session information is down, which means all the shards are not reachable.
- The sessionmgr VM is up, but the mongo process is shutdown gracefully:
  - Using *monit stop aido_client* followed by */etc/init.d/sessionmgr-portno. stop* (for instance)
  - The mongo process is killed abruptly after *monit stop aido_client*

  In this case, all the session shards are not reachable on that secondary member.

- Another variant could be that the shards were created but later dropped from the mongo CLI using:

  `session_cache_*, db.dropDatabase()` command on the secondary members

New `qns.conf` parameters, `skipUnreachableShards` and `skipDbOperOnUnreachableShards` are added. Both the parameters must be set to *true* to skip database operations on unreachable shards.

- `skipUnreachableShards`: By default, the value is set to false. The value must be set to true to generate or clear the alarms.
- `skipDbOperOnUnreachableShards`: By default, the value is set to false. The value must be set to true to enable the feature.

Contact your Cisco Account representative for information on `qns.conf` file parameters.

Restriction

The scope of the feature is restricted to prevent an application from using secondary member session shards when they are unreachable. This feature is not meant to eliminate CRUD operations on other databases such as, SPR, Balance, sk_cache and so on.

All session shards unreachable checks happen on secondary MongoDB members in a replica-set. The existing logic to detect unreachable primaries remain unaltered, which when not available writes to the hot standby database for a brief period. Then they are reconciled to newly elected primaries and flushed from hot standby database. No changes were done to the existing behavior related to primary members when unreachable.

RCS Feature Supressing Gx

Feature Summary and Revision History

Table 16: Summary Data

Applicable Product(s) or Functional Area	CPS
Applicable Platform(s)	Not Applicable
Default Setting	Enabled – Configuration Required
Related Changes in This Release	Not Applicable
Related Documentation	CPS Mobile Configuration Guide

Table 17: Revision History

Revision Details	Release
First introduced	19.4.0

Feature Description

Previous Behavior: In previous releases, PCRF created dedicated bearers during the RCS service. Creating dedicated bearers for each RCS service sometimes exceeded dedicated bearer limit in PGW and hence there was a degradation and interference with VoLTE calls.

New Behavior: In current release, CPS skips dedication bearer creation during AAR for RCS media type 'MESSAGE' and takes no action over the Gx interface for rule-installation.

For the first Rx transaction, there is no existing pre-configured rule. CRD lookup happens on table rx_dedicated_bearer_create. If result is **true**, then dedicated bearer is created. Otherwise, rule installation is skipped. For this feature, CRD must be configured to return false for media-type=6 so that the rule installation is skipped. There is no RAR on Gx interface for rule activation, but any other trigger (for example, new event-trigger added to Gx session based on specific-actions AVP on Rx_AAR) can trigger Gx RAR.

For the second or any upcoming transaction, CRD lookup is done based on existing service object **EvaluateRxDedicatedBearer** for each media sub-components and result is used in evaluating the existing pre-configured rule.

To support this feature, following new configurations have been added in Policy Builder:

- EvaluateRxDedicatedBearerCreate under **rx** service configuration objects

  EvaluateRxDedicatedBearerCreate service object is introduced which should be added under service configuration object. While configuring the service option, new STG table should be linked to this object.

- rx_dedicated_bearer_create search table for CRD lookup
- rx_bearer_evaluation_create search table to evaluate rule for the first transaction

Note As CRD lookup happens for AAR-I, there can be an increase in Rx response time.

For more information, see *EvaluateRxDedicatedBearerCreate* section in the *CPS Mobile Configuration Guide*.

Configuration

Service object **EvaluateRxDedicatedBearerCreate** has been added which must be binded to new STG. In New CRD/STG table, one must set output column to 'false' for media-type:6 so that dedicated bearer is not created.

Sample Policy Builder and CRD changes are mentioned below (table name, column names can be changed).

1. **Policy Builder - Search Table Groups**

   Adding new STG: rx_dedicated_bearer_create. Evaluation order for this STG must be new (evaluation order configured can be different for new CRD version).

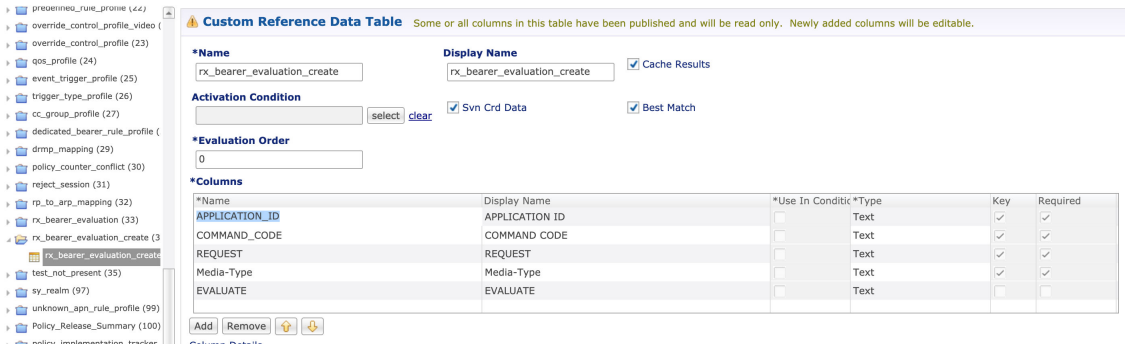

2. Policy Builder - Service Configuration

1. Use case template:

   Add newly created service object EvaluateRxDedicatedBearerCreate in the existing use case template:

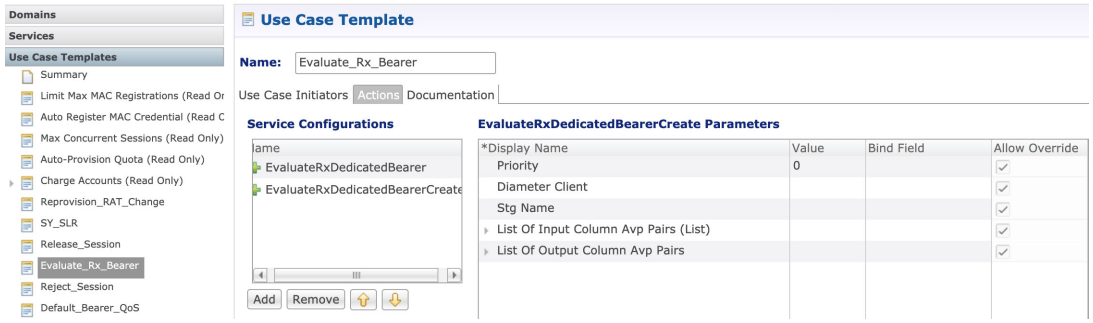

2. Service option:

   Existing:

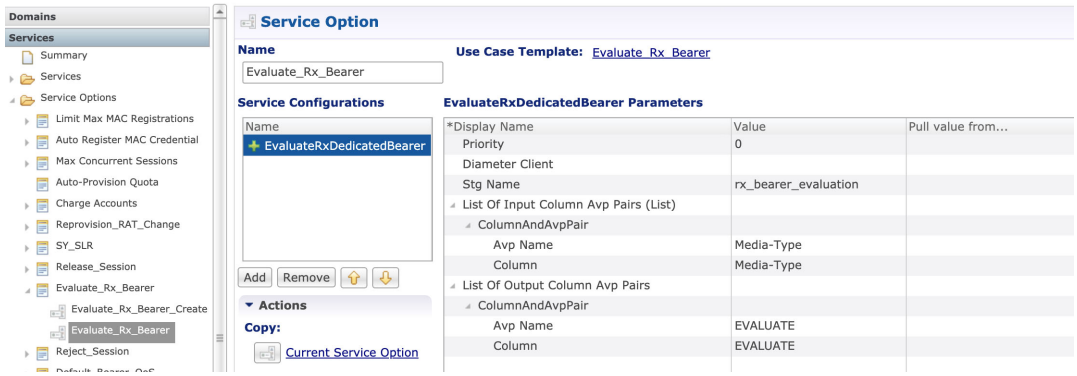

   Add new service object and bind it to new STG. Also bind the input and output column to new CRD table (Media-Type, Evaluate of rx_bearer_evaluation_create).

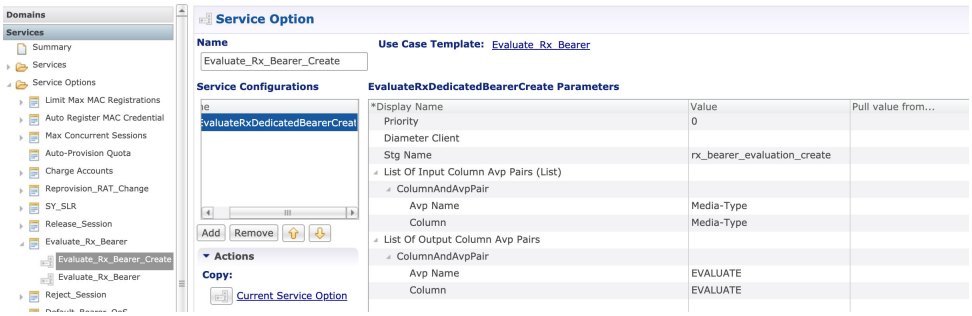

3. Service binding:

   New service option must be added to service.

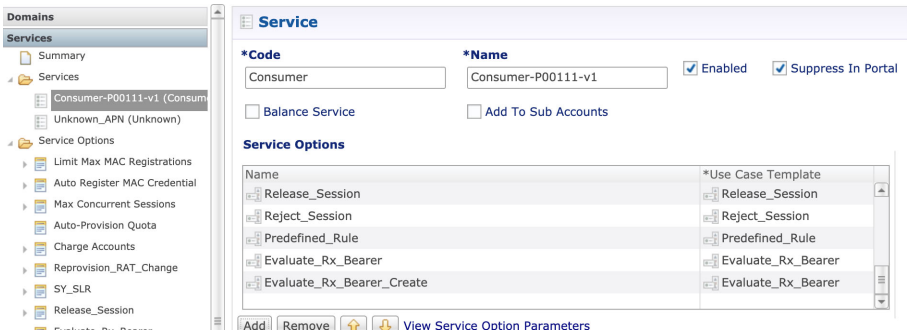

3. CRD Data: rx_bearer_evalution table data

   Existing design contains “Evaluate” as output column which is used only during update (can be CCR-U/2nd AAR/Sy messages)

   1. Existing CRD table:

      The configuration shown below is set Evaluation to false when CPS receive a Sy SSNR or SLA.

      One more entry must be added for media-type: 6 and output should be set to false so that bearer evaluation is skipped (however, bearer is not created).

rx_bearer_evaluation

APPLICATION ID	COMMAND CODE	REQUEST	Media-Type	EVALUATE	imp
16777302	8388636	true	1	false	
16777302	8388635	true	1	false	
16777236	265	true	1	true	
16777238	272	true	1	true	
*	*	*	1	false	
*	*	*	6	false	

   2. New CRD table:

      The **rx_bearer_evaluation_create** search table is created to evaluate rule for the first transaction and decide if dedicated bearer must be created.

      Output must be set to false for media-type: 6 (MESSAGE) so that rule is not created.

      For all other media-type, it should be true. If there are no entries for specific media-type, the value is set to true (by default) and rule is created for each media sub-component if present.

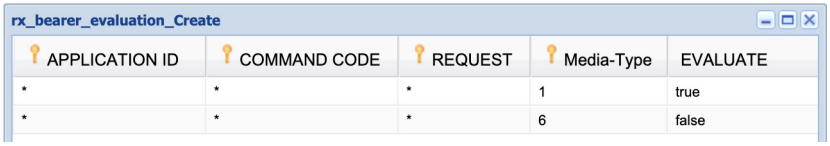
rx_bearer_evaluation_Create

APPLICATION ID	COMMAND CODE	REQUEST	Media-Type	EVALUATE
*	*	*	1	true
*	*	*	6	false

Session Manager NTP Synchronization Improvement

Feature Summary and Revision History

Table 18: Summary Data

Applicable Product(s) or Functional Area	CPS
Applicable Platform(s)	Not Applicable
Default Setting	Not Applicable
Related Changes in This Release	Not Applicable
Related Documentation	Not Applicable

Table 19: Revision History

Revision Details	Release
First introduced	19.4.0

Feature Description

Previous Behavior: In previous releases, clock synchronization did not happen between Session Managers and Policy Director (LB01) for fresh installations. Also, NTP was not in sync and TIME_WAIT increased after sessionmgr VM was restarted or suspended.

New Behavior: In current release, the issue has been fixed and NTP time synchronization happens between Session Managers and Policy Director (LB) VM of HA/GR/Dual cluster during deployment/reboot/powered-on scenarios.

As a part of fix, MongoDB on Session Manager does not start till loop for syncing NTP is completed.

Support for WPS Service on PCRF

Feature Summary and Revision History

Table 20: Summary Data

Applicable Product(s) or Functional Area	UDC
Applicable Platform(s)	Not Applicable

Default Setting	Enabled – Configuration Required
Related Changes in This Release	Not Applicable
Related Documentation	CPS UDC Administration Guide

Table 21: Revision History

Revision Details	Release
First introduced	19.4.0

Feature Description

Important This feature is not fully qualified in this release, and is available only for testing purposes. It is planned to be qualified in the subsequent release. For more information, contact your Cisco Account representative.

CPS now provides support for WPS service on PCRF using LWR. The following new UDC service object is introduced to configure the Rx message AVPs which can be used to generate a LDAP attribute:

- BindUDCAttributeForRx

LWR replicated LDAP attribute generated using service BindUDCAttributeForRx can become stale during certain error scenarios. In such scenarios, a Guard Time can be specified to ignore the AVP if it is older than the configured time. While processing any event, QNS evaluates the staleness of the configured attribute and does not use it for policy evaluations.

CPS also supports modification of QoS in WPS usecase if the UE sessions for different APNs are created on different CPS datacenters.

Note This feature can only be supported when LWR replication is configured for cross-site replication.

For more information, see *Configuring WPS Service on PCRF* section in the *CPS UDC Administration Guide*.

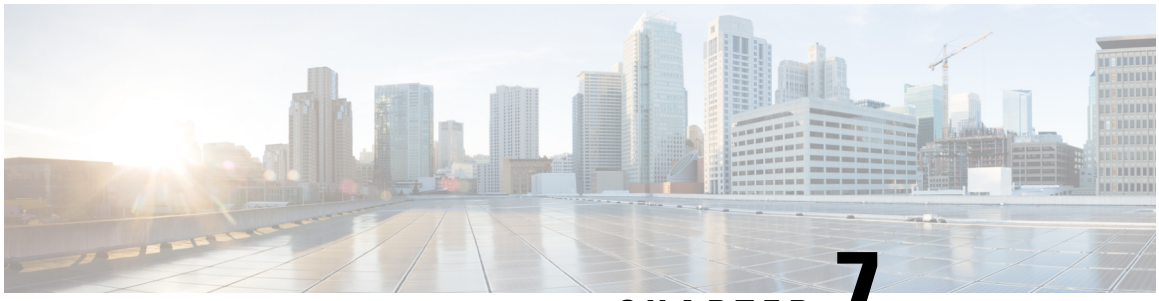

CHAPTER 7

MOG

- MOG, on page 27

MOG

No new features or changes were introduced in this release.

MOG

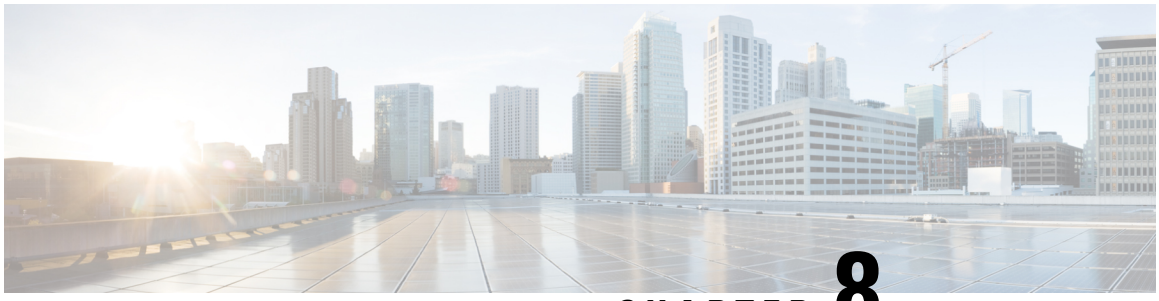

CHAPTER 8

Operations

- API Additions or Changes, on page 29
- MIB Additions or Changes, on page 29
- Log Additions or Changes, on page 29
- SNMP Alarm Additions or Changes, on page 31
- Statistics/KPI Additions or Changes, on page 31

API Additions or Changes

No changes were introduced in this release.

MIB Additions or Changes

No changes were introduced in this release.

Log Additions or Changes

Policy Builder and Control Center Activity Logs

Feature Summary and Revision History

Table 22: Summary Data

Applicable Product(s) or Functional Area	CPS
Applicable Platform(s)	Not Applicable
Default Setting	Enabled - Always ON
Related Changes in This Release	Not Applicable
Related Documentation	Not Applicable

Table 23: Revision History

Revision Details	Release
First introduced	19.4.0

Policy Builder Logging

- Login and logout message in audit logs is now written into separate audit log for easy tracing.
  - File location: `/var/log/broadhop/qns-audit-pb.log`
  - Logs are available in pcrfclient VMs.
  - When the log file reach 20 MB, it gets rotated. Maximum of five latest log files are available at a specific time.

  You need to enable the log information in `/etc/broadhop/logback-pb.xml` file.

```
==============================================================
<!-- UI Activity Loggers -->
<logger name="com.broadhop.client.WorkspaceChooserDialog" level="info"><appender-ref
ref="UI-ACTIVITY" /></logger>
<logger name="com.broadhop.client.ui.framework.handlers.ExitHandler"
level="info"><appender-ref ref="UI-ACTIVITY" /></logger>
<!-- UI Activity Loggers -->
==============================================================
```

- **Policy Builder publish logs in audit database:** User name is updated into the existing audit database entry.

  Enable Audit Database

  1. Enable the audit database logging by configuring the parameter in `/etc/broadhop/pb/pb.conf` file.

     ```
     -Dua.client.submit.audit=true
     ```

  2. Select the checkbox, **Log Read Request** in Policy Builder under **Systems** > ***system name*** > **Plugin Configuration** > **Audit Configuration**. For more information, see *Audit Configuration* sections in the *CPS Mobile Configuration Guide*.

Control Center Logging

- Login and logout message in audit logs is now written into separate audit log for easy tracing.
  - File location: `/var/log/broadhop/qns-audit-1.log`
  - Logs are available in Policy Server (QNS) VM.
  - When the log file reach 20 MB, it gets rotated. Maximum of five latest log files are available at a specific time.

  You need to enable the log information in `/etc/broadhop/logback-pb.xml` file.

```
<!-- CC Login Logout -->
  <logger name="com.broadhop.ui.security.server.SessionConcurrencyManager"
level="info"><appender-ref ref="UI-ACTIVITY" /></logger>
<!-- CC Login Logout -->
```

- By default, Control Center activity logs are captured in Audit database.

SNMP Alarm Additions or Changes

The following table provides information on new/modified alarms:

Table 24: Alarm Additions or Changes

New/Modified Alarms	Release Introduced/ Modified	Applicable Product(s)/
SESSION_SHARD_UNREACHABLE	19.4.0	CPS
ADMIN_DB_MISSING_ SHARD_ENTRIES	19.4.0	CPS
MISSING_SESSION_INDEXES	19.4.0	CPS
MISSING_SPR_INDEXES	19.4.0	CPS
ProcessRestarted	19.4.0	CPS
Database Operation	19.4.0	CPS
SESSION_DB_LIMIT_EXCEEDED	19.4.0	vDRA
IPV6_DB_LIMIT_EXCEEDED	19.4.0	vDRA
IPV4_DB_LIMIT_EXCEEDED	19.4.0	vDRA
IMSIAPN_DB_ LIMIT_EXCEEDED	19.4.0	vDRA
MSISDNAPN_DB_ LIMIT_EXCEEDED	19.4.0	vDRA

For more information, see the following sections:

- *Application Notifications* table in the *CPS SNMP, Alarms, and Clearing Procedures Guide*
- *Clearing Procedures* chapter in the *CPS SNMP, Alarms, and Clearing Procedures Guide*
- *Testing Traps Generated by CPS* in the *CPS Troubleshooting Guide*
- *Application Notifications* table in the *CPS vDRA SNMP and Alarms Guide*

Statistics/KPI Additions or Changes

The following table provides information on new/modified statistics:

Table 25: Statistics Additions or Changes

New/Modified Statistics	Description	Applicable Product(s)/
ccri_db_limit_breached_total	This statistics displays the total number of requests rejected/forwarded due to database record limit breach. Request are forwarded to database in case of database limit breach and bindings are marked as best effort. Field in statistics: binding_type = session/ipv6/ipv4/ imsi_apn/msisdn_apn status = discard/forward	vDRA
db_stats_per_shard	The gauge statistics has been created for getting the number of records in each shard. This can be used to monitor the record distribution across shards (as seen from DRA VNF side). db_stats_per_shard Filter: sharddb – database name (drasessions, ipvbindings and so on) replicaSet – replica-set name for the physical shard (as configured in database at binding VNF side). **Sample query to get the record distribution across shards for session database:** *sum(db_stats_per_shard* *{sharddb="drasessions", replicaSet!=""})* *by (replicaSet)*	vDRA
ipv6_zone_based_ db_message_total	The IPv6 zone based database message total for all (insert/update/read/delete) database operations. Label: zone Type: Counter The source is JMX/Script	vDRA

New/Modified Statistics	Description	Applicable Product(s)/
relay_message_total	This statistics shows the total number of messages relayed on relay channel. Field in statistics: • endpoint: endpoint • relay_peer: relay peer • app_id: Application ID • direction: ingress/egress • message_type: request/answer	vDRA
priority_aar_count	This counter value increase when main thread pool queue gets full and AAR specific thread pool tarts processing AARs	vDRA
db_query_*	This counter is now migrated to the new metric style and two new statistics are added. db_query_total and following labels are added: • db (session/ipv6/ipv4/imsiapn/msisdnapn) • status (success/success.noentry/failed) • read_perf (primary/nearest/nearest+primary) • shard_type (mongo-sharding/app-sharding) db_query_duration_seconds with labels: • db (session/ipv6/ipv4/imsiapn/msisdnapn) • status (success/success.noentry/failed) • shard_type (mongo-sharding/app-sharding)	vDRA
mongodb_heartbeat_update_ success_<replicaSetName>	This counter continues increments when update operation is successful with respect to corresponding replica-set only.	CPS
mongodb_heartbeat_update _error_<replicaSetName>	This counter continues increments when update operation fails with respect to corresponding replica-set only.	CPS

New/Modified Statistics	Description	Applicable Product(s)/
mongo_client_reset _<replicaSetName>	This counter continues increments when mongo client establishes connection to the application re-initiated with respect to corresponding replica-set only.	CPS
mongo_client_reset_ fail_<replicasetSetName>	This counter continues increments when mongo client connection to the application re-initiated fails with respect to corresponding replica-set only. When it reaches the number of retry reset operation this counter stops increments.	CPS

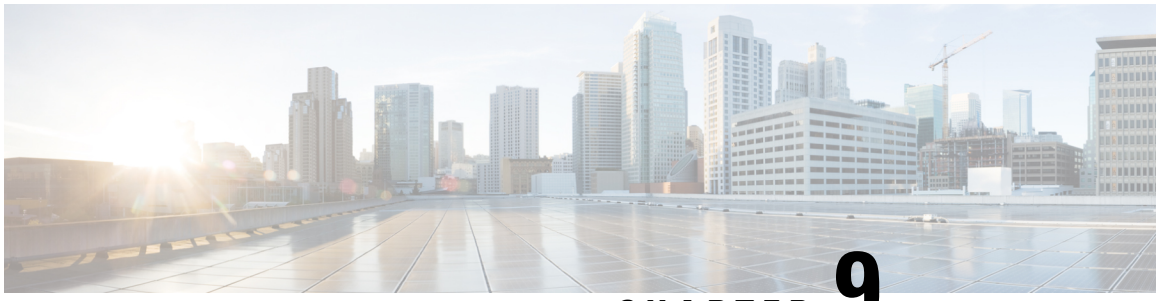

CHAPTER 9

Performance Improvement

• Performance Improvement, on page 35

Performance Improvement

No new features or changes were introduced in this release.

Performance Improvement

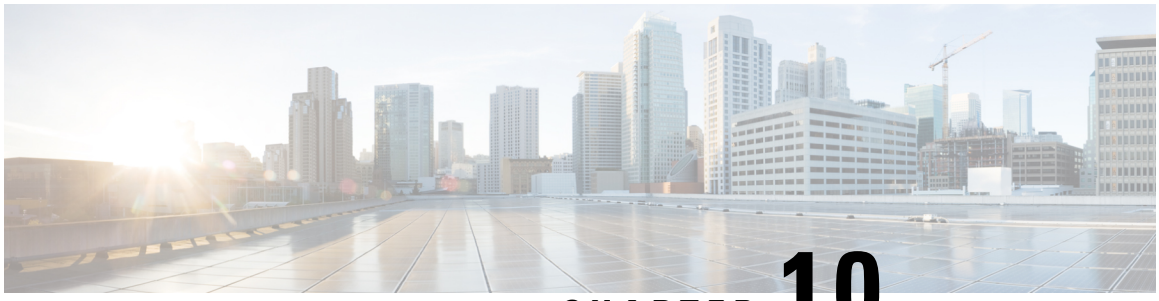

CHAPTER 10

Platform

- Send SNMP Traps for Process Restarts, on page 37
- SNMP Alarm Resynchronization, on page 38
- SR-IOV Support for CPS, on page 39

Send SNMP Traps for Process Restarts

Feature Summary and Revision History

Table 26: Summary Data

Applicable Product(s) or Functional Area	CPS
Applicable Platform(s)	Not Applicable
Default Setting	Enabled - Always-on
Related Changes in This Release	Not Applicable
Related Documentation	CPS SNMP, Alarms, and Clearing Procedures Guide CPS Troubleshooting Guide

Table 27: Revision History

Revision Details	Release
First introduced	19.4.0

Feature Description

CPS now supports a process restart event whenever qns process is restarted on the system.

To support this, a new alarm `ProcessRestarted` is introduced.

CPS must send process restarted snmp traps whenever qns process gets restarted on the system.

For more information, see the following sections:

- *Application Notifications* table in the *CPS SNMP, Alarms, and Clearing Procedures Guide*
- *Clearing Procedures* chapter in the *CPS SNMP, Alarms, and Clearing Procedures Guide*
- *Testing Traps Generated by CPS* in the *CPS Troubleshooting Guide*

SNMP Alarm Resynchronization

Feature Summary and Revision History

Table 28: Summary Data

Applicable Product(s) or Functional Area	CPS
Applicable Platform(s)	Not Applicable
Default Setting	Disabled - Configuration Required
Related Changes in This Release	Not Applicable
Related Documentation	CPS Installation Guide for VMware

Table 29: Revision History

Revision Details	Release
First introduced	19.4.0

Feature Description

CPS now supports holding the alarms when NMS is down or unreachable and sends the alarms to NMS once it is reachable.

If NMS is not available, the alarms are stored in Admin DB (DB Name: CPSAlarmResync, collection names: CpsCompAlarm, CpsAppAlarm) and sent only when the NMS is up.

To support this, a new parameter, *alarm_resync_enabled* has been added in `Configuration.csv` file. By default, the value is false.

Note Currently, alarm synchronization is supported only for VMware environment.

- When alarm_resync feature is enabled:
  - If NMS is up, the NMS receives all the alarms.
  - If NMS is down, NMS does not receive any alarms.

    The alarms are stored in alarm database. Once the NMS is in reachable state, all the alarms stored in alarm database are forwarded to NMS.

- When alarm_resync feature is disabled:
  - If NMS is up/down, the alarms are forwarded to NMS without checking NMS server's availability.
  - If alarm_resync feature is disabled after having been enabled for sometime.

    Any older alarms stored in alarm database are forwarded to NMS without checking its reachability.

    New alarms generated from all VMs are forwarded to NMS without checking its reachability.

For more information, see *General Configuration Parameters* table in the *CPS Installation Guide for VMware*.

Memory and Performance Impact

- If the alarms frequency increases while NMS is down, then additional disk space is required in Session Manager VMs.
- Once NMS is reachable, alarms are forwarded to NMS server sequentially from active Policy Director (LB). This can result in performance impact on the Policy Director (LB) VM.

Limitations

- If critical trap is generated for particular alarm while alarm resync feature is enabled and NMS is UP, the alarm is forwarded to NMS. After sometime when NMS is down and clear trap is generated for the same alarm, the clear alarm does not get stored in alarm database (critical alarms only get inserted into alarm database) and is not sent to NMS. When NMS is UP, the feature sends the existing active alarms to NMS. In such cases, the NMS misses the clear alarm for the critical alarm already received.
- If alarm_resync feature is disabled after enabled, existing alarms stored in database while NMS is down is forwarded to NMS without checking its reachability to avoid stale alarms. This is done to maintain in-order delivery of alarms to NMS. If the alarm_resync feature stores alarms even after feature is disabled, there is a possibility of new alarms reaching the NMS before alarm_resync feature sends the alarms stored in database to NMS. To avoid such stale alarm scenarios, alarm_resync feature sends out the alarms to NMS immediately when it is disabled.
- **VirtualInterfaceDown, VirtualInterfaceUp**: This particluar alarm for internalvip (lbvip02) is not stored in AlarmResync DB. This alarm is sent directly from gen_noti_to_nms function to NMS server.

SR-IOV Support for CPS

Feature Summary and Revision History

Table 30: Summary Data

Applicable Product(s) or Functional Area	CPS
Applicable Platform(s)	Not Applicable
Default Setting	Enabled - Always-on
Related Changes in This Release	Not Applicable

Related Documentation	CPS Installation Guide for OpenStack

Table 31: Revision History

Revision Details	Release
First introduced	19.4.0

Feature Description

CPS now supports SRIOV interface on CPS VMs with XL710 Intel NIC interface. It also supports bonding on SRIOV VFs on guest VM.

- If all the guest VM interfaces are SR-IOV interface then *ifrename.yaml* is not required.
- If multiple drivers are used, then *ifrename.yaml* file must be updated with corresponding driver. For example, I40evf for XL710
- Bonding can be created on two different virtual functions. The virtual functions can be created from same physical function or different physical function in the host based on the requirements.

Note Currently, SR-IOV is only supported for OpenStack deployments. SR-IOV interface is supported only in fresh deployment since all the VMs need to redeployed with new network and interface configured.

For more information, see *SR-IOV Support* section in the *CPS Installation Guide for OpenStack*.

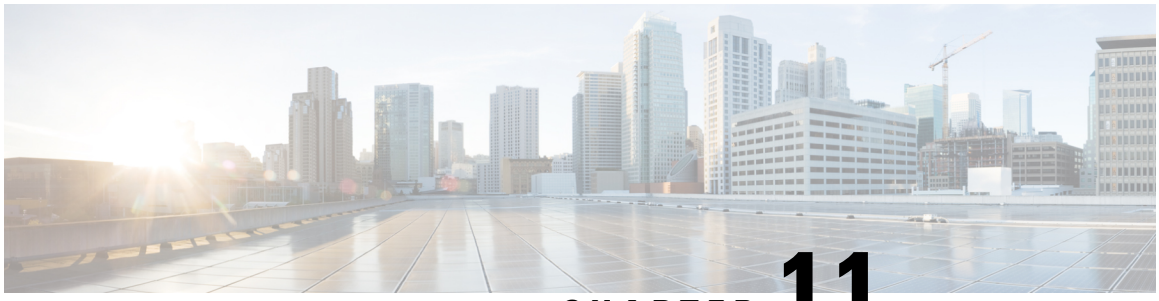

CHAPTER 11

Policy Reporting

• Policy Reporting, on page 41

Policy Reporting

No new features or changes were introduced in this release.

Policy Reporting

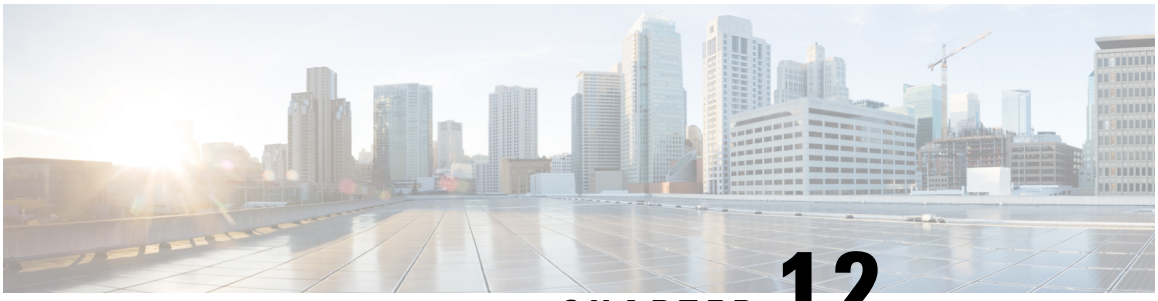

CHAPTER 12

Product Security

• CentOS Security Enhancements/Kernel Upgrade, on page 43

CentOS Security Enhancements/Kernel Upgrade

Feature Summary and Revision History

Table 32: Summary Data

Applicable Product(s) or Functional Area	CPS
Applicable Platform(s)	Not Applicable
Default Setting	Not Applicable
Related Changes in This Release	Not Applicable
Related Documentation	Not Applicable

Table 33: Revision History

Revision Details	Release
Kernel upgraded to 3.10.0-957.12.2.el7 Grafana upgraded to 6.2.2-1	19.4.0
CentOS upgraded to 7.6 (1810) Kernel upgraded to 3.10.0-957.10.1.el7	19.3.0
Kernel upgraded to 3.10.0-957.5.1.el7	19.2.0
Kernel upgraded to 3.10.0-957.e17	19.1.0
First introduced: kernel upgraded to 3.10.0-862.14.4.el7.x86_64	18.5.0

Feature Description

In this release, the following upgrades have been done to fix the vulnerabilities:

- Kernel upgraded from 3.10.0-957.10.1.el7 to 3.10.0-957.12.2.el7
- Grafana upgraded from 4.6.3-1 to 6.2.2-1

For service related issues, you can use `journactl` to get systemctl logs.

The following tables list the vulnerabilities that have been fixed as a part of this release:

Table 34: Cisco Internal Alert Manager (CIAM) CVEs

CVE	Name
CVE-2019-8324	ruby (CESA-2019:1235)
CVE-2019-8322	ruby (CESA-2019:1235)
CVE-2019-8323	ruby (CESA-2019:1235)
CVE-2019-8325	ruby (CESA-2019:1235)
CVE-2019-5953	wget (CESA-2019:1228)
CVE-2019-6974	kernel (CESA-2019:0818)
CVE-2019-7221	kernel (CESA-2019:0818)
CVE-2018-12130	kernel (CESA-2019:1168) (MDSUM/RIDL) (MFBDS/RIDL/ZombieLoad) (MLPDS/RIDL)
CVE-2018-12126	kernel (CESA-2019:1168) (MDSUM/RIDL) (MFBDS/RIDL/ZombieLoad) (MLPDS/RIDL)
CVE-2018-12127	kernel (CESA-2019:1168) (MDSUM/RIDL) (MFBDS/RIDL/ZombieLoad) (MLPDS/RIDL)
CVE-2019-11091	kernel (CESA-2019:1168) (MDSUM/RIDL) (MFBDS/RIDL/ZombieLoad) (MLPDS/RIDL)
CVE-2019-9636	python (CESA-2019:0710)
CVE-2016-10745	python-jinja2 (CESA-2019:1022)
CVE-2018-12130	libvirt (CESA-2019:1177) (MDSUM/RIDL) (MFBDS/RIDL/ZombieLoad) (MLPDS/RIDL)
CVE-2018-12126	libvirt (CESA-2019:1177) (MDSUM/RIDL) (MFBDS/RIDL/ZombieLoad) (MLPDS/RIDL)
CVE-2018-12127	libvirt (CESA-2019:1177) (MDSUM/RIDL) (MFBDS/RIDL/ZombieLoad) (MLPDS/RIDL)
CVE-2019-11091	libvirt (CESA-2019:1177) (MDSUM/RIDL) (MFBDS/RIDL/ZombieLoad) (MLPDS/RIDL)

Table 35: Nessus CVEs

CVE	Name
CVE-2019-8322	CentOS 7 : ruby (CESA-2019:1235)
CVE-2019-8323	CentOS 7 : ruby (CESA-2019:1235)

CVE	Name
CVE-2019-8324	CentOS 7 : ruby (CESA-2019:1235)
CVE-2019-8325	CentOS 7 : ruby (CESA-2019:1235)
CVE-2019-5953	CentOS 7 : wget (CESA-2019:1228)
CVE-2019-6974	CentOS 7 : kernel (CESA-2019:0818)
CVE-2019-7221	CentOS 7 : kernel (CESA-2019:0818)
CVE-2018-12130	CentOS 7 : kernel (CESA-2019:1168)
CVE-2018-12126	CentOS 7 : kernel (CESA-2019:1168)
CVE-2018-12127	CentOS 7 : kernel (CESA-2019:1168)
CVE-2019-11091	CentOS 7 : kernel (CESA-2019:1168)
CVE-2019-9636	CentOS 7 : python (CESA-2019:0710
CVE-2016-10745	CentOS 7 : python-jinja2 (CESA-2019:1022)
CVE-2018-12130	CentOS 7 : libvirt (CESA-2019:1177)
CVE-2018-12126	CentOS 7 : libvirt (CESA-2019:1177)
CVE-2018-12127	CentOS 7 : libvirt (CESA-2019:1177)
CVE-2019-11091	CentOS 7 : libvirt (CESA-2019:1177)
CVE-2018-12099	Grafana Dashboard Links Cross-Site Scripting Vulnerability
CVE-2018-15727	Grafana remember me Cookie Authentication Bypass Vulnerability
CVE-2018-19039	Grafana Arbitrary File Read Information Disclosure Vulnerability

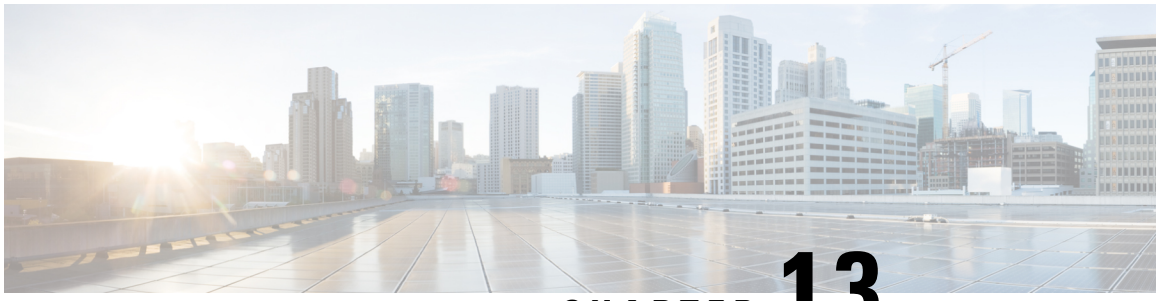

CHAPTER 13

Security Enhancements

- Security Enhancements, on page 47

Security Enhancements

This section lists enhancements introduced to support Cisco Product Security Requirements and the Product Security Baseline (PSB). For more information about Cisco Product Security Requirements, refer to: https://www.cisco.com/c/en/us/about/security-center/security-programs/secure-development-lifecycle/sdl-process.html

PSB Requirements for GUI

Feature Summary and Revision History

Table 36: Summary Data

Applicable Product(s) or Functional Area	CPS
Applicable Platform(s)	Not Applicable
Default Setting	Enabled - Always-on
Related Changes in This Release	Not Applicable
Related Documentation	Not Applicable

Table 37: Revision History

Revision Details	Release
First introduced	19.4.0

SEC-WEB-CLCKJACK-2: Prevent Click-Jacking

CPS web applications, Policy Builder, Control Center, and Central are now protected against click-jacking (SEC-WEB-CLCKJACK-2: Prevent Click-Jacking) attacks to prevent coerced user actions.

Click-jacking is a type of attack wherein the attacker tricks a user to click a button or link which is not visible to them and either redirecting them to attacker-controlled website or do an action that victim is not aware of. Therefore, the attacker is hijacking victim's click.

Note If the Policy Builder, Control Center, and Central applications are embedded to an iframe, then these applications fail to load and display the following message:

```
Blocked by Content Security Policy
```

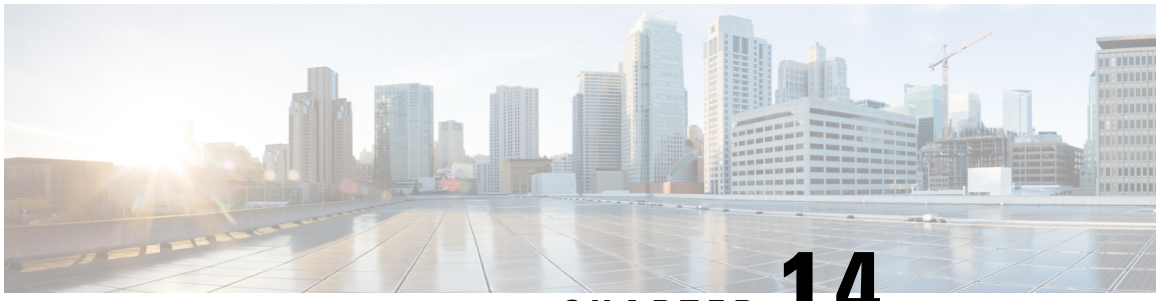

CHAPTER 14

UDC

• UDC, on page 49

UDC

Note When you upgrade/migrate or performing a fresh install for UDC setup, if required, you need to add the cluster ID in `/etc/broadhop/iomanager/qns_other.conf` file.

For example, `-Dcom.broadhop.run.clusterId=cluster-A`

Support for WPS Service on PCRF

Feature Summary and Revision History

Table 38: Summary Data

Applicable Product(s) or Functional Area	UDC
Applicable Platform(s)	Not Applicable
Default Setting	Enabled – Configuration Required
Related Changes in This Release	Not Applicable
Related Documentation	CPS UDC Administration Guide

Table 39: Revision History

Revision Details	Release
First introduced	19.4.0

Feature Description

Important This feature is not fully qualified in this release, and is available only for testing purposes. It is planned to be qualified in the subsequent release. For more information, contact your Cisco Account representative.

CPS now provides support for WPS service on PCRF using LWR. The following new UDC service object is introduced to configure the Rx message AVPs which can be used to generate a LDAP attribute:

- BindUDCAttributeForRx

LWR replicated LDAP attribute generated using service BindUDCAttributeForRx can become stale during certain error scenarios. In such scenarios, a Guard Time can be specified to ignore the AVP if it is older than the configured time. While processing any event, QNS evaluates the staleness of the configured attribute and does not use it for policy evaluations.

CPS also supports modification of QoS in WPS usecase if the UE sessions for different APNs are created on different CPS datacenters.

Note This feature can only be supported when LWR replication is configured for cross-site replication.

For more information, see *Configuring WPS Service on PCRF* section in the *CPS UDC Administration Guide*.

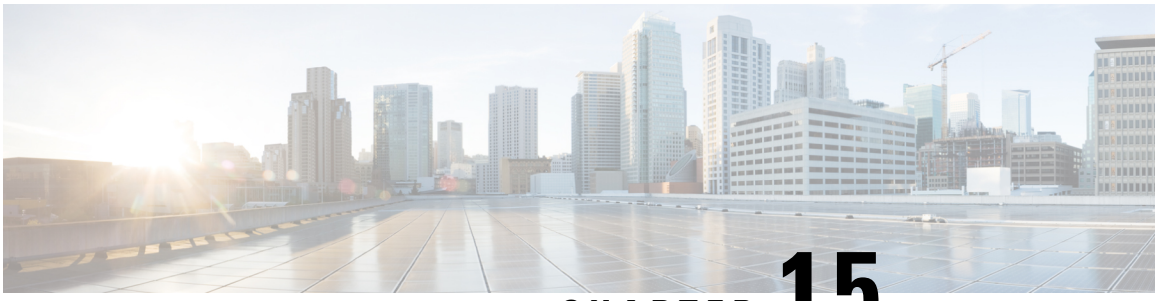

CHAPTER 15

UI Enhancements

- Policy Builder Enhancements, on page 51

Policy Builder Enhancements

Feature Summary and Revision History

Table 40: Summary Data

Applicable Product(s) or Functional Area	CPS
Applicable Platform(s)	Not Applicable
Default Setting	Disabled - Configuration Required
Related Changes in This Release	Not Applicable
Related Documentation	CPS Advanced Tuning Guide

Table 41: Revision History

Revision Details	Release
First introduced	19.4.0

Feature Description

Policy Builder is enhanced to display **Policies** tab for read-only users.

`pb.readOnly.showPolicies` parameter must be configured in `qns.conf` file to display the **Policies** tab.

When the admin user selects the checkbox under **Tools** > **Preferences** in Policy Builder and value of flag (`pb.readOnly.showPolicies`) is set to true, **Policies** tab appears in readonly mode.

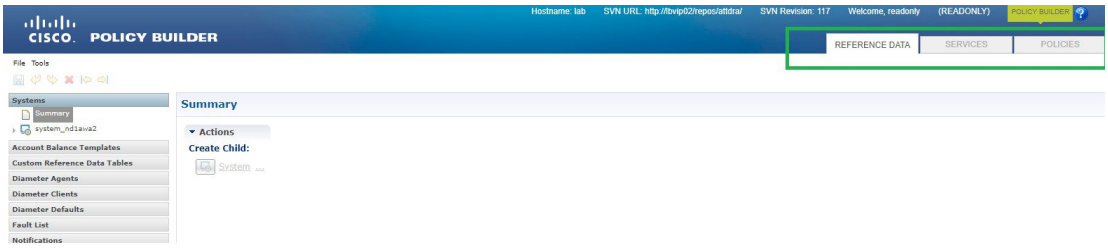

When the admin user selects the checkbox under **Tools** > **Preferences** in Policy Builder and value of flag (`pb.readOnly.showPolicies`) is set to false, **Policies** tab does not appear in readonly mode.

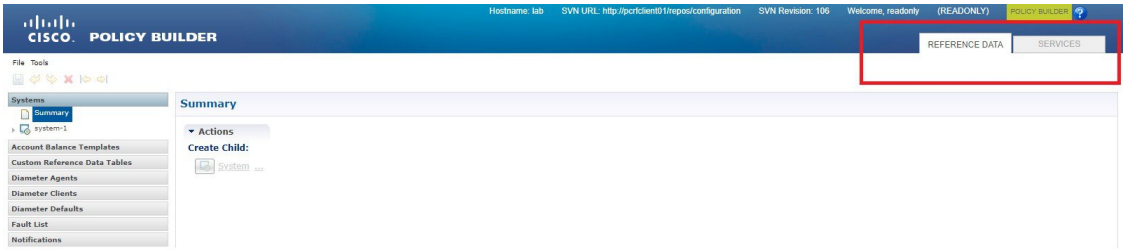

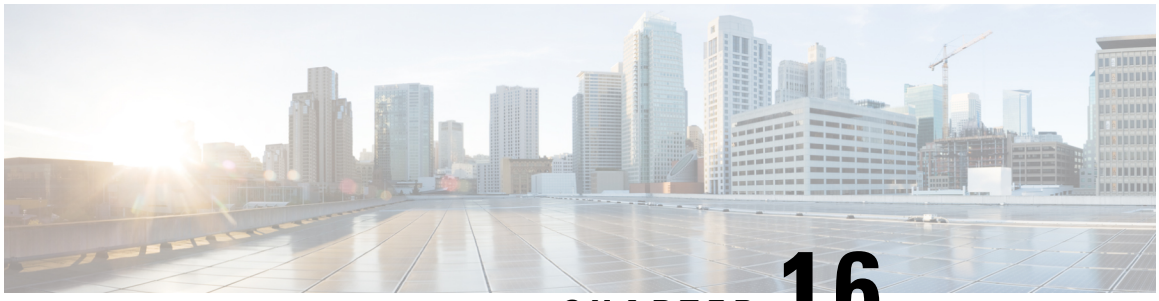

CHAPTER 16

vDRA

- Application based Client Sharding, on page 53
- Binding Database Record Capacity Overload Handling, on page 54
- DRA Distributor, on page 56
- DRA Distributor Usability Improvements, on page 56
- In-Service Migration from MongoDB Sharding to Application Sharding, on page 58
- Linux Utilities for Troubleshooting, on page 59
- Load Balance Diameter Messages across vDRA Relay Connections, on page 60
- MongoDB Service without Authentication Detection, on page 61
- Redis Security Vulnerabilities on vDRA, on page 62
- Routing Method Order Enhancements, on page 63
- Support to Create Mongo Shard/Replica-set, on page 64
- Support to Create Prometheus Statistics Partition, on page 65
- Upgrade MongoDB Version 3.6, on page 66
- vDRA Resiliency Improvements and Database Overload Protection, on page 67
- Zone-Aware Sharding, on page 67

Application based Client Sharding

Feature Summary and Revision History

Table 42: Summary Data

Applicable Product(s) or Functional Area	vDRA
Applicable Platform(s)	Not Applicable
Default Setting	Enabled - Always-on
Related Changes in This Release	Not Applicable
Related Documentation	CPS vDRA Configuration Guide CPS vDRA Operations Guide

Table 43: Revision History

Revision Details	Release
First introduced:	19.4.0

Feature Description

vDRA now supports application-based client sharding to partition/distribute session/binding data across multiple database replica-sets. This feature is implemented to scale larger-scale deployments with large numbers of replica-sets.

vDRA also supports regional IPv6 bindings (per mated-pair) in addition to a national binding database storing IMSI and MSISDN binding records.

For more information, see the following:

- `binding db-read-connection-settings` section in the *CPS vDRA Operations Guide*
- `binding shard-metadata-db-connection` section in the *CPS vDRA Operations Guide*
- `database cluster db-name sharding-db` *`name`* section in the *CPS vDRA Operations Guide*
- `database cluster db-name sharding-db-seed` *`name`* section in the *CPS vDRA Operations Guide*
- `database cluster db-name multi-db-collections` section in the *CPS vDRA Operations Guide*
- *Configuring Application based Sharding* section in the *CPS vDRA Configuration Guide*.

Binding Database Record Capacity Overload Handling

Feature Summary and Revision History

Table 44: Summary Data

Applicable Product(s) or Functional Area	vDRA
Applicable Platform(s)	Not Applicable
Default Setting	Enabled - Always-on
Related Changes in This Release	Not Applicable
Related Documentation	CPS vDRA Configuration Guide CPS vDRA Operations Guide CPS vDRA SNMP and Alarms Guide

Table 45: Revision History

Revision Details	Release
First introduced	19.4.0

Feature Description

vDRA supports enforcing the maximum number of records in a binding database (active + stale) so that vDRA remains stable in the event an excessive number of bindings are required.

To make sure that the system remains stable and continues to work, there is a limit in the session and binding records in a database. If the database record limit is breached for session/bindings, vDRA rejects the call with the configurable error code. For more information about the error code configuration, see *Configure Error Result Code Profile* in the *CPS vDRA Configuration Guide*.

When session and database record limit is breached and the best effort is enabled on binding keys, vDRA forwards call to PCRF without storing best effort binding keys in database.

Using `binding db-max-record-limit` command, you can configure maximum record limit on session and bindings.

If the database is mated pair, maximum record limit values must be same on both the sites.

If the maximum record limit is not configured or if configured to zero, there will be no database record limit on that binding type. By default, there is no database record limit configured on any of the binding types.

The following new alarms have been added:

- SESSION_DB_LIMIT_EXCEEDED
- IPV6_DB_LIMIT_EXCEEDED
- IPV4_DB_LIMIT_EXCEEDED
- IMSIAPN_DB_LIMIT_EXCEEDED
- MSISDNAPN_DB_LIMIT_EXCEEDED

For more information, see the following:

- *binding db-max-record-limit* section in the *CPS vDRA Operations Guide*
- *Application Notifications* table in the *CPS vDRA SNMP and Alarms Guide*
- *Configuring Binding Database Overload* section in the *CPS vDRA Configuration Guide*

Configuration and Restrictions

If vDRA is configured to send error response, in case of database overload error response might not be sent back to PCEF for all the rejected calls.

As the record count thread runs every 10 sec, vDRA does not always have a real time record count. When applying database limit over bindings, database limit breached condition takes effect with few more records in the database than the specified maximum limit.

DRA Distributor

Feature Summary and Revision History

Table 46: Summary Data

Applicable Product(s) or Functional Area	vDRA
Applicable Platform(s)	Not Applicable
Default Setting	Enabled - Always-on
Related Changes in This Release	Not Applicable
Related Documentation	Not Applicable

Table 47: Revision History

Revision Details	Release
First introduced	19.4.0

Feature Description

CPS is enhanced to support a DRA Distributor which is a traffic distributor that transparently does the following:

- Presents a virtual IP address representing a set of up to 16 application servers to application clients. This reduces the number of addresses with which the clients need to be configured.
- Routes packets from the clients that establish new TCP connections to selected application servers.
- Routes packets in existing TCP connections to the correct servers for the connection.

DRA Distributor Usability Improvements

Feature Summary and Revision History

Table 48: Summary Data

Applicable Product(s) or Functional Area	vDRA
Applicable Platform(s)	Not Applicable
Default Setting	Enabled - Always-on
Related Changes in This Release	Not Applicable
Related Documentation	CPS vDRA Operations Guide

Table 49: Revision History

Revision Details	Release
First introduced:	19.4.0

Automate VIP and Netfilter Configurations

Distributor IPv4/IPv6 services require a VIP configuration in the loopback interface and Netfilter rules for each VIP in the Director VMs. Previously, you had to manually add the Distributor IPv4 and IPv6 VIPs and associated Netfilter rules to the Director VMs.

Now vDRA supports an automated mechanism to configure the Director's VIPs, ARP tables, and IPv6 tables rules. Scripts are added in two new containers on the director VMs. These containers are included in the diameter-endpoint module.

- Real-server-monitor
- Real-server containers

Display Output of Distributor ipvsadm Command

New CLI command "*show dra-distributor*" has been added to the Orchestrator container to display the output of ipvsadm from all Distributor VMs.

For more information, see *show dra-distributor* section in the *CPS vDRA Operations Guide*.

Health Check Single Diameter Peers TCP Port

Diameter endpoint IP address and TCP port number are monitored to determine the health of the diameter endpoint at each Director VM supporting a Distributor Service. If the diameter IP address/port status fails, the Director VM is removed from the Distributor Service until the heath check passes.

Improvement in Grafana Dashboard

Two new panels, vDRA distributor statistics, and vDRA distributor VIPs are added to Grafana.

vDRA distributor statistics displays the following for each Distributor VM.

- Incoming Packet and bytes per second
- CPU Utilization
- Memory Utilization
- Connection Info
- Real Server Weights

vDRA distributor VIPs display the following information for each Distributor VM:

- Active connections per VIP
- Inactive connections per VIP

Configuration Parameter to Enable/Disable Preemption

DRA Distributor's Virtual IPs (VIPs) are managed by keepalived. Keepalived implements the Virtual Router Redundancy Protocol. VIPs have a priority of 1 – 254. The VIP with the highest priority becomes the Master. VRRP normally preempts a lower priority DRA Distributor VM when a higher priority DRA Distributor becomes available. In some circumstances it is desirable to not preempt lower priority DRA Distributor VM.

The preempt parameter, if set to false, disables preemption. By default, the preempt parameter is set to true.

Here is an example configuration:

```
network dra-distributor client
sync-id          110
sync-interface   ens192
tracking-service diameter-endpoint
preempt          true
preempt-delay    30
connection-timeout tcp 30
connection-timeout tcpfin 30
host 192.169.21.20
  priority 10
 !
 host 192.169.21.21
  priority 5
 !
 service Gx4
  virtual-router-id 121
  interface         ens224
  service-ip        192.169.22.50
  service-port      3868
  real-server 192.169.22.13
   weight 1
  !
  real-server 192.169.22.14
   weight 1
  !
 !
```

In-Service Migration from MongoDB Sharding to Application Sharding

vDRA now supports an in-service migration from vDRA running MongoDB sharding to a vDRA running Application Sharding so that there is no VoLTE calls drops and/or trigger an FCC or CRTC reportable outage.

Restriction

- During migration of a mated-pair, the GR resiliency is not fully available until both the sites are migrated. For example, if Site1 is migrated and all the new sessions for Site1 are getting stored in new application-sharded databases, then when Site1 goes down, the existing sessions from Site1 are not accessible from Site2 (as Site2 is still using only mongo-sharded cluster).
- vDRA does not support mixed databases with different shard-type. A DRA VNF site should not be configured where session database has mongo-based sharding and binding database has other type of sharding. The connections on all the databases (session and bindings) should be configured at the same time.
- When DRA VNF is configured with two database connections on mongo-sharded as well as application-sharded databases, then the session/binding expiration is performed on the records in the Primary DB (default primary application-sharded DB). MongoDB sharded database can be made as primary using `dra migration enable-mongo-sharded-db-as-primary-db` command

For more information, see *In-Service Migration from MongoDB Sharding to Application Sharding* chapter in the *CPS vDRA Installation Guide for VMware*.

Linux Utilities for Troubleshooting

Feature Summary and Revision History

Table 50: Summary Data

Applicable Product(s) or Functional Area	vDRA
Applicable Platform(s)	Not Applicable
Default Setting	Enabled - Always-on
Related Changes in This Release	Not Applicable
Related Documentation	Not Applicable

Table 51: Revision History

Revision Details	Release
First introduced:	19.4.0

Feature Description

The following linux utilities are added as a part of vDRA base_vmdk:

- traceroute
- pstree
- dig

- nslookup
- netcat
- htop
- ipset
- ndisc6
- rdisc6
- sosreport
- sysstat utilities (iostat, mpstat, pidstat and so on)
- tcpdump (tcpdump is also available in docker-host-info container other than VM. But if you need packet captures after running "system stop", tcpdump is no longer available from the container. The tcpdump utility available in VM can be used.)

Load Balance Diameter Messages across vDRA Relay Connections

Feature Summary and Revision History

Table 52: Summary Data

Applicable Product(s) or Functional Area	vDRA
Applicable Platform(s)	Not Applicable
Default Setting	Enabled - Always-on
Related Changes in This Release	Not Applicable
Related Documentation	Not Applicable

Table 53: Revision History

Revision Details	Release
First introduced:	19.4.0

Feature Description

In this release, vDRA supports load balance inter-PAS diameter requests across relay connections so that the higher TPS requests is achieved.

The following new statistics has been added:

- relay_message_total: This statistics shows the total number of messages relayed on relay channel.

Field in statistics:

- endpoint: endpoint
- relay_peer: relay peer
- app_id: Application ID
- direction: ingress/egress
- message_type: request/answer

MongoDB Service without Authentication Detection

Feature Summary and Revision History

Table 54: Summary Data

Applicable Product(s) or Functional Area	vDRA
Applicable Platform(s)	Not Applicable
Default Setting	Enabled - Always-on
Related Changes in This Release	Not Applicable
Related Documentation	CPS vDRA Configuration Guide CPS vDRA Operations Guide

Table 55: Revision History

Revision Details	Release
First introduced:	19.4.0

Feature Description

vDRA now supports password based authentication for MongoDB. Encrypted password is stored in consul database in key value format.

Note Role base authentication is not supported.

For more information, see the following:

- `db-authentication set-password database mongo password` section in the *CPS vDRA Operations Guide*
- `db-authentication enable-transition-auth database mongo` section in the *CPS vDRA Operations Guide*

- `db-authentication disable-transition-auth database mongo` section in the *CPS vDRA Operations Guide*
- `db-authentication rolling-restart database mongo` section in the *CPS vDRA Operations Guide*
- `db-authentication rolling-restart-status database mongo` section in the *CPS vDRA Operations Guide*
- `db-authentication remove-password database mongo` section in the *CPS vDRA Operations Guide*
- `db-authentication change-password database mongo` section in the *CPS vDRA Operations Guide*
- `db-authentication sync-password database mongo` section in the *CPS vDRA Operations Guide*
- *Configuring MongoDB Authentication* section in the *CPS vDRA Configuration Guide*

Redis Security Vulnerabilities on vDRA

Feature Summary and Revision History

Table 56: Summary Data

Applicable Product(s) or Functional Area	vDRA
Applicable Platform(s)	Not Applicable
Default Setting	Enabled - Configuration Required
Related Changes in This Release	Not Applicable
Related Documentation	CPS vDRA Operations Guide

Table 57: Revision History

Revision Details	Release
First introduced:	19.4.0

Feature Description

In this release, Redis is upgraded to 4.0.10-alpine version which supports password based authentication. The password can be configured through the confD CLI.

After successful upgrade to this release (all DRA's in network), enable Redis authentication. It is mandatory that Redis authentication be enabled only after successful upgrade on all sites in the network.

All the sites in the network must have the same password to connect with local/global control planes/Redis servers.

Important Password provisioning/changes must be done across all the sites in the network in one maintenance window.

For more information on Redis authentication using confD CLI, see the following sections:

- `db-authentication set-password database redis password` section in the *CPS vDRA Operations Guide*
- `db-authentication show-password database redis` section in the *CPS vDRA Operations Guide*
- `db-authentication remove-password database redis` section in the *CPS vDRA Operations Guide*

Routing Method Order Enhancements

Feature Summary and Revision History

Table 58: Summary Data

Applicable Product(s) or Functional Area	vDRA
Applicable Platform(s)	Not Applicable
Default Setting	Enabled - Always-on
Related Changes in This Release	Not Applicable
Related Documentation	Not Applicable

Table 59: Revision History

Revision Details	Release
First introduced:	19.4.0

Feature Description

In this release, routing order has been changed as follows:

- Dest-Host Routing (local topology)
- Dest-Host Routing (global topology)
- SRK routing (local topology)
- SRK routing (global topology)
- Table-Driven routing (use local/global topology to find the peer)

Support to Create Mongo Shard/Replica-set

Feature Summary and Revision History

Table 60: Summary Data

Applicable Product(s) or Functional Area	vDRA
Applicable Platform(s)	Not Applicable
Default Setting	Enabled - Always-on
Related Changes in This Release	Not Applicable
Related Documentation	Not Applicable

Table 61: Revision History

Revision Details	Release
First introduced:	19.4.0

Feature Description

vDRA now supports initializing of replica sets with minimum members present in NOT_INITIALIZED state. As the members from other site comes up, they automatically get added to replica set with appropriate state as Secondary or arbiter.

You need to configure database cluster configurations in each site only one time. Replica set gets initialized which results in available members with Primary, Secondary and Arbiter. Even though other members, which are not available from other sites, can still be in NO_CONNECTION state.

```
admin@orchestrator[pn-bind-master-0]# show database status | tab

                          CLUSTER
ADDRESS         PORT   NAME      STATUS          TYPE           NAME     SHARD   REPLICA SET
-------------------------------------------------------------------------------------------
192.169.67.28  27017  server-a  PRIMARY        replica_set    binding  shard-1   rs-shard-1
192.169.67.31  27017  server-b  SECONDARY      replica_set    binding  shard-1   rs-shard-1
192.169.67.32  27017  server-c  ARBITER        replica_set    binding  shard-1   rs-shard-1
192.169.67.14  27017  server-d  NO_CONNECTION  replica_set    binding  shard-1   rs-shard-1
192.169.67.15  27017  server-d  NO_CONNECTION  replica_set    binding  shard-1   rs-shard-1
```

Note Make sure minimum two members are added in replica set while creating database cluster.

Support to Create Prometheus Statistics Partition

Feature Summary and Revision History

Table 62: Summary Data

Applicable Product(s) or Functional Area	vDRA
Applicable Platform(s)	Not Applicable
Default Setting	Enabled - Always-on
Related Changes in This Release	Not Applicable
Related Documentation	Not Applicable

Table 63: Revision History

Revision Details	Release
First introduced:	19.4.0

Feature Description

During installation, the CPS deployer is used to create a data disk for the master, and two control VMs. Cloud-init partitions the data disk, creates a file system, and mounts the data disk. The storage layout for the Prometheus containers is hard coded in the orchestrator Prometheus module.

The new data disk layout is as follows:

- Partitions: The new data disk is 200GB and contains the following partitions.
  - /data (70 GB)
  - /stats (130 GB)
- Directory Structure

Description	Location
haproxy certificate	/data/certs/proxy.pem
Consul database	/data/consul-[123]/data
Confd network dns entries	/data/dnsmasq/etc/dnsmasq.conf
swarm.json	/data/install/swarm.json
ISOs	/data/isos/
Keepalived metadata	/data/keepalived/
Mongo admin-db	/data/mongo-admin-[abc]/

Description	Location
Mongod-node	/data/mongod-node/
Mongo-s10x	/data/mongo-s10x/
Orchestrator (confd, user config, orchestration db)	/data/orchestration/
Bulkstats	/stats/bulkstats/
Node Exporter Textfile Collector	/stats/node/
Prometheus Textfile Collector	/stats/prometheus/
Prometheus Hi-Res storage db	/stats/prometheus-hi-res/2.0/
Prometheus-hi-res-s10x metadata	/stats/prometheus-hi-res-s10x/data/
Prometheus Planning storage db	/stats/prometheus-planning/2.0/
Prometheus-planning-s10x metadata	/stats/prometheus-planning-s10x/data/
Prometheus Trending storage db	/stats/Prometheus-trending/2.0/
prometheus-trending-s10x metadata	/stats/Prometheus-trending-s10x/data/

Upgrade MongoDB Version 3.6

Feature Summary and Revision History

Table 64: Summary Data

Applicable Product(s) or Functional Area	vDRA
Applicable Platform(s)	Not Applicable
Default Setting	Enabled - Always-on
Related Changes in This Release	Not Applicable
Related Documentation	Not Applicable

Table 65: Revision History

Revision Details	Release
First introduced:	19.4.0

Feature Description

vDRA supports MongoDB version 3.6.9. Currently, vDRA supports only fresh installation of MongoDB v3.6.9.

vDRA Resiliency Improvements and Database Overload Protection

Feature Summary and Revision History

Table 66: Summary Data

Applicable Product(s) or Functional Area	vDRA
Applicable Platform(s)	Not Applicable
Default Setting	Disabled - Configuration Required
Related Changes in This Release	Not Applicable
Related Documentation	CPS vDRA Configuration Guide

Table 67: Revision History

Revision Details	Release
First introduced:	19.4.0

Feature Description

In this release, vDRA worker node supports processing of all AARs when there is a high load. There is no AAR (contains Framed-IPv6/IPv4) message drop during high load. This feature can be enabled or disabled using **A A R Priority Processing** check box through Policy Builder. vDRA worker node must have separate thread pool to execute session and binding deletions. This thread pool is configurable under **Threading Configuration** in Policy Builder.

For more information, see *DRA Feature* section in the *CPS vDRA Configuration Guide*.

Zone-Aware Sharding

Feature Summary and Revision History

Table 68: Summary Data

Applicable Product(s) or Functional Area	vDRA
Applicable Platform(s)	Not Applicable

Default Setting	Enabled - Always-on
Related Changes in This Release	Not Applicable
Related Documentation	CPS vDRA Configuration Guide CPS vDRA Operations Guide

Table 69: Revision History

Revision Details	Release
First introduced:	19.4.0

Feature Description

vDRA now supports creating zones for IPv6 shards based on IPv6 pools. This is done so that the primary member of the replica set for an IPv6 address resides at the same physical location as the PGW assigning addresses from the IPv6 pool. This results in local writes (and reads) for the IPv6 binding database.

For more information, see the following:

- `database cluster` *`db-name`* `ipv6-zone-sharding` section in the *CPS vDRA Operations Guide*
- `database cluster` *`db-name`* `ipv6-zones-range` section in the *CPS vDRA Operations Guide*
- `database cluster` *`db-name`* `shard` section in the *CPS vDRA Operations Guide*
- *Configuring Zone Aware Sharding* section in the *CPS vDRA Configuration Guide*