UPC CUPS Release Change Reference, Release 21.28

First Published: 2022-09-29

Last Modified: 2024-10-30

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

CONTENTS

About this Guide

Note Control and User Plane Separation (CUPS) represents a significant architectural change in the way StarOS-based products are deployed in the 3G, 4G, and 5G networks. This document provides information on the features and functionality specifically supported by this 3G/4G CUPS product deployed in a 3G/4G network. It should not be assumed that features and functionality that have been previously supported in legacy or non-CUPS products are supported by this product. References to any legacy or non-CUPS products or features are for informational purposes only. Furthermore, it should not be assumed that any constructs (including, but not limited to, commands, statistics, attributes, MIB objects, alarms, logs, services) referenced in this document imply functional parity with legacy or non-CUPS products. Please contact your Cisco Account or Support representative for any questions about parity between this product and any legacy or non-CUPS products.

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. 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 RCR describes new and modified feature and behavior change information for the 21.24.x CUPS releases.

• Conventions Used, on page xi

Conventions Used

The following tables describe the conventions used throughout this documentation.

Notice Type	Description
Information Note	Provides information about important features or instructions.
Caution	Alerts you of potential damage to a program, device, or system.

Notice Type	Description
Warning	Alerts you of potential personal injury or fatality. May also alert you of potential electrical hazards.

Typeface Conventions	Description
Text represented as a `screen display`	This typeface represents displays that appear on your terminal screen, for example: `Login:`
Text represented as **commands**	This typeface represents commands that you enter, for example: **show ip access-list** This document always gives the full form of a command in lowercase letters. Commands are not case sensitive.
Text represented as a **command** *variable*	This typeface represents a variable that is part of a command, for example: **show card** *slot_number* *slot_number* is a variable representing the desired chassis slot number.
Text represented as menu or sub-menu names	This typeface represents menus and sub-menus that you access within a software application, for example: Click the **File** menu, then click **New**

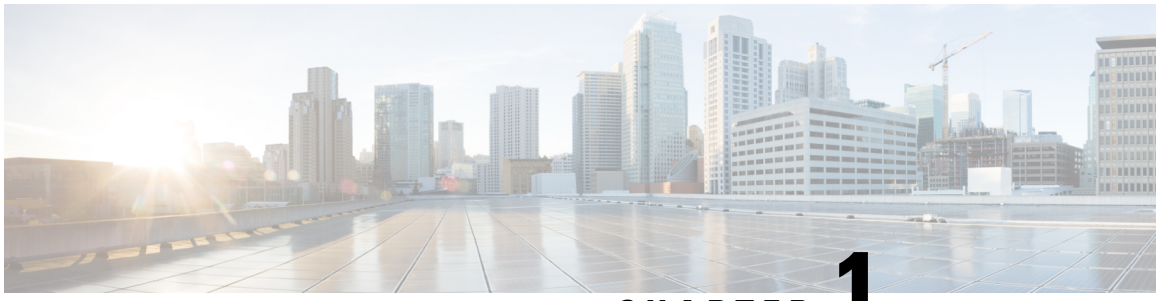

CHAPTER 1

UPC CUPS Release Change Reference

Accurate Traffic Throttling—CSCwc62127

Revision History

Revision Details	Release
First introduced.	21.28.m18

Behavior Change

Previous Behavior: The bit rates of dummy QoS Enforcement Rules (QERs associated with rulebase PDR having valid QoS Flow Identifier (QFI)) that were incorrectly applied for policing led to incorrect throttling.

New Behavior: The bit rates of dummy QERs will be ignored for policing.

Customer Impact: The customer will observe accurate throttling of traffic

Appending Original URL to Redirect URL

Revision History

Revision Details	Release
First introduced.	21.28.m10

Feature Changes

UPF supports dynamic Advice of Charge (AoC) redirections with URL provided by Online Charging System (OCS). This redirection is performed for a particular Service ID/Rating Group combination without affecting the flows mapped to other Service ID/Rating Group combinations.

For redirection to an AoC or top-up server, the UPF appends the original HTTP URL to the redirected session. To append the original URL for redirection, the OCS indicates to the CP and UP by specifying a special "?" character to the end of the AoC redirection. The redirect URL will be appended with the original URL information using the token name configured with the **diameter redirect-url-token** command under the Credit Control Configuration mode. The AoC server redirects the user to the original location on completion of AoC.

For more information, refer to the *FUI Redirection* chapter in the *UPC CUPS User Plane Administration Guide*.

Avoid Sx Flaps during Simultaneous User Plane Registrations—CSCwd39954

Revision History

Revision Details	Release
First introduced.	21.28.4

Feature Changes

In CUPS, Sx association flaps occur on the Control Plane (CP) due to delay in processing Sx queue and messenger queue at CP Sx demux. The flaps occur during simultaneous User Plane (UP) registrations with multiple egress contexts. CUPS overcomes this scenario by establishing Sx in groups after CP reboot.

Previous Behavior: During scaled setup, the Sx flap had simultaneous UP registrations.

New Behavior: CUPS supports UP registrations at the rate of 1000 per second with a pacing of 1 ms between subsequent UP registration requests.

Customer Impact: Sx flaps require reregistration of UP.

Important You must ensure to add the **update ip-pool apn all** CLI command in the CP configuration, and run it after load or reload. The **update ip-pool apn all** command is mandatory and must be part of any CP load or reload instances. Otherwise, IP chunking to UPs may not be successful.

Behavior of Debuffered TCP Packets—CSCwh37204

Revision History

Revision Details	Release
First introduced.	21.28.m14

Behavior Change

During a CUPS collapsed (Sxab) call, the downlink TCP packets are debuffered when UE moves from idle to active state. The downlink buffered TCP packets are sent after rule matching.

Previous Behavior: If the downlink buffered TCP packet was corrupt or invalid (payload length in TCP header lesser than actual payload size), L7 analysis was done when packet was debuffered.

New Behavior: If the downlink buffered TCP packet is corrupt or invalid, L7 analysis will not be done when packet is debuffered.

Behavior of Secondary RAT Usage Reports in CDR—CSCwd20301

Revision History

Revision Details	Release
First introduced.	21.28.3

Feature Changes

In the CUPS Secondary RAT Usage records scenario where the configured number of Secondary RAT Usage records are received at CP, the CP sends the query URR to UP to get the current usage report. The CP generates the PGW-CDR on receipt of the usage report.

Previous Behavior: During load conditions, if more Secondary RAT Usage Reports are received at CP before receiving the usage report from UP, the Secondary RAT Usage records reported in PGW-CDR could be more than the configured limit.

New Behavior: During load conditions, if more Secondary RAT Usage Reports are received at CP before receiving the usage report from UP, the last Secondary RAT Usage record merges with the subsequent records received until the usage report from UP. In normal scenarios, the records reported are within the configured limit. If there is a delay in the usage report due to load conditions, the last record could be bulky.

Customer Impact: There is no impact on charging. With the new behavior, the last Secondary RAT Usage record might be bulky instead of a granular report.

Boot State Assignment Trap—CSCwe40744

Revision History

Revision Details	Release
First introduced.	21.28.m14 21.28.m7

Behavior Change

The new *UPFStateAssigned* trap displays the boot state assignment.

Previous Behavior: RCM did not support the new trap.

New Behavior: RCM supports the new *UPFStateAssigned* trap. This trap displays the assigned state for a newly booted UPF registering with RCM (pending active/standby). This trap also displays the active/standby state of a UPF on controller restart or when RCM becomes HA active, and if a fully active/standby state UPF re-registers.

The existing *UPFBootComplete* trap displays the final active/standby state.

Both these traps display the UPF IP address. On comparing the timestamp of these two traps, the user can estimate the config push time for UPF.

Customer Impact: The traps display additional information about UPF.

Change in the RCM VM Kubernetes Container Runtime CLI-CSCwe08983

Revision History

Revision Details	Release
First introduced.	21.28.mh3

Behavior Change

The RCM VM kubernetes container runtime **docker** is removed and replaced with **containerd**.

Previous Behavior: The CLI **docker** was present in the RCM VM.

New Behavior: The **docker** CLI is replaced with **nerdctl** CLI. However, as a convenience, the term "docker" has been made an alias for **nerdctl --namespace k8s.io**.

Also, two new low-level commands **crictl** and **ctr** are introduced.

Customer Impact: There is no change in the behavior of **containerd**. However, any customized script using the **docker** CLI needs to be changed accordingly.

Clearing and Displaying IP Neighbor VPP using New CLIs - CSCwd92954

Revision History

Revision Details	Release
First introduced.	21.28.mh3

Behavior Change

Previous Behavior: There were no CLIs available for the users to clear or display the neighbor VPPs of any IPv4 or IPv6 address.

New Behavior: 4 new CLIs are introduced to allow the users to clear and display the neighbor VPPs of any IPv4 or IPv6 address.

Customer Impact: The user will be able to clear IP ARPs and neighbors on UPF now.

Command Changes

2 new CLIs are introduced for clearing neighbor VPPs and 2 new CLIs are introduced for displaying the neighbor VPPs.

To clear IPv4 address neighbor VPPs, execute following command:

clear ip neighbors vpp port *port-number* **vlan** *vlan_tagid*

To clear IPv6 address neighbor VPPs, execute following command:

clear ipv6 neighbors vpp port *port-number* **vlan** *vlan_tagid*

NOTES:

- **vpp**—It clears ARP and neighbor table from the VPP for this context.
- **port** *port_number*—Slot or port number of the neighbor VPP.
- **vlan** *vlan_tagid*—It clears the VLAN NPU counters. It must be followed by VLAN tag ID.

To display IPv4 address neighbor VPPs, execute following command:

```
show ip neighbors vpp
```

To display IPv6 address neighbor VPPs, execute following command:

```
show ipv6 neighbors vpp
```

Example Configuration

Following is the example configuration for clearing IPv4 address neighbor VPPs:

```
clear ip neighbors vpp port 1/10 vlan 110 192.168.110.57
```

Following is the example configuration for clearing IPv6 address neighbor VPPs:

```
clear ipv6 neighbors vpp port 1/11 vlan 110 fd4d:5643:2886:6e::59:1
```

Following is the example configuration for displaying IPv4 address neighbor VPPs:

```
show ip neighbors vpp
 Time             IP             Flags     Ethernet             Interface
slot/port
 6493.6096     192.168.110.57     D     00:1b:21:87:13:2d     TenGigabitEthernet0/6/0.110
  1/10
```

Following is the example configuration for displaying IPv6 address neighbor VPPs:

```
show ipv6 neighbors vpp
 Time                     IP             Flags     Ethernet                 Interface
           slot/port
 13089.7977     fe80::250:56ff:fe87:af30    D     00:50:56:87:af:30
```

```
TenGigabitEthernet0/7/0.110     1/11 (ifc-6)
 13291.8022     fe80::be16:65ff:fe3d:6b43   D     bc:16:65:3d:6b:43
TenGigabitEthernet0/7/0.110     1/11 (ifc-6)
 13327.3504     fd4d:5643:2886:6e::59:1     D     00:50:56:87:af:30
TenGigabitEthernet0/7/0.110     1/11 (ifc-6)
```

Conditions for RCM Timers—CSCwe36370

Revision History

Revision Details	Release
First introduced.	21.28

Feature Changes

Previous Behavior: The planned UPF switchover timers were not configurable.

New Behavior: The planned UPF switchover timers are now configurable.

The configured value of the pending standby timer in UP must be greater than or equal to the planned switchover timer value in UPF. The condition for these timers is defined only in UP.

Command Changes

To configure the timers in UPF, use the following comfiguration:

```
configure
  rcm-service rcm_svc_name
    pending-standby-timeout pending_timer_value
    planned-standby-timeout planned_timer_value
    exit
```

NOTES:

- **pending-standby-timeout** *pending_timer_value*—Specify the pending standby timeout, in seconds. This timer is applicable only when UPF is in the Source UPF role. If the source UPF does not receive Standby state from RCM within this timeout period, it will revert back to Active from Pending Standby.

  pending_timer_value is an integer from 300 to 3600. Default: 300 seconds.

- **planned-standby-timeout** *planned_timer_value*—Specify the planned switchover timeout, in seconds. This timer is applicable only when UPF is in the Destination UPF role. Destination UPF will reload if the planned UPF switchover does not complete within this timeout period.

  planned_timer_value must be an integer from 300 to 3600. Default: 300 seconds.

To configure the timers in RCM, use the following configuration:

```
config
  k8 smf profile rcm-config-ep swo-timeouts { pre-switchover
preswitchover_timer_value | stage1-chkpt-flush stage1flush_timer_value |
```

```
stage2-chkpt-flush stage2flush_timer_value }
  exit
```

NOTES:

- **pre-switchover** *preswitchover_timer_value*—Specify the Source UPF preswitchover check timer in seconds. The planned UPF switchover will be aborted on timeout.

  preswitchover_timer_value is an integer from 15 to 3600. Default: 15 seconds.

- **stage1-chkpt-flush** *stage1flush_timer_value*—Specify the stage 1 checkpoint flush from Source UPF to CheckpointMgrs, in seconds. The planned UPF switchover will be aborted on timeout. The Source UPF will be reverted to Active and Destination UPF will be reloaded.

  stage1flush_timer_value is an integer from 15 to 3600. Default: 15 seconds.

- **stage2-chkpt-flush** *stage2flush_timer_value*—Specify the stage 2 checkpoint flush from Source UPF to CheckpointMgrs, in seconds. Failure or timeout of stage 2 checkpoint flush does not fail the planned UPF switchover.

  stage2flush_timer_value is an integer from 10 to 3600. Default: 10 seconds.

Configurable Init Wait Timer and Mass UPF Failure Timer—CSCwb66179

Revision History

Revision Details	Release
First introduced.	21.28.0

Feature Changes

RCM supports the following timers:

- **Init Wait Timer**—The Init Wait timer defers the registration of Init state UPs and registers the active UPs first. This timer starts only when RCM controller starts or when RCM moves to HA MASTER state.

  When RCM controller starts or moves to HA MASTER state, the RCM controller has no state and learns the UP state from the UPs itself. The Init state UPFs should not be assigned HostIDs that are already allocated to Active UPs.

- **Mass UPF Failure Timer**—The Mass UPF Failure timer starts when all UPs lose BFD connectivity with RCM. Depending on the network deployment, there could be network connectivity issue between RCM and UPs. If RCM cannot establish BFD connectivity to any UPF within the timeout period, then RCM HA switchover is performed.

Previous Behavior:

- The Init Wait timer starts only with the first UPF registration.

- The Init Wait Timer was not configurable and fixed to 300 seconds.
- The Mass UPF Failure timer was not configurable and fixed to 3 minutes.

New Behavior:

- The Init Wait timer starts only when the RCM controller starts or when RCM moves to HA MASTER state.
- The Init Wait timer is configurable using the **k8 smf profile rcm-config-ep init-wait-timeout** *init_wait_timeout* command.
- The Mass UPF Failure timer is configurable using the **k8 smf profile rcm-config-ep mass-upf-failure-timeout** *upf_failure_timeout* command.

Customer Impact:

- Reduced wait times for UPF registration.
- No change in behavior if the timer CLI commands are not used.
- The timer CLI commands can be used to change or disable the Init Wait timeout and Mass UPF Failure timeout.

Command Changes

Use the following RCM Ops Center CLI commands to configure the Init Wait timer and Mass UPF Failure timer.

```
k8 smf profile rcm-config-ep init-wait-timeout init_wait_timeout
k8 smf profile rcm-config-ep mass-upf-failure-timeout upf_failure_timeout
```

NOTES:

- **init-wait-timeout** *init_wait_timeout*: Specify the Init Wait timer, in seconds, as an integer from 0 to 300.

  Default: 300 seconds. A value of 0 disables the Init Wait timer.

- **mass-upf-failure-timeout** *upf_failure_timeout*: Specify the Mass UPF Failure timer, in minutes, as an integer from 0 to 60.

  Default: 3 minutes. A value of less than 3 minutes disables the timer. UPFs need at least three minutes to reload. When all UPFs are simultaneously reloaded as part of the UPF-RCM workflow, a value of less than three minutes can potentially cause false positive alarms.

CUPS Support for VMware Release 7

Revision History

Revision Details	Release
First introduced.	21.28.0

Feature Description

CUPS (control plane and user plane) using VPC-SI supports VMware Release 7. CUPS uses the VMware-based deployments as an alternate deployment model to reduce cost and complexity.

For more information, refer to the *Ultra Packet Core CUPS Control Plane Administration Guide* or *Ultra Packet Core CUPS User Plane Administration Guide*.

CUPS Support on VMware ESXi 6.7

Revision History

Revision Details	Release
First introduced.	21.28.0

Feature Description

CUPS supports the VMware-based deployment model on VMware ESXi 6.7, VMXNET3 for GW-C, and PCT-PT for GW-U for Intel 6248R CPU.

For more information, see the *Ultra Packet Core CUPS Control Plane Administration Guide* or *Ultra Packet Core CUPS User Plane Administration Guide*.

DDN Trigger for RA Packets—CSCwm47782

Revision History

Revision Details	Release
First introduced.	21.28.m28

Behavior Change

Previous Behavior: Previously DDN was not getting triggered for RA packets. Due to this, when the UE was in the idle state, it was not receiving any packets. Therefore, the packets and bandwidth were getting wasted at that time.

New Behavior: The DDN functionality is triggered for RA packets as well as it is considered as Downlink data packets by SAEGW-U. Therefore, the UE receives RA packets even in the idle state.

Deletion of ACS Configuration—CSCwf98047

Revision History

Revision Details	Release
First introduced.	21.28.m14

Behavior Change

The deletion of ACS Configuration will be effective for existing offloaded flows on UPF.

Previous Behaviour: When ACS configuration was deleted in the config path, the offloaded flows in UPF were not onloaded.

New Behavior: After deletion of ACS configuration, the offloaded flows will be onloaded and reprogrammed in UPF.

Customer Impact: You will observe stable and defined behaviour on UPF post deletion of the ACS configuration.

DI-Net Encryption

Revision History

Revision Details	Release
First introduced.	21.28.0

Feature Description

The CUPS-DI systems uses Galois/Counter Mode (GCM) encryption algorithm for DI-Net traffic encryption. The GCM algorithm replaces Advanced Encryption Standard Cipher Block Chaining (AES CBC) algorithm for better data protection and integrity.

Note The change in encryption algorithm requires a system reload.

The new encryption algorithm is configurable via boot parameter file. The GCM algorithm supports an authenticated encryption mode. On the decrypting side, GCM uses Additional Authentication Data (AAD) to authenticate the payload.

For more information, see the *DI-Net Encryption* chapter in the *Ultra Packet Core CUPS Control Plane Administration Guide* or *Ultra Packet Core CUPS User Plane Administration Guide*.

Disabling RCM Traps—CSCwe40690

Revision History

Revision Details	Release
First introduced.	21.28.m14

Behavior Change

Previous Behavior: All RCM traps are enabled by default.

New Behavior: The following SNMP traps are disabled by default to prevent flooding of less significant traps in the **rcm show-snmp-trap history** command.

- ActiveSessmgrConnected
- ActiveSessmgrDisconnected
- CheckpointAuditEnded
- CheckpointAuditStarted
- StandbySessmgrConnected
- StandbySessmgrDisconnected

To enable these traps, use the **k8 smf profile rcm-snmp-trapper-ep snmp-trapper clear-dflt-traps** command. After issuing this command, disable the traps using the **disable-trap** command.

Customer Impact: The **rcm show-snmp-trap history** command will not be flooded with traps that are very frequent and less significant.

DNS Snooping and Tethering Detection Bypass Support

Revision History

Revision Details	Release
First introduced.	21.28.0

Feature Description

In this release, the DNS Snooping and Tethering Detection Bypass features for ECSv2 are supported as part of L3, L4, and L7 rule combination in Ruledef.

See the *L3, L4, and L7 Rule Combination in Ruledef* chapter in the *Ultra Packet Core CUPS User Plane Administration Guide* and *Ultra Packet Core CUPS User Plane Administration Guide* for more information.

EDNS Enrichment

Revision History

Revision Details	Release
First introduced.	21.28.m10

Feature Changes

CUPS supports enrichment of EDNS requests to enrich and readdress DNS requests of subscribers who are subscribed to the parental control service.

When a subscriber subscribes to a parental control service, DNS requests by the subscriber are enriched with additional information (IMSI, MSISDN, APN) and readdressed to the dedicated DNS server for appropriate analysis and treatment. This additional information is configurable through an eDNS format that specifies different fields (tag values). These fields are encoded and appended to the DNS request header.

For more information, refer to the *EDNS Enrichment* chapter in the *UPC CUPS User Plane Administration Guide*.

Enhancement to Local Policy Fallback—CSCwk11686

Revision History

Revision Details	Release
First introduced.	2024.04.0

Feature Changes

The the local policy feature in CP is enhanced to support the following scenarios:

- While managing failures under local policy, the CCR-I request sent to the PCRF over the Gx session includes the QoS-Information and Default-EPS-Bearer-QoS AVPs.
- When the Gx session operates under local policy, the CP retrieves usage data from the UP and reports it to the PCRF via CCR-U during retry attempts after timer expiration. To fetch the usage report from UP, you must configure the **fetch-usage-from-up** command in the Local Policy Actiondef configuration mode. For more information on the command, see the *Local Policy Actiondef Configuration Mode Commands* chapter in *CLI Command Reference Guide*.

TCP and UDP Packet Statistics for Application Data in Show Subscriber Output-CSCwj64339

Revision History

Revision Details	Release
First introduced.	21.28.20

Behavior Change

Previous Behavior: The statistics **DNS to EDNS** in the show command **show subscriber user-plane-only full all** and **EDNS Encode Success** in the show command **show user-plane-service statistics analyzer name dns** were incremented for uplink TCP control packets or uplink UDP packets with no application data.

New Behavior: The statistics **DNS to EDNS** and **EDNS Encode Success** do not get incremented for TCP control packets or UDP packets with no application data.

EGPTC Message Validation_CSCwk30287

Revision History

Revision Details	Release
First introduced.	2024.03.0

Behavior Change

Previous Behvaior: If a Modify Bearer Request (MBR) on the S11 interface was received with IMSI, there was no validation check to find whether the tunnel on which the message was received belonged to the same IMSI.

New behavior: A validation check is now available to find if the Modify Bearer Request on the S11 interface is received with IMSI and the tunnel on which the message is received belongs to the same IMSI or not. In case the MBR received with a valid TEID is for a different IMSI, Modify Bearer Response is sent with cause "Context Not Found".

Customer Impact: There is a possibility for Modify Bearer Rejection .

Enabling Standalone RCM Without Keepalived—CSCwc12468

Revision History

Revision Details	Release
First introduced. CDETS ID: CSCwc12468	21.28.0

Feature Changes

This release supports a configurable CLI to use when you run RCM without keepalived achieve MASTER status. The MASTER status is required for UPF to register.

Previous Behavior: The CLI command to enable keepalived in RCM was not required in releases prior to 21.28.x.

New Behavior: This release supports a new CLI command to use when you run RCM without keepalived to achieve MASTER status.

```
k8 smf profile rcm-config-ep ha-standalone { true | false }
```

The default setting is **false**.

Customer Impact: From release 21.28.x onwards, customers running RCM without keepalived must set the CLI command to true.

END MARKER Handling during eNB Path Switchover for Multi-Bearer PDNs—CSCwj13323

Revision History

Revision Details	Release
First introduced.	21.28.20

Behavior Change

In CUPS, SNDEM flag is set in the FAR and sent towards UP. UP sends the GTP END MARKER on the GTPU tunnel for which SNDEM flag is received.

Previous Behavior: All the bearers receive GTP END MARKER even if there is a change in F-TEID in any one of the bearers.

New Behavior: UP sends the GTP END MARKER only for the bearer on which the F-TEID is changed.

Customer Impact: END MARKER is sent to the tunnels that are being switchover to a new eNB/gNB and not to every bearer binded to the same PDN Session.

Encryption of LI Information in RCM

Revision History

Revision Details	Release
First introduced.	21.28.0

Feature Description

The RCM configuration on CUPS UP must include encrypted public and private keys to secure the LI information. When this feature is enabled, the LI information that the RCM receives is in encrypted format.

Important The users of CUPS UP running the trusted builds must enable this feature. Otherwise, recovery is not supported for LI functionality.

All UPs (both active and standby) should have the same set of public and private keys. Active UP uses public key to encrypt the LI information before sending to RCM. Standby UP uses the private key to decrypt the information received from RCM. If keys are not present in active UP running the trusted build, LI information is not sent to RCM and it impacts the LI recovery functionality. Non-trusted build, with no key configuration, continue sending LI information as plain binary.

Important The keys once configured, cannot be removed.

For more information, refer to the *Lawful Intercept in CUPS* chapter in the *CUPS LI Guide*.

Enriching DNS Requests with Additional RRs

Revision History

Revision Details	Release
First introduced.	21.28.m23

Feature Changes

CUPS supports enrichment of DNS requests that contain Additional RRs.

The DNS requests are enriched by adding Option-Codes and Option-Data fields based on the configured EDNS format in the following scenarios:

- Presence of additional RRs of OPT RR type in the incoming DNS request

  If an OPT RR is present in the incoming request, it is deleted, and a new OPT RR is added as the first additional RR based on the configured EDNS format.

- Absence of additional RRs in the DNS request

  If no Additional RRs are present in the DNS request, enrichment is done by adding an OPT RR to the request.

- Presence of additional RRs other than OPT RR type in the DNS request

For more information, see the *EDNS Enrichment* chapter in the *UPC CUPS User Plane Administration Guide, Release 21.28*.

Exact Duration for SGW-CDR Fields—CSCwf12125

Revision History

Revision Details	Release
First introduced.	21.28.m14

Behavior Change

The "duration" and "changeTime" fields display the exact accumulated time in the following scenarios for S-GW CDR when:

- Zero volume CDR suppress configuration is enabled,
- Sx Session Report Request is received without data continuously, and
- Session Manager restarts

Previous Behavior: In the above scenario for SGW-CDR, the duration and changeTime fields in CDR are not accumulated.

New Behavior: In the above scenario for SGW-CDR, the duration and changeTime fields in CDR are accumulated.

Customer Impact: The changeTime and duration fields of CDR ELEMENTS display accurate information with the exact duration.

FAR Buffering Limit

Revision History

Revision Details	Release
First introduced.	21.28.m23

Feature Changes

With this release, the number of packets to be buffered per FAR on UP is configurable using the **buffering-limit far-max-packets** *far_max_packets* CLI in the ACS Configuration mode.

By default, 5 packets will be buffered per FAR. You can configure more number of FAR buffered packets to achieve QoS with fewer packet drops.

The **show user-plane-service statistics all** CLI is enhanced to display the dropped packets per FAR.

Command Changes

Use the following configuration to configure the maximum number of packets to be buffered per FAR:

```
configure
  active-charging service acs_service_name
    buffering-limit far-max-packets far_max_packets
    end
```

NOTES:

- **buffering-limit far-max-packets** *far_max_packets*—Specify the maximum number of packets to be buffered per FAR. *far_max_packets* must be an integer from 1 to 128.

  Default value: 5 packets

The packets received after maximum number of packets are already buffered and the subsequent packets are dropped. To view the number of times that packets in a specified range are dropped, use the **show user-plane-service statistics all** command.

The following is a sample output of this command:

```
[local]qvpc-si# show user-plane-service statistics all
…
  Data Statistics Related To Buffering:
    Packets Buffered:                  0    Bytes Buffered:                     0
    Packets Discarded:                 0    Bytes Discarded:                    0
    Packets Dropped per FAR (<=9)      0    Packets Dropped per FAR (10-19)     0
    Packets Dropped per FAR (20-49)    0    Packets Dropped per FAR (30-39)     0
    Packets Dropped per FAR (40-49)    0    Packets Dropped per FAR (>=50)      0
```

Generating SNMP Trap from User Plane during Warning Interval—CSCwc01756

Revision history

Revision Details	Release
First introduced.	21.28.m10

Feature Changes

In StarOS, if the password is not reset before the expiration date, you get locked out from the configured gateways. You are allowed to log on back only when the administrator resets the password manually.

Previous Behavior: When the StarOS login password is about to expire, only a warning message was displayed during login attempt. This caused disruption in the RCM workflow and the warning message was disabled for the RCM workflow.

New Behavior: A PasswordExpiryNotification SNMP trap is generated daily every 24 hours during the warning interval before password expiry.

Customer Impact: The SNMP trap gets generated every 24 hours along with the warning message.

For more information, see the *Password Expiration Notification* chapter in the *VPC-SI System Administration Guide*.

Handling Downlink Packets during System in Overload Control State—CSCwm71777

Revision History

Revision Details	Release
First introduced.	21.28.m28

Behavior Change

Previous Behavior: When the existing flow (in any state) gets on-loaded due to any control event or any other reason, it again evaluates if the system is in overload or self-protection mode. If the system is in either of states, the stream goes into the state. Therefore, the packets were getting droppped.

New Behavior: When the system is in Self-protection (overload) mode and the existing flow gets on-loaded due to any control event, in such cases there will not be any change to flow/streams.

Handling Non-SYN TCP and UDP Packets for NAT Subscribers—CSCwf64696

Revision History

Revision Details	Release
First introduced.	21.28.m10

Behavior Change

To avoid generation of empty EDRs, CUPS supports the following functions:

- For TCP, the non-SYN packets for NAT subscribers will be dropped without creating a new flow.
- For UDP, the packets will be buffered while IP allocation is in progress and will be processed once IP allocation is complete.

Previous Behavior: Each non-SYN TCP packet created a flow and generated an empty EDR when the packet got dropped due to NAT or Firewall, and the flow was cleared. The UDP packets of new flows were dropped while NAT IP allocation was in progress for a subscriber.

New Behavior: For a NAT-enabled subscriber, the non-SYN TCP packet will be dropped without creating a flow. The UDP packets of new flows will not be dropped while NAT IP allocation is in progress for a subscriber. These packets will be buffered and processed once NAT IP allocation is successful.

Handling Simultaneous Gy RARs from Different DRAs with Different RGs

Revision History

Revision Details	Release
First introduced.	21.28.m1

Feature Description

CUPS supports multiple Diameter Routing Agents (DRA) to prevent the abort of pending Credit Control Request–Update (CCR-U) requests from previous Reauthorization Requests (RAR) with a different host or peer on the Gy interface.

P-GW accepts different rating-groups (RG) from different peers by configuring the **diameter pending-ccau allow-on-rar-peer-switch** CLI command in the ACS configuration mode. This command allows you to configure the DCCA client to prevent the abort of a pending CCR-U request.

For more information on the multiple DRA support in P-GW, see the *Support for Multiple DRA over Gy Interface* chapter in the *P-GW Administration Guide*.

How it Works

This section describes how the multiple DRA feature works in CUPS.

P-GW and CUPS handle the collision scenarios differently. In legacy P-GW, each CCR-U with FORCED REAUTHORIZATION is sent to the corresponding DRAs.

In CUPS, the user plane fetches every CCR-U that is sent along with the current usage report. During collision, if more than one specific RAR is received at the same time from different DRAs for the respective rating groups, the control plane marks the Gy-URR buckets, and sends Sx Session Modification Request to the user plane. The user plane sends back the current usage reports to the control plane for the requested Gy-URR bucket in Sx Session Modification Response. If RAR is received from different DRAs, the peer switch happens. In CUPS, each CCR-U with FORCED REAUTHORIZATION for the requested rating groups is sent to the peer DRA of the latest path switched.

The following call flow illustrates how P-GW accepts both RGs from different peers.

Figure 1: Multiple DRA Call Flow in CUPS

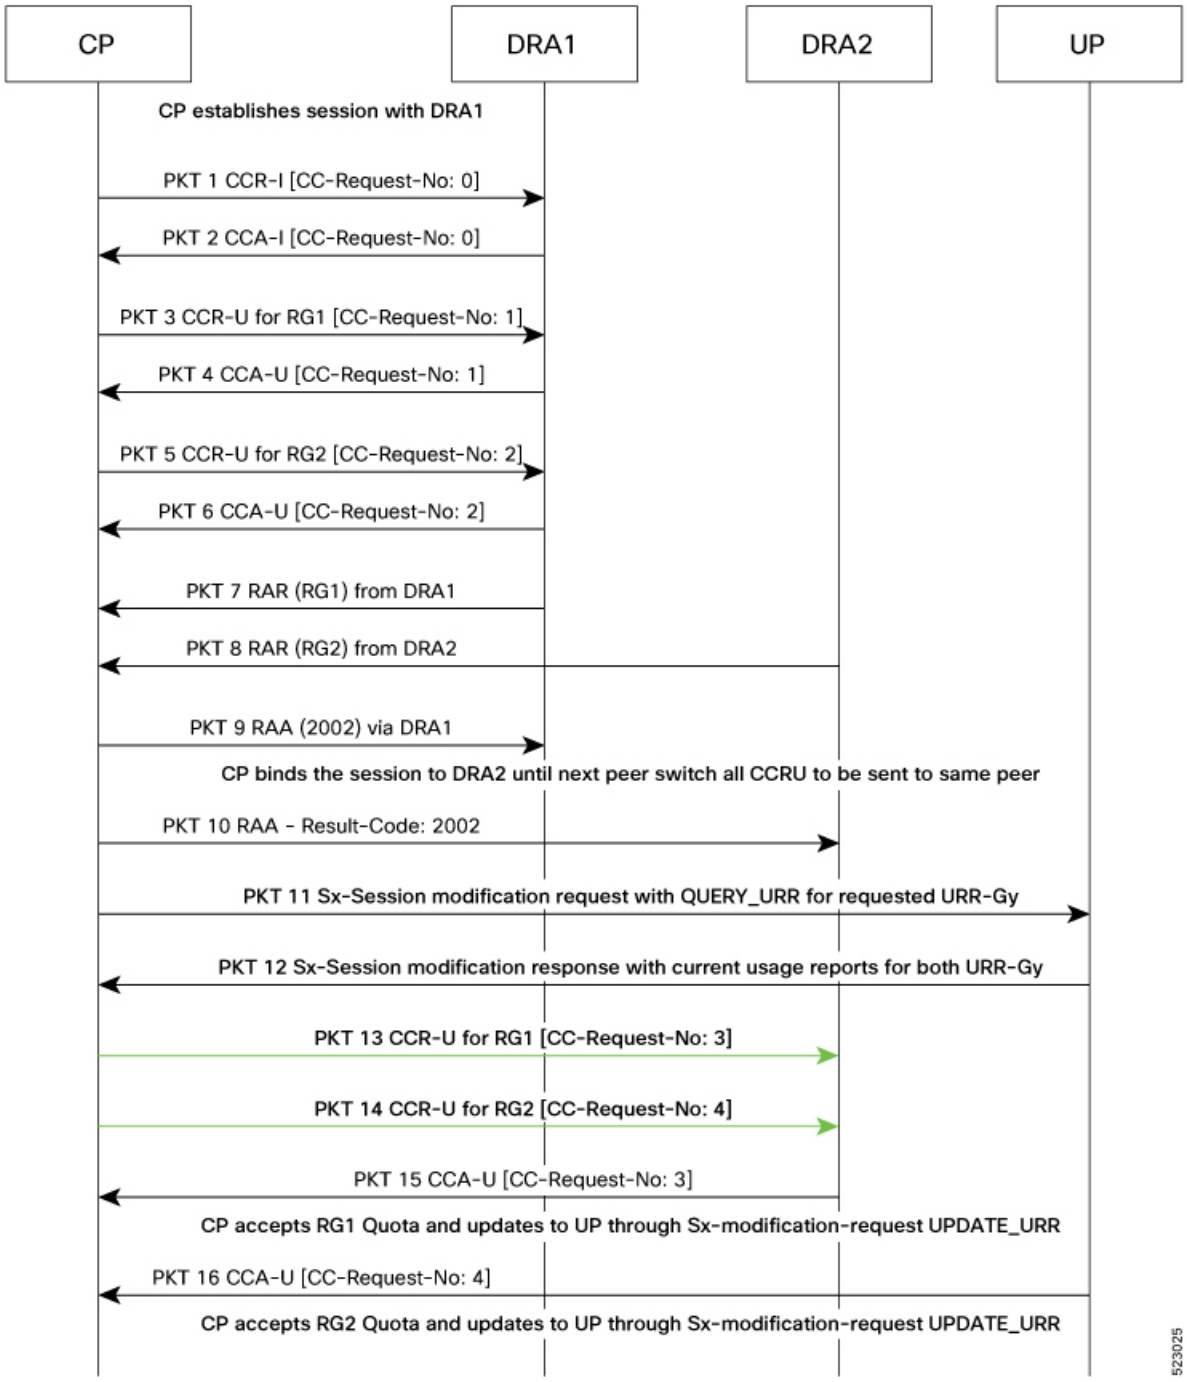

Configuring the Feature

To configure the handling of multiple RAR requests involving multiple DRAs, use the following configuration:

```
configure
  context context_name
  active-charging service acs_service_name
    credit-control [ group cc_group_name ]
      diameter dictionary dictionary
        [ no ] diameter pending-ccau allow-on-rar-peer-switch
        end
```

NOTES:

- **diameter dictionary** *dictionary*: Set the diameter dictionary to handle different DRAs.

  For example: **diameter dictionary** *dcca-custom-26*

- **diameter pending-ccau allow-on-rar-peer-switch**: Allow the DCCA client to prevent the abort of pending CCAU requests.
- **no diameter pending-ccau allow-on-rar-peer-switch**: Disable the DCCA client from preventing the abort of pending CCAU requests.

Monitoring and Troubleshooting

This section provides the monitoring and troubleshooting information for the multiple DRA feature.

Show Commands and Outputs

This section provides information regarding show commands and outputs in support of this feature.

show active-charging service all

Table 1: show active-charging service all

Field	Description
pending ccau:	
allow-on-rar-peer-switch	Displays "Enabled or "Disabled" to indicate the abort of pending CCA-U request if RAR is received from different host or peer on the Gy interface. If this feature is enabled, the functionality is applicable only to new Diameter sessions.

HTTP Request Methods during Redirection—CSCwi06981

Revision History

Revision Details	Release
First introduced.	21.28.m18

Behavior Change

Previous Behavior: HTTP redirection was applied only on HTTP GET and not applied on other methods such as HTTP POST.

New Behavior: HTTP redirection will now be applied on all HTTP Request methods.

Intercept Provisioning Method in UP—CSCwh28931

Revision History

Revision Details	Release
First introduced.	21.28.m14

Behavior Change

Previous Behavior: The Provisioning Method field in the output of the **show lawful-intercept active-only** and **show lawful-intercept camp-on-triggers** CLI commands was displayed in UP.

New Behavior: The user plane does not provision any intercepts as all provisioning is on the control plane. Therefore, the Provisioning Method field on UP will be restricted.

International Roaming

Revision History

Revision Details	Release
First introduced.	21.23.22 21.28.0

Feature Description

Previous Behavior: During the Initial Attach procedure, the SGW-S5U interface is configured with IPv4v6 and the PGW-S5U interface is configured with IPv4 as part of the GTPU configurations. The PGW-C does not recognize the interface details of PGW-S5U. The Sx Session Establishment Request that is initiated from CP includes the GTP-U/UDP/IPv6 Outer Header Removal (OHR) in Create PDR and a similar Outer Header Creation (OHC) in Create Far. This results in a mismatch of IE in UP and the Sx Session Establishment Request gets rejected due to which the call fails with the OUTER_HDR_REMOVAL value as PFCP_CAUSE_MANDATORY_IE_INCORRECT.

New Behavior: PGW-U supports common OHR and OHC types for IPv4 and IPv6 based on PGW-U interfaces. GTP-U/UDP/IPv6 support is added for OHR and IPv4v6 support is added for OHC.

Customer Impact: None.

For more information, see the *International Roaming* chapter in the *Ultra Packet Core CUPS Control Plane Administration Guide* or *Ultra Packet Core CUPS User Plane Administration Guide*.

Keepalived Track Interface and Virtual Routes Support in RCM—CSCwb69008

Revision History

Revision Details	Release
First introduced.	21.28.0

Feature Changes

RCM supports CLI commands to configure track interface and IPv4 virtual routes for the Keepalived pod.

When VIP gets attached to the service interface (when RCM moves to MASTER state), addition of IPv4 routes is required for non-host networking Bfdmgr.

Previous Behavior: RCM did not support CLI commands to configure track interface and IPv4 virtual routes for the Keepalived pod.

New Behavior: RCM supports Ops Center CLI commands to configure additional tracking interfaces and IPv4 virtual routes for the Keepalived pod.

Customer Impact: There is no impact if the CLI commands are not used. The CLI commands are backward compatible

Command Changes

Use the following RCM Ops Center CLI commands to configure the track interface and IPv4 virtual routes in the Keepalived pod:

```
k8 smf profile rcm-keepalived-ep vrrp-config group group_name
   ipv4-route route_serial_number
      destination host_network_ipv4 mask ipv4_mask gateway host_ipv4 device
interface_name
   track-interface track_interface
   exit
```

NOTES:

- **ipv4-route** *route_serial_number*: Configures the Keepalived IPv4 virtual routes.
- **track-interface** *track_interface*: Configures the Keepalived track interface.

Kubernetes Version Upgrade for VM-based RCM

Revision History

Revision Details	Release
First introduced.	• 2024.03.0 • 21.28.m25

Feature Description

In a VM-based RCM deployment the Kubernetes version is upgraded to 1.30.1.

Monitoring and Troubleshooting

If the Calico CNI pod does not start post Kubernetes upgrade on the VM-based RCM, perform the workaround steps:

1. Check if the interface specified in `/etc/netplan/50-cloud-init.yaml` is assigned to IPv4 address on the boot.

Note It is recommended to assign the IPv4 address through DHCP.

2. If the calico pod does not start, run **ifconfig -a** to find the RCM VM physical interface.
3. Add a manual IP address through cloud-init network configuration.
4. Redeploy the RCM VM with manual IP address allocation.

For example, If the RCM VM physical interface is *enp1s0*, create file *network-config.yaml* with the following information.

```
version: 2
ethernets:
  enp1s0:
    dhcp4: false
    dhcp6: false
    addresses:
      - 10.105.201.206/24
    gateway4: 10.105.201.1
    nameservers:
      addresses: [72.163.128.140]

and then
Example:
cloud-localds -H rcm-1 -N network-config.yaml -m local rcm-k8s-1.30.1-seed.iso
user-data-lab-1.30.yml
```

Create the Cloud Init ISO (CDROM) with option "-N network-config.yaml"

```
cloud-localds -H rcm-1 -N network-config.yaml -m local rcm-k8s-1.30.1-seed.iso
user-data-lab-1.30.yml
```

LI Keepalive Message Support

Revision History

Revision Details	Release
First introduced.	21.28.5 21.28.m6

Feature Changes

CUPS supports S8HR LI application heartbeat messages to be sent from BBIFF to LMISF at periodic intervals. This feature avoids IPsec tunnel termination or TCP socket clearance in the firewall present between BBIFF and LMISF. The keepalive message is sent on both the Xia and Xib interfaces.

This feature is configurable and the command must be executed only on CP. If executed on UP, it throws an error.

Previous Behavior: An IPsec tunnel connects the BBIFF to the firewall when in between BBIFF and LMISF. When the idle timer expires, the firewall clears the port map, resulting in termination of the IPsec tunnel or clearance of TCP sockets.

New Behavior: BBIFF sends a new unidirectional keepalive message to LMISF periodically after a configurable time interval. This avoids IPsec tunnel termination or TCP socket clearance in the firewall present between BBIFF and LMISF.

For more information, contact your Cisco Account representative.

MBR and UBR Collision Handling—CSCwh47513

Revision History

Revision Details	Release
First introduced.	21.28.m14

Behavior Change

For Pure-S calls, S-GW drops the PGW-initiated UBR during Modify Bearer Request (MBR) and Update Bearer Request (UBR) collision in the network.

Previous Behavior: If S-GW receives the PGW-initiated UBR while processing MBR for ULI change, then S-GW rejects UBR with cause "No Resource Available".

New Behavior: If S-GW receives the PGW-initiated UBR while processing MBR for ULI, MME, or eNodeB change, then S-GW drops UBR silently. P-GW then retries UBR and S-GW processes the request after MBR.

Monitor Subscriber Trace–CSCwk67358

Revision History

Revision Details	Release
First introduced.	2024.04.0

Behavior Change

Previous Behavior: The "W - UP PCAP Trace (ON)" option was visible on both the CUPS CP and UP when running the **monitor subscriber** command. However, the upcoming call and IMSI monitoring outputs were not captured.

New Behavior: The "W - UP PCAP Trace (ON)" option is now visible and manageable only on the CUPS CP. Users cannot toggle the "W - UP PCAP Trace" option (ON/OFF) on the CUPS UP.

Customer Impact: This change allows subscribers to capture traces on the CP independently from the UP, eliminating the need to start traces on both planes.

Note When the "W - UP PCAP Trace" is enabled on the CP, do not run the **monitor subscriber** command on the UP.

Namespace Option in RCM Script—CSCwd79932

Revision History

Revision Details	Release
First introduced.	21.28.m14 21.28.m3 21.28.3

Feature Changes

The **apply_config_v2** RCM script is modified to support the namespace option.

Previous Behavior: The script did not support the namespace option.

New Behavior: The script supports the "N" option to provide the namespace argument. By default, the value is **rcm**.

```
bash -x ./apply_config_v2.sh -g 1 -n -c UPCommon.cfg -P
/var/lib/smi/data/common_config/config4.cfg -G 4 -p connect_file -k
password -N rcm
```

NNRF Service for RCM—CSCwc49421

Revision History

Revision Details	Release
First introduced.	21.28.0

Feature Changes

RCM host configuration includes a new service type "NNRF". In the string-based approach, RCM acts as a configurator and pushes the configuration of all services including the NNRF service.

In the Yang-based approach, NSO acts as a configurator and allows configuration of only the service names in RCM. NSO pushes the whole configuration including the NNRF service.

Important The script support does exist currently for NNRF service type. So, you must manually configure this service type.

New Messaging or IE on Sx Interface_CSCwk37340

Revision History

Revision Details	Release
First introduced.	21.28.m25

Behavior Change

Previous Behavior: When the Update Bearer Response includes User Location Information (ULI) and the **gtpp trigger uli** is enabled, there is no Query URR going in the Sx Session Modification Request towards the User Plane.

New Behavior: When the Update Bearer Response includes User Location Information (ULI) or TimeZone (TZ) and the **gtpp trigger uli** or **gtpp trigger ms-timezone-change** is enabled, the Query URR is triggered in the Sx Session Modification Request towards the User Plane.

Customer Impact: The Control Plane could send a new IE of Query URR in the existing message while sending the Sx Session Modification Request on receiving the Update Bearer Response.

If there is nothing to update to the User Plane and only ULI or TZ is present, the Control Plane can explicitly trigger a Sx Session Modification Request towards the User Plane to query URR.

Optimizing GTPU Memory Usage—CSCwf21120

Revision History

Revision Details	Release
First introduced.	21.28.m14 21.28.m10

Behavior Change

Previous Behavior: For user plane service and Sx-u service, the GTPU peers were freed only after reaching one million peers.

New Behavior: For user plane service and Sx-u service, the GTPU peer entries will be removed when they become inactive. This happens after the last session associated with a GTPU peer is terminated.

Customer Impact: The user might observe differences in the output of CLI commands related to GTPU peers. Since only active peers remain in the system, the information displayed will pertain to that of active peers. This change will reduce the memory usage of GTPU manager.

Planned Switchover Timers on RCM—CSCwd35392

Revision History

Revision Details	Release
First introduced.	21.28.m14 21.26.17

Feature Changes

UPF supports the following timers for planned switchover through RCM:

- Preswitchover timer that defaults to 15 seconds

- Stage 1 checkpoint flush timer from old Active UPF to Checkpointmgrs that defaults to 15 seconds
- Stage 2 Checkpoint flush timer (non critical) from old Active UPF to Checkpointmgrs that defaults to 10 seconds

Command Changes

Use the following RCM OpsCenter Configuration mode CLIs to configure the following timers:

- Preswitchover timer:

  k8 smf profile rcm-config-ep swo-timeouts pre-switchover *preswitchover_timeout*

- Stage 1 Checkpoint Flush timer:

  k8 smf profile rcm-config-ep swo-timeouts stage1-chkpt-flush *stage1_flush_timeout*

- Stage 2 Checkpoint Flush timer:

  k8 smf profile rcm-config-ep swo-timeouts stage2-chkpt-flush *stage2_flush_timeout*

NOTES:

- **k8 smf profile rcm-config-ep swo-timeouts pre-switchover** *preswitchover_timeout*: Specify the timeout for preswitchover, in seconds. *preswitchover_timeout* must be an integer from 15 to 3600.

  Default value: 15 seconds

- **k8 smf profile rcm-config-ep swo-timeouts stage1-chkpt-flush** *stage1_flush_timeout*: Specify the timeout for stage 1 checkpoint flush from old Active UPF to checkpointmgrs, in seconds. *stage1_flush_timeout* must be an integer from 15 to 3600.

  Default value: 15 seconds

- **k8 smf profile rcm-config-ep swo-timeouts stage2-chkpt-flush** *stage2_flush_timeout*: Specify the timeout for stage 2 checkpoint flush (non-critical) from old Active UPF to checkpointmgrs, in seconds. *stage2_flush_timeout* must be an integer from 15 to 3600.

  Default value: 10 seconds

Port Number Behavior in EDR Module Configuration—CSCwi42567

Revision History

Revision Details	Release
First Introduced.	2024.02.0 (21.28.m23)

Behavior Change

After Cisco SSH or SSL upgrade, the port number behavior has changed for the EDR-module configuration.

Previous Behavior: In the EDR-module configuration, the default SFTP port number "0" was selected automatically and connected to port 0 when colon was specified.

For example:

cdr transfer-mode push primary url sftp://root:starent@192.0.2.1:/root/EDR/ via local-context

New Behavior: The EDR-module configuration allows the new default SFTP port number "22" or disallow the port number without specifying a colon.

If colon is specified after host, the port number is mandatory. The default SFTP port number is 22. If colon is not specified after host, the port number need not be entered and the default SFTP port number is used.

For example:

cdr transfer-mode push primary url sftp://root:starent@192.0.2.1:22/root/EDR/ via local-context(with default SFTP port number)

cdr transfer-mode push primary url sftp://root:starent@192.0.2.1/root/EDR/ via local-context (without colon)

For more information about **cdr** command, see the EDR Module Configuration Mode Commands chapter in the *Command Line Interface Reference Guide*.

Prioritizing IMEI over MAC Address

Revision History

Revision Details	Release
First introduced.	21.28.m22

Feature Changes

Previous Behavior: If IMEI was received in Create Session Request for a Wi-Fi call, the ePDG CDR encoded MAC address in the servedIMEISV field. The MAC address was given preference over IMEI and encoded in servedIMEISV. With this behavior, CDR processing was affected.

New Behavior: IMEI will be prioritized over MAC address and will be sent in the servedIMEISV field. The servedIMEIsv field in CDR is optional. If Create Session Request for a Wi-Fi call has both IMEI and MAC address, then IMEI is encoded in servedIMEIsv.

This behavior is configurable using the **gtpp prioritize-imei-over-mac-address** CLI in the GTPP Server Group Configuration mode. If the CLI is not configured, the existing behaviour will take effect.

Command Changes

Use the following configuration to prioritize IMEI over MAC address and encode IMEI in the servedIMEISV field of the CDR. If IMEI is not available, the servedIMEISV field will not be present in the CDR.

```
configure
  context context_name
    gtpp group group_name
      [ default | no ] gtpp prioritize-imei-over-mac-address
      end
```

NOTES:

- **default | no**—Specify either one of the options to prioritize MAC over IMEI and encode it in servedIMEISV field. This is the existing behavior.

  If MAC is not available, then IMEI is encoded in the servedIMEISV field.

QCI 67 Support

Revision History

Revision Details	Release
CUPS supports Standardized Mission Critical and Push-to-Talk (MC/PTT) QCI-67 application standard, which is used in mission critical communications.	2024.03.0

Standard QCI-67 Support

CUPS supports the new standard Quality of Service Class Identifier value (QCI) 67 in addition to the existing values 65, 66, 69 and 70.

This QCI 67 feature supports the following functionality:

- Creates, delets, and updates Dynamic rules from PCRF.
- Creates and deletes Pre-defined rules.
- Allows LTE to Wi-Fi HO for Bearer with QCI.
- Supports S2B/S2A HO with MC QCIs.
- Supprts Bulkstats for APN and SAEGW.
- MC-QCI values supports X-header insertion, DSCP marking, and EDRs.

For more information, refer to the Standard QCI support chapter in the *Ultra Packet Core CUPS User Plane Administration* and *Ultra Packet Core CUPS Control Plane Administration Guides*.

Limitations

The following are the known limitations of this feature:

- Does not support the overall eMPS functionality.
- If **require ecs credit-control session-mode per-subscriber** is configured, then URR is treated for entire subscriber session including secondary bearers which can lead to a problem in some applications. In CUPS, use the **credit-control-client override session-mode per-sub-session** command at the APN level to override the session mode configuration.

Configure QCI

Use the **qci-qos-mapping** CLI command to configure the DSCP marking for QCI.

```
config
  qci-qos-mapping name
    qcinum { gbr }  [ { downlink | uplink } [ user-datagram dscp-marking
dscp-marking-value ][ encaps-header { copy-inner | copy-outer | dscp-marking
dscp-marking-value ] ]
  end
```

NOTES:

- **qci***num*: Specify standard QoS Class Identifier between 1-9, 65,66,67,69,70,80,82,83 - integer 1..83.
- **gbr**: Specifies that this QCI type is Guaranteed Bit Rate (GBR).
- **downlink**: Configures parameters for downlink traffic.
- **uplink**: Configures parameters for uplink traffic.
- **encaps-header { copy-inner | copy-outer | dscp-marking** *dscp-marking-value***}**: Specifies that the DSCP marking must be set on the encapsulation header for IP-in-IP, GRE, or GTP encapsulation.
  - **copy-inner** : Specifies that the DSCP marking is to be acquired from the UDP headers within the encapsulation.
  - **copy-outer** used to copy the DSCP value coming in the data packet from S1u interface to the data packet sent on the S5 interface and vice-versa.
  - **dscp-marking**dscp-marking-value : Specifies that the DSCP marking is to be defined by this keyword.

    dscp-marking-value is expressed as a hexadecimal number from 0x00 through 0x3F.
- **user-datagram dscp-marking** *dscp-marking-value*: Specifies that the IP DSCP marking is to be defined by this keyword.

  dscp-marking-value is expressed as a hexadecimal number from 0x00 through 0x3F.

Monitoring and Troubleshooting

This section provides information regarding monitoring and troubleshooting the feature.

Show Command(s) and/or Outputs

This section provides information regarding show commands and/or their outputs in support of this feature.

show subscribers saegw-only imsi *<imsi>*

The output of this command includes QCI field under Bearer QOS and PCRF Authorized Bearer QOS as shown in the following example:

```
Bearer QoS        :
   QCI      : 67
   ARP              : 0x054
   PCI              : 1 (Disabled)
   PL               : 5
   PVI              : 0 (Enabled)
   MBR Uplink(bps)  : 0                          MBR Downlink(bps)   : 0
   GBR Uplink(bps)  : 0                          GBR Downlink(bps)   : 0


PCRF Authorized Bearer QoS:
   QCI: 67
   ARP: 0x054
   PCI: 1 (Disabled)
   PL : 5
   PVI: 0 (Enabled)
   MBR uplink (bps): n/a                      MBR downlink (bps): n/a
   GBR uplink (bps): n/a                      GBR downlink (bps): n/a
   Uplink APN AMBR (bps): 200000              Downlink APN AMBR (bps): 300000
```

show user-plane-service gtpu statistics

The output of this command includes the following fields under QCI67 as shown in the following example:

```
QCI 67:
Uplink Packets: 0 Uplink Bytes: 0
Downlink Packets: 0 Downlink Bytes: 0
Packets Discarded: 0 Bytes Discarded: 0
```

show apn statistics

The output of this command includes the following fields:

QCI 67:

- Bearer Active
- Bearer setup
- Bearer Released
- Bearer Rehected
- Uplink Bytes forwarded
- Downlink Bytes forwarded
- Uplink pkts forwarded
- Downlink pkts forwarded
- Uplink Bytes dropped
- Downlink Bytes dropped

- Uplink pkts dropped
- Downlink pkts dropped
- Uplink Bytes dropped(MBR Excd)
- Downlink Bytes dropped(MBR Excd)
- Uplink pkts dropped(MBR Excd)
- Downlink pkts dropped(MBR Excd)

```
QCI 67:
    Bearer Active:                              0   Bearer setup:
      0
    Bearer Released:                            0   Bearer Rejected:
      0

    Uplink Bytes forwarded:                     0   Downlink Bytes forwarded:
      0
    Uplink pkts forwarded:                      0   Downlink pkts forwarded:
      0
    Uplink Bytes dropped:                       0   Downlink Bytes dropped:
      0
    Uplink pkts dropped:                        0   Downlink pkts dropped:
      0
    Uplink Bytes dropped(MBR Excd):             0   Downlink Bytes dropped(MBR Excd):
      0
    Uplink pkts dropped(MBR Excd):              0   Downlink pkts dropped(MBR Excd):
      0
```

show qci-qos-mapping table name *name*

The output of this command displays QCI 67 is configured in the **qci-qos-mapping table**.

```
[local]qvpc-si# show qci-qos-mapping table
all       name
[local]qvpc-si# show qci-qos-mapping table name qci-qos1
QCI-QOS Table Name: qci-qos1
Qci:                  67
  uplink:   user-datagram dscp-marking 0x2e encaps-header dscp-marking 0x2e
  downlink: user-datagram dscp-marking 0x2e  encaps-header dscp-marking 0x2e
  pre-rel8-qos-mapping:  n/a         qci type:              gbr
```

show active-charging sessions full all

The output of this command includes the following fields:

```
Dynamic Charging Rule Definition(s) Configured:
Name                 Prior Content-Id Chrg-Type Rule Parameters
-------------------- ----- ---------- --------- --------------------------
   mcptt_audio_rule9   129      13009      Both Gate Status:      Allow All
                                                QoS Class Identifier:    67
                                                ARP Priority Level:       2
                                                Reporting Level: Rating Grp
                                                Metering Method: Durn + Vol
                                                Uplink MBR:           49000
                                                Downlink MBR:         49000
                                                Uplink GBR:           49000
                                                Downlink GBR:         49000
```

Bulk Statistics

Run the following CLI to check the counters available for APN and SAEGW schema.

show bulkstats variables saegw | grep qci67

The following list of bulkstats variables are added in **saegw** schema.

- sgw-totepsbearact-qci67
- sgw-totepsbearset-qci67
- sgw-totepsbearrel-qci67
- sgw-totepsbearmod-qci67
- sgw-totepsbearrel-dedrsn-pgw-qci67
- sgw-totepsbearrel-dedrsn-slerr-qci67
- sgw-totepsbearrel-dedrsn-s5err-qci67
- sgw-totepsbearrel-dedrsn-s4err-qci67

Example:

```
[local]qvpc-si# show bulkstats variables saegw | grep qci67
189 saegw %sgw-totepsbearact-qci67%               Int32    0   Gauge
204 saegw %sgw-totepsbearset-qci67%               Int32    0   Counter
221 saegw %sgw-totepsbearrel-qci67%               Int32    0   Counter
236 saegw %sgw-totepsbearmod-qci67%               Int32    0   Counter
252 saegw %sgw-totepsbearrel-dedrsn-pgw-qci67%    Int32    0   Counter
269 saegw %sgw-totepsbearrel-dedrsn-slerr-qci67%  Int32    0   Counter
285 saegw %sgw-totepsbearrel-dedrsn-s5err-qci67%  Int32    0   Counter
301 saegw %sgw-totepsbearrel-dedrsn-s4err-qci67%  Int32    0   Counter
```

Similarly, you can use the **show bulkstats variables apn | grep qci67** command to view list of bulkstats variables in **apn** schema.

RCM Endpoint Statistics—CSCwf06065

Revision History

Table 2: Revision History

Revision Details	Release
First introduced.	21.28.m14

Behavior Change

The RCM checkpoint manager now supports aggregate counters for endpoint statistics and Prometheus metrics.

Previous Behavior: RCM supported checkpoint statistics at the session level only.

New Behavior: RCM supports aggregate counters at the checkpoint-manager instance level under the**rcm endpointstats checkptmgr** CLI command and Prometheus metrics.

Customer Impact: The customer can view the new counters in the output of the endpoint statistics command.

RCM Helm Version Upgrade

Revision History

Revision Details	Release
First introduced.	21.28.0

Feature Changes

The RCM Helm version is upgraded from Helm2 to Helm3 for Kubernetes (K8s) pod deployment.

Previous Behavior: RCM used Helm2 and tiller for K8s pod deployment.

New Behavior: RCM uses only Helm3 for K8s pod deployment.

RCM Security Enhancements

Revision History

Revision Details	Release
First introduced.	21.28.0

Feature Changes

As part of security enhancement, RCM supports the following functionality:

- Partition Usage in RCM VM—In RCM VM, the */tmp* and */var/tmp* directories are mounted as separate partitions to prevent privilege escalation attacks.
- RCM provides flexibility to configure the host-networking mode for SNMP trapper pod. The **k8 smf profile rcm-snmp-trapper-ep snmp-trapper host-networking { false | true }** CLI command configures the SNMP trapper pod in host networking mode and non-host networking mode.
- RCM supports the conversion of host networking pods to non-host networking mode for restricting pod access to host network namespace. The CLI commands **k8 smf profile rcm-bfd-ep host-networking { true | false }** and **k8 smf profile rcm-bfd-ep node-port-enabled { true | false }** can be configured to run BFDmgr in non-host networking mode.
- RCM supports the tracking interface and IPv4 virtual-routes configuration for the Keepalived pod. The IPv4 virtual-routes configuration installs routes when RCM moves to MASTER state.

Command Changes

Use the following RCM Ops Center CLI commands to configure the following functionality:

- To configure the SNMP trapper pod in host networking mode and non-host networking mode:

  k8 smf profile rcm-snmp-trapper-ep snmp-trapper host-networking { false | true }

- To configure host networking mode and non-host networking mod in BFDmgr:

  k8 smf profile rcm-bfd-ep host-networking { true | false }

  Default value: **true**

- To configure node port:

  k8 smf profile rcm-bfd-ep node-port-enabled { true | false }

  Default value: **false**

  The node port must be set to **true** when host networking is set to **false**.

- To configure tracking interfaces in Keepalived pod:

  k8 smf profile rcm-keepalived-ep vrrp-config group *vrrp_group_name*
  track-interface *interface_name*
  exit

- To configure IPv4 virtual routes in Keepalived pod:

  k8 smf profile rcm-keepalived-ep vrrp-config group *vrrp_group_name*
  ipv4-route *route_serial_number*
  destination *host_network_ipv4* **mask** *ipv4_mask* **gateway** *host_ipv4* **device** *interface-name*
  exit
  exit

RCM SNMP Traps History

Revision History

Revision Details	Release
First introduced.	21.28.0

Feature Changes

The **rcm show-snmp-trap history** CLI command displays the history of SNMP event traps.

Previous Behavior: RCM did not support any command to display the SNMP trap history.

New Behavior: RCM supports the **rcm show-snmp-trap history** CLI command to display the SNMP trap history. This command displays details for the latest 5000 SNMP traps.

Customer Impact: This command eases debugging with the detailed history of SNMP traps.

RCM Statistics Information—CSCwe42938

Revision History

Revision Details	Release
First introduced.	2023.02.1

Behavior Change

The output of the rcm show-statistics bfdmgr command displays additional information.

Previous Behavior: The **rcm show-statistics bfdmgr** command displayed the following information:

- The minRx and minTx values in microseconds
- The locally configured multipier
- No down detect time

New Behavior: The updated **rcm show-statistics bfdmgr** command displays the following information:

- The minRx and minTx values in milliseconds
- The remotely configured multiplier when the BFD state is STATE_UP
- The down detect time in milliseconds

Customer Impact: The updated software displays the additional information. The **rcm show-statistics bfdmgr** command also displays values negotiated by both locally configured Tx/Rx and remotely configured Rx/Tx.

RCM Support for Cisco SSL—CSCwd32422

Revision History

Revision Details	Release
First introduced.	21.28.m14

Behavior Change

RCM uses Cisco SSL instead of OpenSSL.

Previous Behavior: The RCM VM and rcm-strongswan pod used the following OpenSSL version:

`OpenSSL 1.1.1f 31 Mar 2020`

New Behavior: The RCM VM and rcm-strongswan pod uses the following Cisco SSL version:

`CiscoSSL 1.1.1q.7.2.440`

Returning Correct PFCP Cause Code—CSCwh00402

Revision History

Revision Details	Release
First introduced.	21.28.m14

Behavior Change

UP now sends the PFCP error cause code PFCP_CAUSE_NO_RESOURCE_AVAILABLE in the following failure scenarios during:

- Sx Establishment Request and Sx Modification request message processing
  - Bearer stream creation failure for Sxa
  - TEP row add failure for Sxa
  - Local local_gtpu_endpt address mismatch or unavailable
  - Above failure scenarios for N4 visited call
- PFCP_IE_QGR_INFO IE processing memory failures
- NAT rulebase change or policy change cases and failure due to
  - FW-and-NAT policy initilization failure during call setup or rulebase change
  - Invalid clp destination context
  - Memory allocation failure
- Sx Establishment or Sx Modification message processing – local GTPU TEID allocation failure

Previous Behavior: UP sent the error cause PFCP_CAUSE_REQUEST_REJECTED for the above failure scenarios.

New Behavior: UP sends the error cause PFCP_CAUSE_NO_RESOURCE_AVAILABLE instead of PFCP_CAUSE_REQUEST_REJECTED for the above failure scenarios.

Route Map Configuration—CSCwf54987

Revision History

Revision Details	Release
First introduced.	21.28.6 21.28.m10

Behavior Change

CUPS supports a new CLI command to add a route-map under VPNv6 address-family.

Previous Behavior: CUPS did not support the capability to add a route-map under VPNv6 address-family.

New Behavior: CUPS supports the **route-map** option CLI command under the BGP Address-Family Configuration mode to apply a route-map.

To apply the route-map to a neighbor, use the following configuration:

```
configure
  context context_name
    router bgp as_number
      address-family vpnv6
        neighbor route-map map_name { in | out }
        end
```

NOTES:

- **address-family vpnv6**: Configure the IPv6 VPN address family configuration parameters for BGP router.
- **neighbor route-map** *map_name* **{ in | out }**: Specify the route map to apply to a neighbor. *map_name* must be the name of an existing route-map in the current context.
  - **in**: Indicates that the route map applies to incoming advertisements.
  - **out**: Indicates that the route map applies to outgoing advertisements.

Rule Match after TCP Teardown Initiation

Revision History

Revision Details	Release
First introduced.	21.28.0

Feature Changes

After TCP teardown is initiated, there is a change in rule match and charging of the received data packets.

Previous Behavior: Data packets received during or after TCP teardown initiation were matched to the previous rule that was matched for packets in that direction.

New Behavior: Data packets received during or after TCP teardown initiation will be rule matched. Depending on the rule configuration, the packets may match a different rule than that was matched for the previous packet in that direction.

Customer Impact: Difference in rule match for packets received during or after TCP teardown initiation. To be able to match the data packets correctly to the L7 rule, the TCP flag-based and packet length-based L4 rule will have to be configured to exclusively match the TCP control packets.

S8HR LI TCP Connection Timeout

Revision History

Revision Details	Release
First introduced.	21.28.7

Feature Description

CUPS LI supports the TCP connection timeout feature for an S8 Home Routing (S8HR) roaming user in the S-GW service.

For more information, contact your Cisco account representative.

Security Enhancement

Revision History

Revision Details	Release
First introduced.	21.28.1

Feature Description

During upgrade or downgrade, it is recommended to use the compatible configuration files to avoid lockout. The configuration files saved from a new trusted build will not work on older builds (trusted or regular) and new regular builds.

Customer Impact: Possible impact during upgrade or downgrade activities.

Session Checkpoint Compression Algorithms

Revision History

Revision Details	Release
First introduced.	21.28.0

Feature Changes

The RCM supports a combination of both ZLib and LZ4 compression algorithms. For M:N redundancy model, ZLib is the default algorithm that is used to compress the session checkpoint information from UPF to RCM.

In the SRP-based model, the user can choose either of the two compression algorithms for session checkpointing.

Important For data compression and decompression to work, both active and standby UPs must be configured with the same algorithm.

Command Changes

Use the **checkpoint session compression lz4** CLI command in RCM configuration mode to enable the use of LZ4 compression algorithm. You can also revert the compression algorithm to zlib using the **checkpoint session compression zlib** CLI command.

The following command sequence enables the use of LZ4 compression:

```
configure
  context context_name
    redundancy-configuration-module rcm_name
      checkpoint session compression lz4
      end
```

NOTES:

- **checkpoint session compression**: Enables compression of checkpointed session information.
- **checkpoint session compression lz4**: Compresses the checkpointed session information using algorithm - lz4.

For detailed configuration steps, see the *Configuring LZ4 Compression Algorithm* section in the *UPC CUPS User Plane Administration Guide*.

SNMP Trap for Keepalived Status Change Update Failure—CSCwc10141

Revision History

Revision Details	Release
First introduced.	21.28.0

Feature Changes

Previous Behavior: The SNMP trap for status change update failure was not supported by the keepalived pod.

New Behavior: The keepalived pod generates and raises the **RCMControllerStateUpdateFailure** SNMP trap when RCM status change request through HTTP POST from RCM keepalived to RCM controller fails.

Customer Impact: The new SNMP trap eases diagnosis of issues in the keepalived pod.

SNMP Traps to Debug Sx Endpoints Misconfiguration—CSCwf04861

Revision History

Revision Details	Release
First introduced.	21.28.m14

Behavior Change

Previous Behavior: If there was any invalid monitor configuration in Sx endpoints, SNMP traps do not get triggered.

New Behavior: UPF supports the following two SNMP traps to detect misconfiguration in Sx endpoints:

- **NotAllSxMonitorsUp**—This SNMP trap is sent whenever Sx monitor is configured.

  The following conditions apply at the time of Sx monitor installation:

  1. If an Sx peer is not configured or misconfigured in the Control Group (CG), the generated trap remains as NotAllSxMonitorsUp.
  2. If an Sx peer is configured in the CG and Sx association is Down, the generated trap remains as NotAllSxMonitorsUp.

3. If an Sx peer is configured in the CG and Sx association is Up, then configured Sx monitors are checked. If any Sx monitor has the NOT STATUS_UP status, the trap generated remains as NotAllSxMonitorsUp. Otherwise, the AllSxMonitorsUp trap is generated.
4. If an Sx peer is not configured or misconfigured in the CG but the peer is subsequently configured in CG, then the preceding points 2 and 3 will apply.

- **AllSxMonitorsUp**—When an Sx monitor sends the Sx association status (up or down), the configured Sx monitors are reviewed. If an Sx monitor has the NOT STATUS_UP status, the NotAllSxMonitorsUp trap is generated. Otherwise, the AllSxMonitorsUp trap is generated.

Customer Impact: The new SNMP traps allow easy debugging.

Spoofing Detection Support

Revision History

Revision Details	Release
First introduced.	21.28.m1

Feature Changes

The X-Header Enrichment feature appends headers to HTTP or WSP GET and POST request packets, and HTTP Response packets. This feature is used by end applications for mobile advertisement insertion (MSISDN, IMSI, IP address, and so on).

This release supports spoofing detection in X-header fields using a configurable CLI. The **delete-existing** keyword option is added under the **xheader-format** command to enable spoofing detection.

For more information, see the *X-Header Insertion and Encryption* chapter in the *Ultra Packet Core CUPS Control Plane Administration Guide* and *Ultra Packet Core CUPS User Plane Administration Guide*.

Command Changes

The **delete-existing** option is added to the **insert** command in the X-Header Format Configuration mode.

The **delete-existing** option enables spoofing detection in X-header fields. The X-header field configured with this keyword will be removed from the HTTP header if it already exists, and only the gateway inserted field will remain. By default, anti-spoofing is disabled. and if required, should be enabled at a field level.

To configure an X-header format, use the following configuration:

```
configure
  active-charging service ecs_service_name
    xheader-format xheader_format_name
      insert xheader_field_name string-constant xheader_field_value | variable
{ bearer { 3gpp { apn | charging-characteristics | charging-id | imei |
imsi | qos | rat-type | s-mcc-mnc | sgsn-address } | acr | customer-id |
 ggsn-address | mdn | msisdn-no-cc | radius-string |
```

```
radius-calling-station-id | session-id | sn-rulebase |
subscriber-ip-address | username } [ encrypt ] [ delete-existing ] | http
 { host | url } }
       end
```

Sxa Tunnel Retained till DSR on SAEGW—CSCwe80030

Revision History

Revision Details	Release
First introduced.	21.28.m14 21.28.6 21.28.m10 21.25.17

Feature Changes

During X2/S1 handover with S-GW relocation, support is added to retain the Sxa tunnel endpoints of source SGW-U. This tunnel retention enables uplink data to flow over SGW-U until the path switches. The new CLI command **sxa-tunnel-del-at-dsr-on-sgw-change** helps SAEGW/PGW-C to retain the Sxa tunnel of source SGW-U until a Delete Session Request (DSR) is sent from MME.

Previous Behavior: During X2/S1-based handover with S-GW relocation, the SAEGW/PGW-C sent Sx Session Modification Request to SAEGW/PGW-U to remove traffic endpoints of source S-GW (Sxa). Due to this, Sxa traffic endpoints were deleted.

New Behavior: During X2/S1-based handover with S-GW relocation, when you configure the **sxa-tunnel-del-at-dsr-on-sgw-change** CLI, it helps the SAEGW/PGW-U to retain Sxa traffic endpoints of source S-GW until DSR is received.

Customer Impact: Data passed over source SGW-U during X2/S1 based handover will have GTP-U error indication.

Command Changes

To enable or disable Sxa tunnel deletion, use the following configuration:

```
configure
  context context_name
    saegw-service service_name
      [ no ] sxa-tunnel-del-at-dsr-on-sgw-change
      end
```

NOTES:

- **sxa-tunnel-del-at-dsr-on-sgw-change**: Enable Sxa tunnel deletion at DSR during X2/S1-based handover with S-GW relocation.

- **no sxa-tunnel-del-at-dsr-on-sgw-change**: Disable Sxa tunnel deletion at DSR during X2/S1-based handover with S-GW relocation.
- By default, the configuration is disabled.
- The configuration is applied to all current and new sessions.

SU URR Association with PDRs—CSCwk49621

Revision History

Revision Details	Release
First introduced.	21.28.mx

Behavior Change

In CUPS, when the OCS server sends failure response during CCR-I, then Control Plane moves to Server Unreachable state. In this state, if the Volume Quota is exhausted, the Control Plane doesn't renew the Quota for predefined rules.

Previous Behavior: When OCS server sends failure response during CCR-I, the Sx Establishment Request message towards the User Plane does not associate SU URR with PDRs or other GY URRs.

New Behavior: When OCS server sends failure response during CCR-I, the Sx Establishment Request message towards the User Plane associates SU URR with PDRs or other GY URRs.

Customer Impact: The Control Plane sends a Server Unreachable URR associated with PDR in existing message while sending the Sx Session Establishment Request. There is no performance or memory impact.

TCP Hardening between RCM and UPF—CSCwc25287

Revision History

Table 3: Revision History

Revision Details	Release
First introduced.	21.28.m14

Feature Changes

TCP hardening between RCM and UPF is supported with this release. As part of RCM checkpoint manager hardening, UPF supports the heartbeat mechanism between UP sessmgr and RCM checkpoint manager. This feature provides CLI support to enable or disable TCP hardening between RCM and UPF.

Previous Behavior: TCP hardening was not supported and not configurable.

New Behavior: Use the following CLI commands to configure the heartbeat mechanism:

- To enable or disable sending the heartbeat from UP sessmgr to RCM checkpointmgr, use the following command in the Context > Redundancy-Configuration-Module mode. This command is disabled by default.

  up-sm-heartbeat { disable | enable }

  To verify the configuration, use the **show config context** *context_name* command.

- To enable or disable heartbeat from RCM to active or standby UPF, use the following command in RCM Ops-center. This command is disabled by default.

  k8 smf profile rcm-config-ep enable-up-heartbeat { false | true }

Customer Impact: The heartbeat mechanism addresses the intermittent issues of TCP connectivity with UP sessmgr and RCM checkpoint managers.

For more information, refer to the *UCC 5G RCM Configuration and Administration Guide*.

TEID Collision Handling during MOCN

Revision History

Revision Details	Release
First introduced.	21.28.m7

Feature Description

Tunnel Endpoint Identifier (TEID) collisions support the Multiple Operator Core Network (MOCN) scenario on CUPS. During TEID collision, P-GW or GGSN allocates a TEID to a home subscriber. In case of a stale session, in an S-GW or SGSN, the same TEID that is allocated by P-GW or GGSN, is allocated to a roaming subscriber.

To eliminate this scenario, CUPS supports TEID Collision with User Location Information (ULI) change to reject a request by configuring P-GW and GGSN when TEID collision occurs. This feature allows comparison only with Mobile Country Code (MCC) instead of comparison with MCC and Mobile Network Code (MNC). This feature supports the MOCN scenario on GGSN, P-GW, and SAEGW.

For more information, refer to the *TEID Collision Handling during MOCN* chapter in the *Ultra Packet Core CUPS Control Plane Administration Guide*.

Updated TCP Heartbeat Timestamps—CSCwe60240

Revision History

Revision Details	Release
First introduced.	21.28.m14

Behavior Change

The output of the **rcm show-statistics checkpointmgr-endpointstats** RCM Ops Center CLI command related to the TCP heartbeat feature is updated in this release.

Previous Behavior: If TCP heartbeat was disabled, the "Lasthbrcvd" and "Lasthbsend" fields printed junk values of the last sent and last received timestamps.

New Behavior: The "Lasthbsend" field is printed only when TCP heartbeat is enabled on RCM.

Customer Impact: When the TCP heartbeat feature is enabled, the output of the **rcm show-statistics checkpointmgr-endpointstats** CLI command is updated with the correct timestamp.

URR Volume Quota Calculation—CSCwd61752

Revision History

Revision Details	Release
First introduced.	21.28.m7 21.26.h5

Behavior Change

Previous Behavior: UPF recalculated the URR volume quota values as per the usage after UP recovery.

New Behavior: The volume quota values provided by OCS must be the same after UP recovery.

Warning Message for Traffic Data Checking Options on CP—CSCwe61210

Revision History

Revision Details	Release
First introduced.	21.28.m10

Feature Changes

Previous Behavior: The **show subscribers [tx-data | rx-data | idle-time]** and **clear subscribers [tx-data | rx-data | idle-time]** CLI commands were executed on both user plane and control plane.

New Behavior: The **show subscribers [tx-data | rx-data | idle-time]** and **clear subscribers [tx-data | rx-data | idle-time]** commands are specific to user plane. If these commands are executed on control plane, they will not be processed and CP displays a warning message.

The following is a sample warning message that displays when you configure the **show subscribers rx-data** or **clear subscribers rx-data** command:

Warning: rx-data option not relevant on CUPS-CP platform. Please use this option on the UP side

UCS C220 M6 Server Support for VPC-DI CP

Revision History

Table 4: Revision History

Revision Details	Release
First introduced.	21.28.mh14

Feature Description

The VPC-DI Control Plane supports Cisco UCS C220 M6 server on the RHOSP.

For more information about the Hardware and Software configurations, refer the *VPC-DI System Administration Guide*.