



CPS vDRA Advanced Tuning Guide, Release 23.1.0

First Published: 2023-01-24

Last Modified: 2023-02-24

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 NON-INFRINGEMENT 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.

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. 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 standards documentation, or language that is used by a referenced third-party product.

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)

© 2023 Cisco Systems, Inc. All rights reserved.



CONTENTS

PREFACE

Preface v

About This Guide v

Audience v

Additional Support vi

Conventions (all documentation) vi

Communications, Services, and Additional Information vii

Important Notes viii

CHAPTER 1

vDRA Advanced Tuning 1

Threading Configuration 2

Database Capacity 4

Database Alert Expressions 4

Database Connection Settings 5

Audit Rate Limiter 7

Control Plane Tuning Configuration 7

IPC Queue Send Thread Tuning Configuration 8

IPC Queue Send Thread Priority Tuning Configuration 9



Preface

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

About This Guide



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

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

For information about available documentation, see the *CPS Documentation Map* for this release at [Cisco.com](https://www.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 .
<i>italic font</i>	Document titles, new or emphasized terms, and arguments for which you supply values are in <i>italic font</i> .
[]	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 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

vDRA Advanced Tuning

The configuration values in this document are based on the type and size of DRA deployment.

The deployment can be categorized into two sizes:

- Small Deployment - 2 Directors and 4 workers
- Large Deployment - 8 directors and 10 workers



Note Recommended values are based on:

- Local Latency:
 - 20 ms for small deployments
 - 25 ms on large deployments
- Remote Latency:
 - 55 ms for small deployments
 - 50 ms for large deployments

In case the latency changes to a higher value, recommended values for all the configurations highlighted in this document needs to be re-characterized again.

- [Threading Configuration, on page 2](#)
- [Database Capacity, on page 4](#)
- [Database Alert Expressions, on page 4](#)
- [Database Connection Settings, on page 5](#)
- [Audit Rate Limiter, on page 7](#)
- [Control Plane Tuning Configuration, on page 7](#)
- [IPC Queue Send Thread Tuning Configuration, on page 8](#)
- [IPC Queue Send Thread Priority Tuning Configuration, on page 9](#)

Threading Configuration

Thread pool configurations are applicable to java processes that run-in diameter-endpoint-s1xx on Director VM and binding-s1xx container on Worker VMs in DRA.

For more information on Threading Configuration, refer to the *CPS vDRA Configuration Guide*.

Table 1: Thread Pool Names

Thread Pool Name	Description
bindings	This thread pool is used for IMSI_APN/MSISDN_APN AAR lookup (MOG) and binding audit updates.
bindings.lookup	This thread pool is used only if appsharding.reader thread pool is overloaded. This thread pool is applicable only for AAR VoLTE calls.
bindings.delete	This thread pool is used to query session database and enqueue session and all the binding for deletion.
bindings.response	This thread pool is used to receive database storage response and forwards it as bind store response. This is applicable for CCR-I requests.
msgtimeouts	This thread pool is used to process timed out requests.
qprocessor	This thread pool is used to send routed messages outside on the specific TCP connection.
receivers	Not Applicable
senders	Not Applicable
qprocessor.priority	This thread pool is used to send only priority routed messages outside on the specific TCP connection.
localcpPublishers	This thread pool is used to publish messages to local control plane.
globalcpPublishers	This thread pool is used to publish messages to global control plane.
cpSubscriberWorker	This thread pool is used to process messages received on local and global control plane.
pcrfQuery	This thread pool is used for PCRF session queries used only for Rx AAR VoLTE calls.
pcrf.restapi	This thread pool is used to send request and receive response to/from pcrf endpoints and used only for RX AAR VoLTE calls. Note Scale By CPU Core is not applicable for this thread pool.
stackStopHandler	This thread pool is used to gracefully disconnect diameter peers on process shutdown.

Thread Pool Name	Description
appsharding.reader	This thread pool is specific to Rx AAR VoLTE IPv6 binding lookup and processing. This thread pool is created for each shard.

The following tables lists recommended values when single DRA installation TPS is above 200 K.



Important The highlighted values must be explicitly configured.

Table 2: Recommended Values

Thread Pool Name	Default Values		Recommended Values	
	Threads	Queue Size	Threads	Queue Size
bindings	30	1000	130	30000
bindings. lookup	10	1000	30	5000
bindings. delete	10	1000	100	50000
bindings. response	10	1000	10	1000
msgtimeouts	3	1000	3	1000
qprocessor	30	5000	150	25000
localcpPublishers	15	2000	15	2000
globalcpPublishers	40	3000	40	3000
cpSubscriberWorker	30	10000	30	10000
pcrfQuery	20	400	20	400
pcrf.restapi	10	2000	10	2000
stackStopHandler	50	0	50	0
appsharding. reader	1	1000	2	2000

Database Capacity



Note Rated database capacity has been arrived at after performing capacity tests.

The following table lists the database capacity for small and large deployments.

Table 3: Database Capacity

Deployment	Cluster	No. of Primary Shards (dual DB enabled)	Network Latency (in ms)	Database Capacity (in k)	Database Capacity Per Shard	85% Threshold Value
Small Deployment	Session_IPv6	8	20	30	$30000/8 = 3750$	3200
	IMSI_MSISDN	4	55	15	$15000/4 = 3750$	3200
Large Deployment	Session_IPv6	48	25	180	$180000/48 = 3750$	3200
	IMSI_MSISDN	48	50	145	$145000/48 = 3020$	2600

The following is the test call model used to determine the capacity for each cluster:

1. CCR_I, CCR_T and AAR call model for Session-IPv6 database cluster capacity.
2. CCR-I and CCR_T call model for IMSI-APN/MSISDN-APN database cluster capacity.
3. Database capacity is based on network latency mentioned in the [Table 3: Database Capacity, on page 4](#). If network latency is changed to higher value, then database capacity needs to be re-characterized.

Database Alert Expressions

Database alerts refer to alerts that can be configured to be triggered if the capacity of the database exceeds beyond a certain threshold.

IMSI_MSISDN Cluster

alert rule DRA_IMSI_MSISDN_DB_TPS_EXCEEDED

expression

```
"sum(rate(mongo_operation_total {state='primary',type='mongo',op=~'update|query|delete',cluster='IMSI_MSISDN'}[5m])) > (2500 * sum (mongo_node_state_primary {cluster='IMSI_MSISDN',type='mongo'}))"
```

event-host-label instance

message "{{ \$labels.instance }}" Persistence DB TPS exceeded , current value is {{ \$value }} !"

snmp-severity critical

snmp-clear-message "{{ \$labels.instance }}" Persistence DB TPS in control, current value is {{ \$value }} !"

Session_IPv6 Cluster

alert rule DRA_SESS_IPV6_DB_TPS_EXCEEDED

expression

```
"sum(rate(mongo_operation_total{state='primary',type='mongo',op=~'update|query|delete',cluster=~'SES_IPV6_*'}[5m]))
> (3200 * sum(mongo_node_state_primary{cluster=~'SES_IPV6_*',type='mongo'}))"
```

event-host-label instance

message "{{ \$labels.instance }}" Persistence DB TPS exceeded , current value is {{ \$value }} !"

snmp-severity critical

snmp-clear-message "{{ \$labels.instance }}" Persistence DB TPS in control, current value is {{ \$value }} !"

Database Connection Settings

The following configurations vary based on deployment size (Small or Large) and should be configured accordingly. The worker VM java processes connect to the databases and create two database connection pools based on the database operation:

- Create/Update, and Delete
- Read

For more information on configuration syntax and examples, refer to the *CPS vDRA Operations Guide*.

binding db-connection-settings

The recommended values listed in the table are for write operations (Create/Update, Delete):



Important The highlighted values must be explicitly configured.

Table 4: Recommended Value for Write Operations

Deployment	Bindings	connection-per-host	
		Default Values	Recommended Values
Small Deployment	drasession	10	14
	ipv6	10	10
	imsiapn	10	35
	msisdnapn	10	10
	ipv4	10	10

Deployment	Bindings	connection-per-host	
		Default Values	Recommended Values
Large Deployment	drasession	10	10
	ipv6	10	10
	imsiapn	10	20
	msisdnapi	10	20
	ipv4	10	2

binding db-read-connection-settings

The recommended values listed in the table are for read operations:



Important The highlighted values must be explicitly configured.

Table 5: Recommended Value for Read Operations

Deployment	Bindings	connections-per-host-for-read	
		Default Values	Recommended Values
Small Deployment	drasession	5	5
	ipv6	5	5
	imsiapn	5	3
	msisdnapi	5	5
	ipv4	5	5
Large Deployment	drasession	5	20
	ipv6	5	20
	imsiapn	5	10
	msisdnapi	5	10
	ipv4	5	1

binding cluster-binding-dbs imsiapn-msisdnapn

The following configuration is applicable only for large deployments. With this configuration IMSI-APN and MSISDN-APN bindings databases use the same connection pool for database transactions. This configuration helps in reducing the overall threads count in the worker VM.

Table 6: Recommended Values

Deployment	Recommended Value
Small Deployment	Disable
Large Deployment	Enable

Audit Rate Limiter

Database Audit is an important functionality in DRA. To ensure that the database does not fill up with unnecessary entries or expired entries are cleared, a house keeping process is always functional. The following parameters should be used based on the deployment size and rated capacity of the system.

Table 7: Audit Rate Limiter Recommended Values

Deployment	Rate Limiter	Stale Session Expiry Count	Binding DB Read Preference
Small Deployment	100	6	Nearest
Large Deployment	51	6	Nearest

Control Plane Tuning Configuration

Director nodes periodically advertise the status of all its peer connections over local and global control plane. All DRA nodes in the network use these peer status messages to keep their peer topology view updated. Peer topology is used to route the messages to appropriate director across different sites. When there are large number of peers connected across different sites, the load on control plane increases. This load increase can cause delay in processing peer status updates.

The following configurations are recommended to handle the increase in the control plane traffic process.

Table 8: Control Plane Traffic Tuning Commands

Command	Parameters	Default Values	Recommended Value
control-plane timers peer-status-update-interval <time-in-ms> peer-expiration-duration <duration-in-ms> Note To reflect the peer expiration duration change, application should be restarted in both director and worker nodes.	<i>time-in-ms</i> , <i>duration-in-ms</i>	2000, 10000	4000, 12000
control-plane remote-peer-policy mated-system id <system-id>	<i>system-id</i> (mated pair system ID)	NA	Mated Pair System ID
control-plane remote-peer-policy global accept diameter-applications [Application Type]	<i>Application Type</i>	All	Rx
control-plane ipc-endpoint update-interval <time-in-milliseconds>	< <i>time-in-milliseconds</i> >	100	1000

IPC Queue Send Thread Tuning Configuration

DRA maintains an IPC message queue where all the internal messages exchanged between Directors or Workers are stored. There are a set of IPC Queue processor threads that process these messages. During slow network conditions, when messages get delivered to one or a subset of peers, an IPC Queue processor thread that sends the message to these peers can eventually degrade the rate at which the IPC messages drop. These slow down conditions can cause message drops and 3002 errors not only to peers having degraded network but also to other peers.

The following configurations are recommended to handle the slow network peers.

Table 9: IPC Queue - send thread Tuning Commands

Deployment	Command	Parameters	Default Values	Recommended Value
Large	dra ipc-send-thread-limit <thread-limit> lock-sla-timeout <time-in-ms> message-throttle-duration <duration-in-ms> timeout-sample-to-throttle <max-samples>	<i>thread-limit</i> <i>time-in-ms</i> <i>duration-in-ms</i> <i>ms</i> <i>max-samples</i>	50 200 30000 150	50 (1/3rd of IPC threads) 250 30000 150

Deployment	Command	Parameters	Default Values	Recommended Value
Small	dra ipc-send-thread-limit <thread-limit> lock-sla-timeout <time-in-ms>message-throttle-duration <duration-in-ms> timeout-sample-to-throttle <max-samples>	<i>thread-limit</i> <i>time-in-ms</i> <i>duration-in-ms</i> <i>max-samples</i>	50 200 30000 150	10 (1/3rd of IPC threads) 50 30000 30

**Note**

- To disable the SLA timeout, configure **lock-sla-timeout** to 2000 ms. Disabling SLA timeout leads to ipc message drop in slow network.
- To disable throttling, configure **message-throttle-duration** to 0.
To disable thread throttling(limit), configure it to the value same as IPC Processor threads. Increasing **timeout-sample-to-throttle** reduces the chance of throttling the peer.

IPC Queue Send Thread Priority Tuning Configuration

The enhancement of the slow peer handling logic is applicable for priority and nonpriority messages. Once a slow peer is identified, the application must be notified. The further messages towards that peer will be throttled for the configured duration (*message-throttle-duration*). Once the configured time elapsed, the peer will be marked as normal peer and all the messages towards that peer will be processed.

The following configurations are recommended to handle the slow network peers with extra optional parameters.

Table 10: IPC Queue - send thread priority Tuning Commands

Deployment	Command	Parameters	Default Values	Recommended Value
Large	dra ipc-send-thread-priority limit <thread-limit> priority lock-sla-timeout <time-in-ms>priority message-throttle-duration <duration-in-ms>priority timeout-sample-to-throttle <max-samples>	<i>thread-limit</i> <i>time-in-ms</i> <i>duration-in-ms</i> <i>max-samples</i>	5 200 30000 150	5 (1/2 of IPC threads) 250 30000 150
Small	dra ipc-send-thread-priority limit <thread-limit> priority lock-sla-timeout <time-in-ms> priority message-throttle-duration <duration-in-ms> priority timeout-sample-to-throttle <max-samples>	<i>thread-limit</i> <i>time-in-ms</i> <i>duration-in-ms</i> <i>max-samples</i>	2 200 30000 150	2 (1/3rd of IPC threads) 50 30000 30

