Cisco Policy Suite 22.2.0 Release Notes

First Published: August 25, 2022

Last Updated: January 25, 2024

Introduction

This Release Note identifies installation notes, limitations, and restrictions, and open and resolved CDETS in Cisco Policy Suite (CPS) software version 22.2.0. Use this Release Note in combination with the documentation listed in the *Related Documentation* section.

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.

This Release Note includes the following sections:

New and Changed Feature Information

For information about a complete list of features and behavior changes associated with this release, see the *CPS Release Change Reference*.

Installation Notes

Download ISO Image

Download the 22.2.0 software package (ISO image) from:

https://software.cisco.com/download/home/284883882/type/284979976/release/22.2.0

Md5sum Details

PCRF

233c414331b48e14137e7e7c08c0f7c1	CPS_22.2.0_Base.release.qcow2_signed.tar.gz
fd73360c8bad97a3b4a755755f893944	CPS_22.2.0_Base.release.vmdk_signed.tar.gz
39e2602557767e9d92fec9a27e10beed	CPS_22.2.0.release.iso_signed.tar.gz

DRA

1de0ec2789415cce5c15de7e0bf460cc	CPS_Microservices_DRA_22.2.0_Base.release.vmdk_signed.tar.gz
98b2b8315426556359d10acd34f420ca	CPS_Microservices_DRA_22.2.0_Deployer.release.vmdk_signed.tar.gz
3e4291a17d1e803df7a9a3cf2904785e	CPS_Microservices_DRA_22.2.0.release.iso_signed.tar.gz
94bc0e0762393c39232e8bbb6674f3aa	CPS_Microservices_DRA_Binding_22.2.0.release.iso_signed.tar.gz

Component Versions

The following table lists the component version details for this release.

Table 1 - Component Versions

Component	Version
API Router	22.2.0.release
Audit	22.2.0.release
Balance	22.2.0.release
Cisco API	22.2.0.release
Cisco CPAR	22.2.0.release
Congestion Reference Data	22.2.0.release
Control Center	22.2.0.release
Core	22.2.0.release
CSB	22.2.0.release
Custom Reference Data	22.2.0.release
DHCP	22.2.0.release
Diameter2	22.2.0.release
DRA	22.2.0.release
Fault Management	22.2.0.release
IPAM	22.2.0.release
ISG Prepaid	22.2.0.release
LDAP	22.2.0.release
LDAP Server	22.2.0.release
LWR	22.2.0.release
Microservices Enablement	22.2.0.release
Notification	22.2.0.release
Policy Intel	22.2.0.release
POP-3 Authentication	22.2.0.release
Recharge Wallet	22.2.0.release
SCE	22.2.0.release
Scheduled Events	22.2.0.release

Component	Version
SPR	22.2.0.release
UDC	22.2.0.release
UDSN Interface	22.2.0.release
Unified API	22.2.0.release

Additional security has been added in CPS to verify the downloaded images.

Image Signing

Image signing allows for the following:

- Authenticity and Integrity: Image or software has not been modified and originated from a trusted source.
- Content Assurance: Image or software contains code from a trusted source, like Cisco.

Software Integrity Verification

To verify the integrity of the software image you have from Cisco, you can validate the md5sum checksum information against the checksum identified by Cisco for the software.

Image checksum information is available through **cisco.com Software Download Details**. To find the checksum, hover the mouse pointer over the software image on cisco.com.

If md5sum is correct, run *tar -zxvf* command to extract the downloaded file.

The files are extracted to a new directory with the same name as the downloaded file name without extension (.tar.gz).

The extracted directory contains the certificate files (.cer), python file (cisco_x509_verify_release.py), digital certificate file (.der), readme files (*.README), signature files (.signature) and installation files (.iso .vmdk, .qcow2 and .tar.gz).

Certificate Validation

To verify whether the installation files are released by Cisco System Pvt. Ltd and are not tampered/modified or infected by virus, malware, spyware, or ransomware, follow the instruction given in corresponding *.README file.

NOTE: Every installation file has its own signature and README file. Before following the instructions in the README file, make sure that cisco.com is accessible from verification server/host/machine/computer. In every README file, a Python command is provided which when executed connects you to cisco.com to verify that all the installation files are released by cisco.com or not. Python 2.7.4 and OpenSSL is required to execute cisco_x509_verify_release.py script.

New Installations

- Pre-requisites
- VMware Environment
- OpenStack Environment

Pre-requisites

- Ensure to use the grafana_update_query.sh script to update the Grafana queries. As part of Python3 upgrade activity, graphite package is upgradedr to add an escape char ('\') before pipe ('|') for the graphs to get displayed. This script replaces the "|" with "\|" to make the query work. The default credentials will be taken as admin:admin. For more information, see the *Migrate the Cluster Manager VM* section in the *CPS Migration and Upgrade Guide*.

- **Migration and Rollback**: At the time of migrating a Cluster Manager VM, for ISSM migration, you need to update the Grafana queries. Run the following command:

  /var/qps/bin/support/grafana_update_query.sh

  /var/qps/bin/support/grafana_update_query.sh [-h]

  var/qps/bin/support/grafana_update_query.sh [-c]

This script execution is mandatory if you are performing iSSM from CPS 22.1 on previous versions. For more information, see the Migrate the Cluster Manager VM section in the CPS Migration and Upgrade Guide and also grfana_update_query.sh CPS Commands section in the CPS Operations Guide.

During Rollback, If the restore step is already performed, it is not necessary to revert the grafana queries as the updated grafana graphs will work in the previous releases also.

Migration and Rollback Requirements:

For *migrate.sh backup cluman* step, you need to mount the existing 22.1 ISO. This is because the system does not support Python 3 ISO due to compatibility issues.

For rollback, mount the 22.1 ISO on old cluster manager and then run enable set 1. Since the old cluster manager will be running on Python 2 and mounting 22.2 ISO for rollback would cause compatibility issues.

VMware Environment

To perform a new installation of CPS 22.2.0 in a VMware environment, see the *CPS Installation Guide for VMware*.

NOTE: After installation is complete, you need to configure at least one Graphite/Grafana user. Grafana supports Graphite data source credential configuration capability. Graphite data source requires common data source credential to be configured using Grafana for Grafana user. Data source credential must be configured after fresh installation. If you fail to add the user, then Grafana will not have access to Graphite database, and you will get continuous prompts for Graphite/Grafana credentials.

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

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

OpenStack Environment

To perform a new installation of CPS 22.2.0 in an OpenStack environment, see the *CPS Installation Guide for OpenStack*.

NOTE: After installation is complete, you need to configure at least one Graphite/Grafana user. Grafana supports Graphite data source credential configuration capability. Graphite data source requires common data source credential to be configured using Grafana for Grafana user. Data source credential must be configured after fresh installation. If you fail to add the user, then Grafana will not have access to Graphite database, and you will get continuous prompts for Graphite/Grafana credentials.

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

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

Support to Replace CentOS with Alma Linux on CPS

CentOS version 8.1 is replaced with Alma Linux 8.5 with latest rpm packages. With Alma Linux 8.5, the kernel version is modified to: # rpm -qa | grep kernel-[0-9] kernel-4.18.0-348.7.1.el8_5.x86_64 ## cat /etc/redhat-release AlmaLinux release 8.5

(Arctic Sphynx) # uname -a Along with the OS upgrade and Kernel upgrade many of the dependent third-party packages are also upgraded.

Support for Python 3.9.6 Version

The python is upgraded to the latest 3.9.6 stable version.

Support for MongoDB 4.2

In CPS 22.2.0, MongoDB version is upgraded to 4.2.20.

Migrate an Existing CPS Installation

To migrate an existing CPS installation, see the *CPS Migration and Upgrade Guide*. CPS migration is supported from CPS 22.1.0/CPS 22.1.1 to CPS 22.2.0.

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

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

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

NOTE: As CPS 22.2.0 supports ESXi 6.7/7.0 (version until 7.0.2), make sure OVF tool version 4.3.0 is installed in CPS 22.2.0 from where you are migrating.

Version 4.3.0 for VMware 6.7/7.0: VMware-ovftool-4.3.0-13981069-lin.x86_64.bundle

You can download the OVF tool version 4.3.0 from https://code.vmware.com/web/tool/4.3.0/ovf.

For more information, consult your Cisco Technical Representative.

Upgrade an Existing CPS Installation

In-Service Software Upgrade (ISSU) is not supported when migrating from CPS 21.1.0/CPS 21.2.0 to CPS22.2.0.

Post Migration/Upgrade Steps

Re-Apply Configuration Changes

After the migration/upgrade is complete, compare your modified configuration files that you backed up earlier with the newly installed versions. Re-apply any modifications to the configuration files.

Verify Configuration Settings

After the migration/upgrade is finished, verify the following configuration settings.

NOTE: Use the default values listed below unless otherwise instructed by your Cisco Account representative.

NOTE: During the migration/upgrade process, these configuration files are not overwritten. Only during a new install will these settings be applied.

- /etc/broadhop/qns.conf
  - -Dmongo.client.thread.maxWaitTime.balance=1200
  - -Dmongo.connections.per.host.balance=10
  - -Dmongo.threads.allowed.to.wait.for.connection.balance=10
  - -Dmongo.client.thread.maxWaitTime=1200
  - -Dmongo.connections.per.host=5
  - -Dmongo.threads.allowed.to.wait.for.connection=10
  - -Dcom.mongodb.updaterIntervalMS=400
  - -Dcom.mongodb.updaterConnectTimeoutMS=600
  - -Dcom.mongodb.updaterSocketTimeoutMS=600
  - -DdbSocketTimeout.balance=1000
  - -DdbSocketTimeout=1000
  - -DdbConnectTimeout.balance=1200
  - -DdbConnectTimeout=1200
  - -Dcontrolcenter.disableAndsf=true
  - -DnodeHeartBeatInterval=9000
  - -DdbConnectTimeout.balance=1200
  - -Dstatistics.step.interval=1
  - -DshardPingLoopLength=3
  - -DshardPingCycle=200
  - -DshardPingerTimeoutMs=75
  - -Ddiameter.default.timeout.ms=2000
  - -DmaxLockAttempts=3
  - -DretryMs=3
  - -DmessageSlaMs=1500
  - -DmemcacheClientTimeout=200
  - -Dlocking.disable=true

NOTE: The following setting should be present only for GR (multi-cluster) CPS deployments:

```
-DclusterFailureDetectionMS=1000
```

NOTE: In an HA or GR deployment with local chassis redundancy, the following setting should be set to true. By default, it is set to false.

```
-Dremote.locking.off
```

- /etc/broadhop/diameter_endpoint/qns.conf
  - -Dzmq.send.hwm=1000
  - -Dzmq.recv.hwm=1000

Reconfigure Service Option

After upgrading from previous release to the current CPS release, Service option configured with Subscriber-Id becomes invalid and you need to reconfigure multiple Subscriber Id in SpendingLimitReport under Service Configurations.

Verify logback.xml Configuration

Make sure the following line exists in the logback.xml file being used. If not, then add the line:

<property scope="context" name="HOSTNAME" value="${HOSTNAME}" />

To ensure logback.xml file changes are reflected at runtime, the scanPeriod must be explicitly specified:

<configuration scan="true" scanPeriod="1 minute">

NOTE: In case scanPeriod is missing from already deployed logback.xml file, the application needs to be restarted for the updated scanPeriod configuration to be applicable.

After completing the updates in logback.xml, execute the following command to copy the file to all the VMs:

SSHUSER_PREFERROOT=true copytoall.sh /etc/broadhop/logback.xml /etc/broadhop/logback.xml

Change Mongo Storage Engine from MMapV1 to WiredTiger in CPS Product

Starting from CPS 22.1.1 release, MongoDB Storage Engine is changed from MMAPv1 to WiredTiger.

WiredTiger storage engine change in MongoDB Server requires additional CPU resources of ~15% and additional memory (RAM) resources of ~40% in the Session Manager VMs. WiredTiger consumes up to ~40% extra memory from total memory(RAM) than MMapV1.

For example, If the sessionmgr VM (150GB) with MMapV1 uses 60GB, then WiredTiger requires 120GB(MMapV1 usage 60GB + 40% of total memory).

As per mongo documentation, the wiredtigercachegb can be configured as [50% of (RAM - 1 GB)] in the VM.

If “n” mongo processes are running in the VM, the wiredtigercachegb can be configured as [50% of (RAM - 1 GB)]/n per mongo process.

For example, in the setup:

- Sessionmgr VMs configured RAM: 157GB
- The number of mongo processes will be running on VM: 6
- Each process cache size can be configured: [50% of (157GB-1GB)]/6 ==> 78/6 = 13GB(can rounded to 12 GB)

NOTE: OS can consume 40-50GB of buffer/cache memory towards system/kernel operations.

The following values must be configured in mongoConfig.cfg:

- WT_CACHESIZEGB=12
- WT_CACHEARBSIZEGB=1

Additional Notes

This section provides additional notes necessary for proper installation/working of CPS.

- Session Manager Configuration: After a new deployment, session managers are not automatically configured.

  a. Edit the */etc/broadhop/mongoConfig.cfg* file to ensure all the data paths are set to /var/data and not /data.

  b. Then execute the following command from pcrfclient01 to configure all the replication sets:

  /var/qps/bin/support/mongo/build_set.sh --all --create

- Default gateway in lb01/lb02: After the installation, the default gateway might not be set to the management LAN. If this is the case, change the default gateway to the management LAN gateway
- By default, pending transaction feature is enabled. If you are not using it, Cisco recommends disabling pending transaction feature post deployment.

  To disable pending transaction, the following parameter can be configured in */etc/broadhop/qns.conf* file:

  com.broadhop.diameter.gx.pending_txn.attempts=0

  After adding the parameter in qns.conf file, restart all VMs using *stopall.sh*/*startall.sh* or *restartall.sh* command.

- Add support to disable syncing carbon database and bulk stats files (ISSM)

  Add the following flags in */var/install.cfg* file:

  SKIP_BLKSTATS

  SKIP_CARBONDB

  Example to disable synching:

  SKIP_BLKSTATS=1

  SKIP_CARBONDB=1

- Add the following parameters in */var/install.cfg* file to skip installation type selection and initialization steps during ISSU/ISSM:

  INSTALL_TYPE

  INITIALIZE_ENVIRONMENT

  Example:

  INSTALL_TYPE=mobile

  INITIALIZE_ENVIRONMENT=yes

- Inconsistency in DPR sent by CPS on executing `monit stop` command

  Issue: When `monit stop all` is executed on Policy Director (LB) VMs with active VIP, DPR is not sent to all the diameter peers.

  Conditions: `monit stop all` executed on Policy Director (LB) VMs with active VIP

  Cause: DPR is sent to all the connected diameter peers. However, since `monit stop all` is executed , all the processes on the Policy Director (LB) go down including corosync/haproxy. As a result, some of the DPR messages go out and some are not delivered based on the order of the services going down.

  Workaround: Instead of `monit stop all`, you can stop all the qns process on Policy Director (LB) VMs by executing `monit stop qns-2/3/4` and then issue a `monit stop all` comand.

With this workaround, processes such as, haproxy/coronsync are up when DPR messages are generated, CPS makes sure that all DPR messages generated by the Policy Directors are delivered.

CSCvr34614: Prometheus Containers stuck in started state after recovering from site failover

Prometheus is the third-party code, used in DRA and Binding VNFs.

For more information related to the issue, see https://github.com/prometheus/prometheus/issues/4058

Issue: Prometheus database blocks contain corrupted data and does not have *meta.json* file to initialize the database when Prometheus comes up.

Solution: Prometheus doesn't have enough capability to repair the corrupted database blocks. Currently, the solution is to manually delete the corrupted block and start the Prometheus process manually.

NOTE: If the Prometheus containers having issue are from Master VM, then some data will not be available and Grafana displays some gap in the data. It is expected behavior as corrupted folders have been deleted. One can access the missing data by adding the data source with another Prometheus container present on control-0 and control-1 VMs (HA for master Prometheus).

The following steps must be performed to delete the corrupted block and start the Prometheus process manually:

NOTE: If there are more than one failed Prometheus containers, the steps need to be repeated for each corrupted block.

1. Connect to the container which has failed to come up.

   docker connect prometheus-hi-res-s101

2. From container, check whether Prometheus process is in FATAL state or not.

   supervisorctl status prometheus

3. If the process is in "FATAL" state, remove the data folder from container.

   *rm -rf /data-2/**

   NOTE: The command deletes the data folder. As Prometheus data is available between master/control-0/control-1 VMs, data can be restored.

4. Inside container, start the Prometheus process again.

   supervisorctl start prometheus

5. From inside container, check again whether Prometheus process is in RUNNING state or not.

   supervisorctl status Prometheus

CSCvr21943: After site resiliency the consul gets struck in STARTED state

Issue: Consul containers remain in STARTED state when a site failure scenario is executed. After the failure scenario is executed, the system does not come up again in the expected state.

Condition: After multiple VM (or) site power off/on cycle, consul containers are stuck in STARTED/STARTING (non-HEALTHY) state.

admin@orchestrator[an-master]# *show scheduling status | tab | include consul*

consul 1 50 infrastructure SCHEDULING false

admin@orchestrator[an-master]# *show docker service | tab | include consul*

consul 1 consul-1 19.4.5-2019-10-01.8115.4fb2b4a an-master consul-1 STARTED true Pending health check

consul 1 consul-2 19.4.5-2019-10-01.8115.4fb2b4a an-control-0 consul-2 STARTED true Pending health check

consul 1 consul-3 19.4.5-2019-10-01.8115.4fb2b4a an-control-1 consul-3 STARTED true Pending health check

Solution:

- Prepare **peers.json** file: Connect to the consul-1 container.

  root@consul-1:/# consul info

  Get the "latest_configuration" value under **raft**:

  Sample output of consul info:

  ....

  raft:

  ...

  last_snapshot_term = 1083

  *latest_configuration = [{Suffrage:**Voter ID**:bb7e19b5-e709-3c8c-686f-e839e941773f **Address**:10.42.0.1:8300} {Suffrage:**Voter ID**:66a6756f-49ac-b2a7-74c6-07922e8c2f81 **Address**:10.40.0.3:8300} {Suffrage:**Voter ID**:7b62389e-af67-d0f3-79d9-95bb356ea52c **Address**:10.47.128.3:8300} {Suffrage:**Voter ID**:b753a43f-4278-6f45-27f1-d2f88081b6d3 **Address**:10.38.0.30:8300} {Suffrage:**Voter ID**:ad423368-98bd-d87a-4d73-99520091321b **Address**:10.45.0.26:8300} {Suffrage:**Voter ID**:b916b8d1-b2dd-4799-db95-09a1e1144380 **Address**:10.37.0.11:8300} {Suffrage:**Voter ID**:543ba9f7-110a-7559-3607-ea6d5d1ef83b **Address**:10.37.192.2:8300}]*

  latest_configuration_index = 2503803

  num_peers = 6

  ...

- **latest_configuration:** This is a list of dictionaries. The number of dictionaries is equal to the **num_peers** field. Each dictionary has 2 keys, which are **Voter ID** and **Address**.

  In the sample output above, the number of dictionaries is 7 (num_peers + self) corresponding to num_peers=6.

  Each dictionary represents the **Voter ID** and **Address** corresponding to each Consul Node (consul-1, consul-2, consul-3, and so on) not in any particular order.

  So, fetch the **Voter ID**/**Address** corresponding to consul-1, consul-2 and consul-3 from the latest_configuration as mentioned below.

  root@consul-1:/# ifconfig

  Get the inet addr: value (IP adress) corresponding to ethwe: interface.

  Compare this IP address from ifconfig command against the **Address** field in **latest_configuration**. Make a note of the corresponding **Voter ID** field of the matching **Address** field.

  Identify the values of **Voter ID** and **Address** fields corresponding to consul-1 that need to be populated into peers.json file

  NOTE: Mapping between latest_configuration and peers.json.

  Table 2 - Mapping Table

latest_configuration	peers.json
Address (should be same as IP address got from Consul container's ifconfig command)	address
Voter ID	id

  Similarly, connect to consul-2 and consul-3 containers and get the **Voter ID** for the matching **Address**.

  Identify the details of **Address** and **Voter ID** corresponding to consul-2 and consul-3 containers, they must be populated into peers.json file.

  Now peers.json file should be populated with details corresponding to consul-1, consul-2 and consul-3 containers as identified above.

- Create peers.json file on Master VM.

  NOTE: The sample peers.json file should not be used. The file is for reference purposes only. Add "id" and "address" fields based on your deployment.

Sample peers.json

```
----------.
[
 {
  "id": "bb7e19b5-e709-3c8c-686f-e839e941773f",
  "address": "10.42.0.1:8300",
  "non_voter": false
 },
 {
  "id": "66a6756f-49ac-b2a7-74c6-07922e8c2f81",
  "address": "10.40.0.3:8300",
  "non_voter": false
 },
 {
  "id": "7b62389e-af67-d0f3-79d9-95bb356ea52c",
  "address": "10.47.128.3:8300",
  "non_voter": false
 }
]
```

- Restart the service after copying peers.json file:

  peers.json is created on the Master VM.

Copy peers.json file from Master VM to the Control VM's.

- Stop the services:

  Stop all the services on all the consul containers of Master and Control VM's.

  From Orchestrator CLI:

  admin@orchestrator[an-master]# docker connect consul-1

  root@consul-1:/# supervisorctl stop all

  admin@orchestrator[an-master]# docker connect consul-2

  root@consul-2:/# supervisorctl stop all

  admin@orchestrator[an-master]# docker connect consul-3

  root@consul-3:/# supervisorctl stop all

- Copy peers.json file:

  On Master VM, copy peers.json file onto "/data/raft" of the consul-1 container.

  sudo cp peers.json /data/consul-1/data/raft/

  On Control-0 VM, copy peers.json file onto "/data/raft" of the consul-2 container.

  sudo cp peers.json /data/consul-2/data/raft/

  On Control-1 VM, copy peers.json file onto "/data/raft" of the consu-3 container.

  sudo cp peers.json /data/consul-3/data/raft/

- Start the services:

  Start all the services on all the consul containers of Master and Control VM's.

  From Orchestrator CLI:

  admin@orchestrator[an-master]# docker connect consul-1

  root@consul-1:/# supervisorctl start all

  admin@orchestrator[an-master]# docker connect consul-2

  root@consul-2:/# supervisorctl start all

  admin@orchestrator[an-master]# docker connect consul-3

  root@consul-3:/# supervisorctl start all

All the consul containers will be restored to HEALTHY state.

admin@orchestrator[an-master]# show docker service | tab | include consul

consul 1 consul-1 19.4.5-2019-10-01.8115.4fb2b4a an-master consul-1 HEALTHY false -

consul 1 consul-2 19.4.5-2019-10-01.8115.4fb2b4a an-control-0 consul-2 HEALTHY false -

consul 1 consul-3 19.4.5-2019-10-01.8115.4fb2b4a an-control-1 consul-3 HEALTHY false -

admin@orchestrator[an-master]# show scheduling status | tab | include consul

consul 1 50 infrastructure RUNNING false

Limitations

This section lists the limitations of this release:

- Solicited Application Reporting

  The following are some restrictions on configuration for the new service options:

  - The pre-configured ADC rule generated by CRD lookup has ADC-Rule-Install AVP definition with support for only three AVPs ADC-Rule-Name, TDF-Application-Identifier, Mute-Notification.
  - For AVPs that are multi-valued, CRD tables are expected to have multiple records - each giving the same output.
  - Comma(,) is not a valid character to be used in values for referenced CRD column in SdToggleConfiguration.
  - AVP Table currently only supports OctetStringAvp value for AVP Data-type.
- During performance testing, it has been found that defining many QoS Group of Rule Definitions for a single session results in degraded CPU performance. Testing with 50 QoS Group of Rule Definitions resulted in a 2x increase in CPU consumption. The relationship appears to be a linear relationship to the number of defined QoS Group of Rule Definitions on a service.
- Hour Boundary Enhancement

  Change in cell congestion level when look-ahead rule is already installed:

  If a cell congestion value changes for current hour or any of the look-ahead hours, there will be no change in rule sent for the rules that are already installed.

  No applicability to QoS Rules:

  The look-ahead works for PCC rules only where we have rule activation/deactivation capabilities and can install upcoming changes in advance. However, if the RAN Congestion use case is changed to use the QoS-Info AVP instead of using PCC rules, we need to fall back to the current RAR on the hour boundary implementation for that use case since the standard do not let us install QoS-info changes ahead of time like we can with PCC rules.
- The Cluster Manager's internal (private) network IP address must be assigned to the host name "installer" in the `/etc/hosts` file. If not, backup/restore scripts (`env_import.sh`, `env_export.sh`) will have access issues to OAM (pcrfclient01/pcrfclient02) VMs.
- CSCva02957: Redis instances continue to run, even after Redis is disabled using the parameter `-DenableQueueSystem=false` in qns.conf (`/etc/broadhop/`) file and `/etc/broadhop/redisTopology.ini` file.
- CSCva16388: A split-brain scenario (that is, VIPs are up on both nodes) can still occur when there is connectivity loss between lb01 and lb02 and not with other hosts.

Open and Resolved CDETS

The following sections list open and resolved CDETS for this release. For your convenience in location CDETS in Cisco's Bug Toolkit, the caveat titles listed in this section are drawn directly from the Bug Toolkit database. These caveat titles are not intended to be read as complete sentences because the title field length is limited. In the caveat titles, some truncation of wording or punctuation might be necessary to provide the most complete and concise description.

NOTE: If you are a registered cisco.com user, view Bug Toolkit on cisco.com at the following website: https://tools.cisco.com/bugsearch

To become a registered cisco.com user, go to the following website: https://tools.cisco.com/RPF/register/register.do?exit_url=

Open CDETS

The following table lists the open CDETS in this release.

CPS Open CDETS

Table 3 - CPS Open CDETS

CDETS ID	Headline
CSCwc54243	Session manager VM reachability from QNS or UDC shard pinger timing out intermittently
CSCwc58310	httpd getting restarted at regular intervals in pcrfclient01 and pcrfclient02 VMs
CSCwc73519	F4226, PCRF is not storing the AVP in the session

vDRA Open CDETS

Table 4 - vDRA Open CDETS

CDETS ID	Headline
CSCwa98187	vPAS - 3002 Error/timeouts observed during longevity randomly
CSCwb96525	Site ID missing for KPI diameter_request_total with label late
CSCwc07906	Binding storage failed message, mongo exception in consolidated QNS logs
CSCvx14701	Gx / Rx Timeout dashboard shows incorrect message processing time

Resolved CDETS

This section lists the resolved/verified CDETS in this release.

CPS Resolved CDETS

Table 5 - CPS Resolved CDETS

CDETS ID	Headline
CSCwc47887	tpm2-tss package upgrade is missing in all the VMs from QPS repo
CSCwb89576	Security Vulnerabilities detected in Nessus Scan: -Apache Log4j 1.x
CSCwa90700	When https enabled, orchestration API is failing with Type Error

vDRA Resolved CDETS

Table 6 - vDRA Resolved CDETS

CDETS ID	Headline
CSCwb23885	Metadata DB health check fail during worker down scenario
CSCwc30675	EFK local forwarding causes memory to increase over time due to overflow of logs
CSCwb17404	peer routing table not taking latest values from CRD
CSCwc32334	DB record count KPIs reporting multiple values from different containers in Grafana
CSCwb02113	Disable Dynamic Rate limiting during run time not clearing old cache
CSCwb27614	Make the Prometheus data retention period configurable by user
CSCwb53438	Add GTAC access for CLI mode in control VMs
CSCwb22705	Some of the regex expression search not working in CRD GUI
CSCwa97125	database fcv check does not include fpas IMSI_MSISDN DB shard members
CSCwb09318	CPS vDRA DRD's SSL Certificate signed using a weak hashing algorithm SHA1
CSCwa01913	Message type label not updated for 3002 error message KPI in some scenarios
CSCwc61822	Install tcpwrapper utility on VM
CSCwc59390	Filter group record in Routing by Peer Group dashboard
CSCwb52054	SSH slowness due to more entries into audit logs

Related Documentation

This section contains information about the documentation available for Cisco Policy Suite.

Release-Specific Documents

Refer to the following documents for better understanding of Cisco Policy Suite.

- *CPS Advanced Tuning Guide*
- *CPS Backup and Restore Guide*
- *CPS CCI Guide for Full Privilege Administrators*
- *CPS CCI Guide for View Only Administrators*
- *CPS Central Administration Guide*
- *CPS Documentation Map*
- *CPS Geographic Redundancy Guide*
- *CPS Installation Guide - OpenStack*
- *CPS Installation Guide – VMware*
- *CPS Migration and Upgrade Guide*
- *CPS Mobile Configuration Guide*
- *CPS Operations Guide*
- *CPS Policy Reporting Guide*
- *CPS Release Change Reference*

- *CPS Release Notes*
- *CPS SNMP, Alarms, and Clearing Procedures Guide*
- *CPS Troubleshooting Guide*
- *CPS Unified API Reference Guide*
- *CPS vDRA Administration Guide*
- *CPS vDRA Advanced Tuning Guide*
- *CPS vDRA Configuration Guide*
- *CPS vDRA Installation Guide for VMware*
- *CPS vDRA Operations Guide*
- *CPS vDRA SNMP and Alarms Guide*
- *CPS vDRA Troubleshooting Guide*

These documents can be downloaded from https://www.cisco.com/c/en/us/support/wireless/policy-suite-mobile/products-installation-and-configuration-guides-list.html.

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, using the Cisco Bug Search Tool (BST), submitting a service request, and gathering additional information, see What's New in Cisco Product Documentation, at: http://www.cisco.com/c/en/us/td/docs/general/whatsnew/whatsnew.html.

Subscribe to What's New in Cisco Product Documentation, which lists all new and revised Cisco technical documentation, as an RSS feed and deliver content directly to your desktop using a reader application. The RSS feeds are a free service.