Ultra Cloud Core 5G User Plane Function, Release 2021.02 - Configuration and Administration Guide

First Published: 2021-05-31

Last Modified: 2021-10-29

C O N T E N T S

CHAPTER 28

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 5G UPF product deployed in a 5G 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 preface describes the *5G User Plane Function Guide*, how it is organized and its document conventions.

This guide describes the Cisco User Plane Function (UPF) and includes infrastructure and interfaces, feature descriptions, specification compliance, session flows, configuration instructions, and CLI commands for monitoring and troubleshooting the system.

Conventions Used

The following tables describe the conventions used throughout this documentation.

Notice Type	Description
Information Note	Provides information about important features or instructions.

Notice Type	Description
Caution	Alerts you of potential damage to a program, device, or system.
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**

5G Architecture

Feature Summary and Revision History

Summary Data

Table 1: Summary Data

Applicable Product(s) or FunctionalArea	• PCF • SMF • UPF
Applicable Platform(s)	SMI
Feature Default Setting	Not Applicable
Related Documentation	Not Applicable

Revision History

Table 2: Revision History

Revision Details	Release
First introduced.	Pre-2020.02.0

Overview

The Ultra Cloud Core is Cisco's solution supporting 3GPP's standards for 5G new radio (NR) standalone (SA) mode. These standards define various network functions (NFs) based on the separation of control plane (CP) and user plane (UP) (e.g. CUPS) functionality for increased network performance and capabilities.

Control Plane Network Functions

The CP-related NFs that comprise the Ultra Cloud Core are based on a common architecture designed around the following tenants:

- Cloud-scale — Fully virtualized for simplicity, speed, and flexibility

- Automation and orchestration — Optimized operations, service creation, and infrastructure

- Security — Multiple layers of security across the deployment stack from the infrastructure through the NF applications

- API exposure — Open and extensive for greater visibility, control, and service enablement

- Access agnostic — Support for heterogeneous network types (e.g. 5G, 4G, 3G, Wi-Fi, etc.)

These CP NFs are each designed as containerized applications (e.g. microservices) for deployment via the Subscriber Microservices Infrastructure (SMI).

The SMI defines the common application layers for functional aspects of the NF such as life cycle management (LCM), operations and management (OAM), and packaging.

Figure 1: Ultra Cloud Core CP Architectural Components
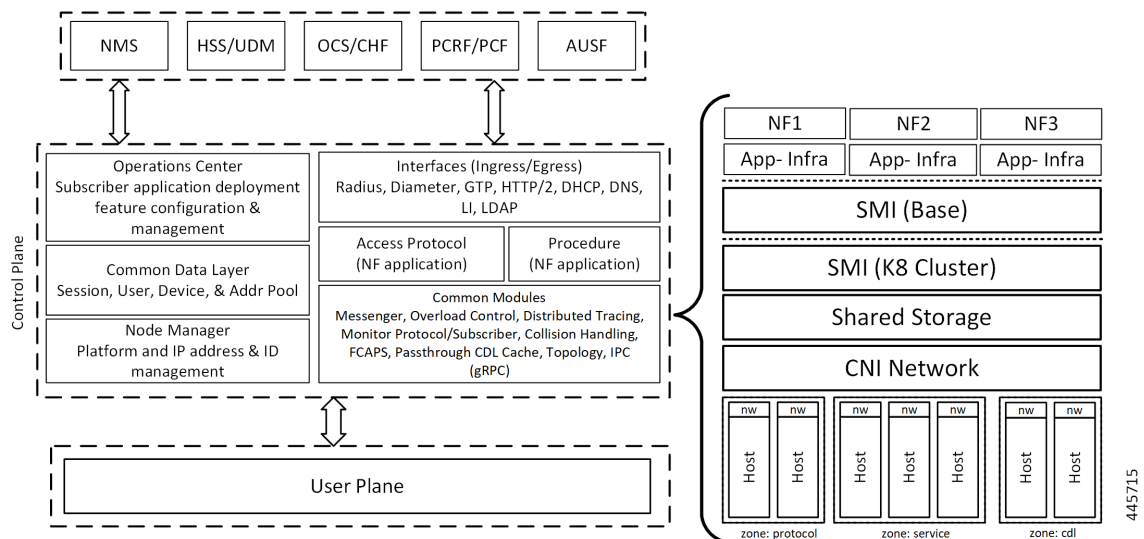
User Plane Network Function

The 5G UP NF within the Ultra Cloud Core is the User Plane Function. Unlike the CP-related NFs, the 5G UPF leverages the same Vector Packet Processing (VPP) technology currently in use by the user plane

component within Cisco's 4G CUPS architecture. This commonality ensures the delivery of a consistent set of capabilities between 4G and 5G such as:

- Ultra-fast packet forwarding

- Extensive integrated IP services such as Subscriber Firewall, Tethering, Deep-Packet Inspection (DPI), Internet Content Adaption Protocol (ICAP), Application Detection and Control (ADC), and header enrichment (HE)

- Integrated third-party applications for traffic and TCP optimization

For more information on UPF, refer to *Ultra Cloud Core 5G UPF Configuration and Administration Guide*.

Subscriber Microservices Infrastructure Architecture

The Ultra Cloud Core Subscriber Microservices Infrastructure (SMI) is a layered stack of cloud technologies that enable the rapid deployment of, and seamless life-cycle operations for microservices-based applications.

The SMI stack consists of the following:

- SMI Cluster Manager—Creates the Kubernetes (K8s) cluster, creates the software repository, and provides ongoing LCM for the cluster including deployment, upgrades, and expansion.

- Kubernetes Management—Includes the K8s master and etcd functions, which provide LCM for the NF applications deployed in the cluster. This component also provides cluster health monitoring and resources scheduling.

- Common Execution Environment (CEE)—Provides common utilities and OAM functionalities for Cisco cloud native NFs and applications, including licensing and entitlement functions, configuration management, telemetry and alarm visualization, logging management, and troubleshooting utilities. Additionally, it provides consistent interaction and experience for all customer touch points and integration points in relation to these tools and deployed applications.

- Common Data Layer (CDL)—Provides a high performance, low latency, stateful data store, designed specifically for 5G and subscriber applications. This next generation data store offers HA in local or geo-redundant deployments.

- Service Mesh—Provides sophisticated message routing between application containers, enabling managed interconnectivity, additional security, and the ability to deploy new code and new configurations in low risk manner.

- NB Streaming—Provides Northbound Data Streaming service for billing and charging systems.

- NF/Application Worker nodes—The containers that comprise an NF application pod.

- NF/Application Endpoints (EPs)—The NF's/application's interfaces to other entities on the network.

- Application Programming Interfaces (APIs)—SMI provides various APIs for deployment, configuration, and management automation.

The following figure depicts how these components interconnect to comprise a microservice-based NF/application.

Figure 2: SMI Components
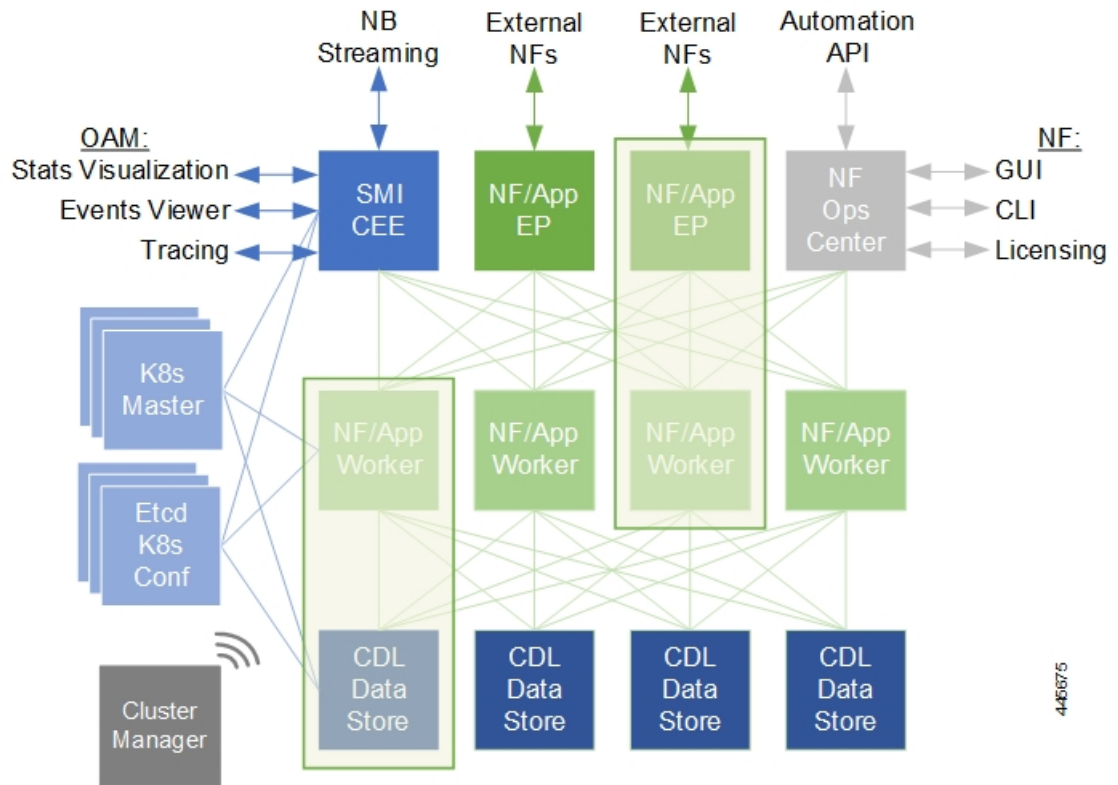
For more information on SMI components, refer to the Ultra Cloud Core Subscriber Microservices Infrastructure documentation— *Deployment Guide > Overview* chapter.

Control Plane Network Function Architecture

CP NFs are designed around a three-tiered architecture that take advantage of the stateful/stateless capabilities afforded within cloud native environments.

The architectural tiers are as follows:

- Protocol Load Balancer Services — These are stateless microservices that are primarily responsible for dynamic discovery of application containers as well as for protocol proxy and/or termination. These include traditional 3GPP protocols and new protocols introduced with 5G.

- Applications Services — Responsible for implementing the core application/business logic, these are the stateless services that render the actual application based on the received information. This layer may contain varying degrees of microservice granularity. Application services are stateless.

- State management services — Enable stateless application services by providing a common data layer (CDL) to store/cache state information (e.g. session and subscriber data). This layer supports various data storage technologies from in-memory caches to full-fledge databases.

Figure 3: Control Plan Network Function Tiered Architecture
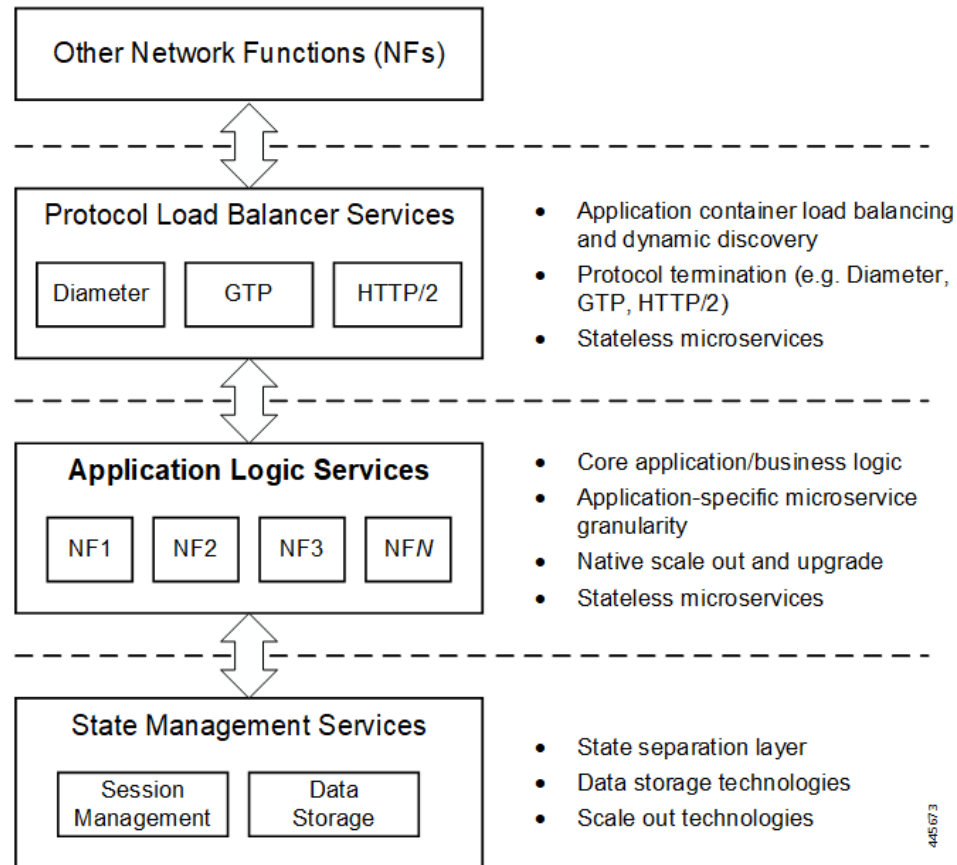
The three-tiered architecture on which Cisco's CP NFs are designed fully support the 5G core (5GC) Service-based Architecture (SBA) defined by 3GPP. These NFs communicate with each other and with third-party NFs over the Service-based Interface (SBI) using HTTP/2 over TCP as defined by 3GPP.

Figure 4: Cisco CP NF Service-based Architecture Support
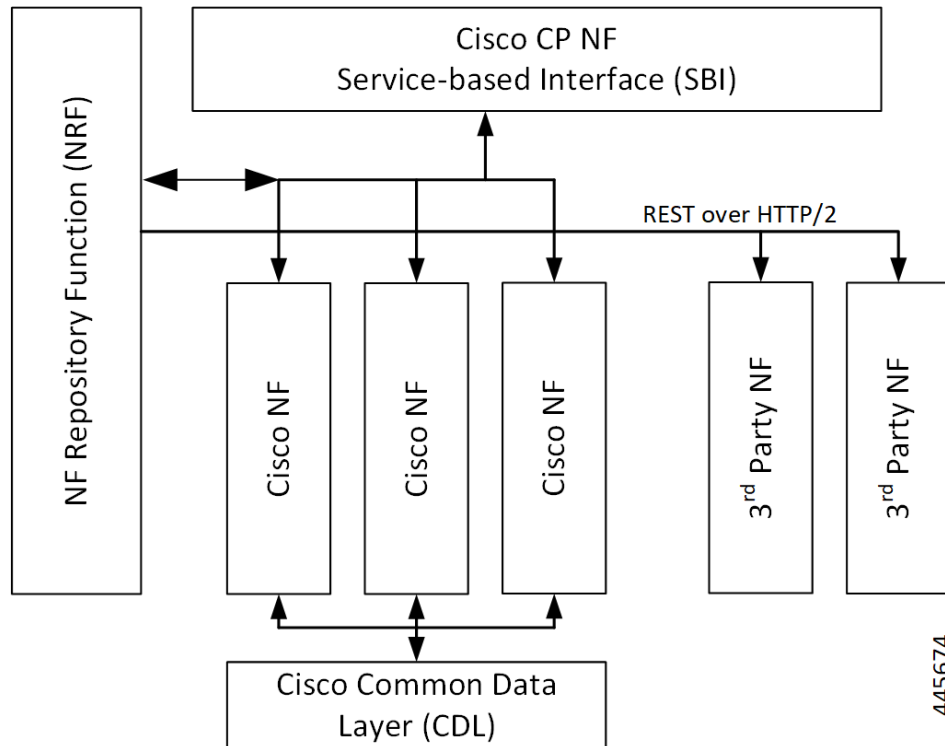
For more information on the Cisco network functions, refer to the corresponding network function documentation.

5G-UPF Overview

This chapter covers the following topics:

Feature Summary and Revision History

Summary Data

Table 3: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI
Feature Default Setting	Disabled – License Required
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Table 4: Revision History

Revision Details	Release
First introduced.	2020.02.0

Product Description

The User Plane Function (UPF) is one of the network functions (NFs) of the 5G core network (5GC). The UPF is responsible for packet routing and forwarding, packet inspection, QoS handling, and external PDU session for interconnecting Data Network (DN), in the 5G architecture.

UPF is a distinct Virtual Network Function (VNF) that offers a high-performance forwarding engine for the user traffic. Using Vector Packet Processing (VPP) technology, the UPF achieves ultra-fast packet forwarding while retaining compatibility with all the user plane functionality. For instance, Shallow Packet Inspection(SPI)/Deep Packet Inspection (DPI), traffic optimization, and inline services (NAT, Firewall, DNS snooping, and so on). UPF is currently designed to offer Integrated Deep Packet Based Inspection (DPI) Services.

A single instance of UPF provides some or all the following functionalities:

- Anchor point for Intra-RAT and Inter-RAT mobility (when applicable).

- External PDU session point of interconnect to Data Network.

- Packet routing and forwarding.

- Packet inspection. For example, Application detection that is based on the service data flow template and the optional PFDs received from the SMF in addition.

- User Plane part of policy rule enforcement. For example, Gating, Redirection, Traffic steering.

- Lawful intercept (UP collection).

- Traffic usage reporting.

- QoS handling for User Plane. For example, Uplink (UL) and Downlink (DL) rate enforcement, Reflective QoS marking in DL, and so on.

- Uplink Traffic verification (SDF to QoS Flow mapping).

- Transport level packet marking in the Uplink and Downlink.

- Downlink packet buffering and Downlink Data Notification triggering.

- Sending and forwarding of one or more "End Marker" to the source NG-RAN node.

The UPF also provides support for an enterprise mobile virtual network operator (MVNO) model, which enables a mobile network operator (MNO) to perform secondary authentication for the leased MVNO subscribers.

Use Cases and Features

Configuration and Deployment Requirement for UPF

With 5G deployment, interoperability is required between Cisco UPF with non-Cisco SMF, and Cisco SMF with non-Cisco UPF. Also, decoupling of configuration-related messaging between SMF and UPF has the following benefits:

- Alignment with 3GPP standards for configuration bifurcation between User Plane and Control Plane.

- Reduced complexity for configuration management on SMF.

- Simplicity and efficiency for the configuration and change management for User Plane related configuration, as it does not require SMF to manage and distribute the configuration.

- Can be enhanced to achieve interworking between non-Cisco SMF and UPFs.

The Cisco UPF supports 3GPP-specified attributes on the N4 interface. In the current architecture, only UPF associates with the SMF.

The following features are related to this use case:

Anchor Point for Intra-RAT and Inter-RAT Mobility

The UPF is the anchor point between the mobile infrastructure and the Data Network (DN). That is, the encapsulation and decapsulation of GPRS Tunneling Protocol for the User Plane (GTP-U). Intra-RAT mobility like Xn handover and inter-RAT mobility like 4G to 5G and 5G to 4G handover are supported for this use case.

The feature is related to this use case.

External PDU Session Point of Interconnect to Data Network

The UPF acts as an external PDU session point of interconnect to Data Network and supports N3, N4, and N6 interfaces. The PDU layer corresponds to the PDU that is transported between the UE and the PDN during a PDU session. The PDU session can be of type IPv4 or IPv6 for transporting IP packets. The GPRS tunneling protocol for the user plane (GTP-U) supports multiplexing of the traffic from different PDU sessions by tunneling user data over the N3 interface (between a 5G access node and the UPF) in the core network. The GTP encapsulates all end-user PDUs and provides encapsulation per-PDU session. This layer also transports the marking associated with the QoS flow. The 5G encapsulation layer supports multiplexing the traffic from different PDU sessions over the N9 interface (an interface between different UPFs). It provides encapsulation per PDU session and carries the marking associated with the QoS flows.

The following features are related to this use case:

Packet Inspection

The Cisco UPF performs L3/L4 and L7 inspection for the user traffic that is received. L3/L4 inspection involves IP-address/port matching and Deep Packet Inspection involves matching of L7 header fields.

The Deep Packet Inspection and Inline Services, on page 89 feature is related to this use case.

User Plane Part of Policy Rule Enforcement

Cisco UPF provides different enforcement mechanisms based on policy received from the SMF. The UPF is the boundary between the Access and IP domains and is the ideal location to implement policy-based enforcement. The pcc-rules provided by the PCF and the pre-defined rules on the SMF are uploaded over the N4 interface and installed on the UPF on a per-DNN basis. This allows for dynamic policy changes that enable differentiated charging and QoS enforcement.

- Dynamic and Static PCC Rules, on page 157

- Voice over New Radio, on page 289

Lawful Intercept

Lawful Interception (LI) enables a LEA to perform electronic surveillance on an individual (a target) as authorized by a judicial or administrative order. To facilitate the lawful intercept process, certain legislation and regulations require service providers and Internet service providers to implement their networks to explicitly support authorized electronic surveillance. Actions taken by the service providers include: provisioning the target identity in the network to enable isolation of target communications (separating it from other users' communications), duplicating the communications for the purpose of sending the copy to the LEA, and delivering the Interception Product to the LEA.

For information about the support of Lawful Intercept by UPF, contact your Cisco Account representative.

Traffic Usage Reporting (Charging)

The usage measurement and reporting function in UPF is controlled by the SMF. The SMF controls these functions by:

- Creating the necessary PDRs to represent the service data flow, application, bearer or session (if they are not existing already).

- Creating the URRs for each Charging Key and combination of Charging Key and Service ID. Also, creating URRs for a combination of Charging Key, Sponsor ID, and Application Service Provider ID.

  Please note that, for static rules, the UPF creates the URR ID. The URR ID is created based on the online/offline and Content ID+Service ID combination that is configured on UPF.

- Associating the URRs to the relevant PDRs defined for the PFCP session, for usage reporting at SDF, Session or Application level.

- For online charging, the SMF provisions Volume and Time quota, if it receives it from the Online Charging Server (OCS).

The Charging Support, on page 49 feature is related to this use case.

QoS Handling for User Plane

The 5G QoS model allow classification and differentiation of specific services, based on subscription-related and invocation-related priority mechanisms. These mechanisms provide abilities such as invoking, modifying, maintaining, and releasing QoS Flows with priority, and delivering QoS Flow packets according to the QoS characteristics under network congestion conditions.

The Dynamic and Static PCC Rules, on page 157 feature is related to this use case.

Downlink Packet Buffering and Data Notification Triggering

A Buffering Action Rule (BAR) provides instructions to control the buffering behavior of the UPF. The BAR controls the buffering behavior for all Forwarding Action Rules (FARs) of the Packet Forwarding Control Protocol (PFCP) session. This control is applicable when the PFCP session is set with an Apply Action parameter, which requests packets to be buffered and associated with the respective BAR.

The Idle Mode Buffering and Paging, on page 179 feature is related to this use case.

Forwarding End Markers to the Source NG-RAN Node

At the time of the handover procedure, the PDU session for the UE – which comprises of UPF node – acts as a PDU session anchor and an intermediate UPF terminating N3 reference point. The SMF sends an N4 Session Modification Request message with the new AN Tunnel Info of NG-RAN to specify the UPF to switch to the N3 paths. In addition, the SMF also specifies the UPF to send the End Marker packets on the old N3 user plane path. After the UPF receives the indication, the End Markers are constructed and sent to each N3 GTP-U tunnel toward the source NG-RAN, after sending the last PDU on the old path.

The N4 Session Management, Node Level, and Reporting Procedures, on page 235 feature is related to this use case.

MVNO Support

The UPF provides support for an enterprise MVNO model. A mobile network operator can perform secondary authentication for the leased MVNO subscribers and also support any additional features related to the AAA server.

The following features are related to this use case:

- APN ACL Support, on page 35

  A configurable mechanism to apply traffic classification and policy enforcement on selective subscriber sessions.

- Dynamic and Static PCC Rules, on page 157

  Increase in maximum number of groups per bandwidth policy.

- Virtual Routing and Forwarding, on page 281

  Support for Overlapping IP Pools and IP Pool chunks.

Deployment Architecture and Interfaces

Cisco UPF is part of the 5GC network functions portfolio (AMF/SMF/NRF/PCF/NSSF/UPF) with a common Mobile Core Platform architecture.

UPF Architecture

The User Plane Function (UPF) is a fundamental component of a 3GPP 5G core infrastructure system architecture. The UPF represents the data plane evolution of a Control and User Plane Separation (CUPS) strategy, first introduced as an extension to existing Evolved Packet Cores (EPCs) by the 3GPP in Release 14 specifications. The CUPS decoupless Packet Gateway (P-GW) Control and User Plane functions, enabling the data forwarding component (PGW-U) to be decentralized. This allows packet processing and traffic aggregation to be performed closer to the network edge, increasing bandwidth efficiencies while reducing network load. The P-GW handling signaling traffic (PGW-C) remains in the core, northbound of the Mobility Management Entity (MME).

The primary goal of CUPS is to support 5G New Radio (NR) implementations enabling early IoT applications and higher data rates. Committing to a complete implementation of CUPS is a complex proposition as it only provides a subset of advantages to the operator adopting a 5G User Plane Function (5G-UPF), offering network slicing. Deployed as a Virtual Machine (VM), the User Plane Function delivers the packet processing foundation for Service-Based Architectures (SBAs).

The UPF identifies User Plane traffic flow that is based on information received from the SMF over the N4 reference point. The N4 interface employs the Packet Forwarding Control Protocol (PFCP), which is defined in the 3GPP technical specification 29.244 for use on Sx/N4 reference points in support of CUPS. The PFCP is similar to OpenFlow but can be limited to only the functionality that is required to support mobile networks. The PFCP sessions, which are established with the UPF, define how packets are identified (Packet Detection Rule / PDR), forwarded (Forwarding Action Rules / FARs), processed (Buffering Action Rules / BARs), marked (QoS Enforcement Rules / QERs) and reported (Usage Reporting Rules / URRs).

UPF Deployment Architecture

The following diagram illustrates, at a high-level, the deployment architecture of UPF along with other NFs.

Virtualized Packet Core—Single Instance (VPC-SI)

VPC-SI consolidates the operations of physical Cisco ASR 5500 chassis running StarOS into a single Virtual Machine (VM) able to run on commercial off-the-shelf (COTS) servers. VPC-SI can be used as a stand-alone single VM within an enterprise, remote site, or customer data center. Alternatively, VPC-SI can be integrated as a part of a larger service provider orchestration solution.

VPC-SI only interacts with supported hypervisors KVM (Kernel-based Virtual Machine) and VMware ESXi. It has little or no knowledge of physical devices.

The UPF functions as user plane node in 5G-based VNF deployments. UPF is deployed as a VNFC running a single, stand-alone instance of the StarOS. Multiple UPF VNFCs can be deployed for scalability based on your deployment requirements.

Hypervisor Requirements

VPC-SI has been qualified to run under the following hypervisors:

- Kernel-based Virtual Machine (KVM) - QEMU emulator 2.0. The VPC-SI StarOS installation build includes a libvirt XML template and ssi_install.sh for VM creation under Ubuntu Server14.04.

- KVM - Red Hat Enterprise Linux 7.2: The VPC-SI StarOS installation build includes an install script called qvpc-si_install.sh.

- VMware ESXi 6.7: The VPC-SI StarOS installation build includes OVF (Open Virtualization Format) and OVA (Open Virtual Application) templates for VM creation via the ESXi GUI.

vNIC Options

The supported vNIC options include:

- VMXNET3—Paravirtual NIC for VMware

- VIRTIO—Paravirtual NIC for KMV

- ixgbe—Intel 10 Gigabit NIC virtual function

- enic—Cisco UCS NIC

- SR-IOV—Single-root input/output virtualization

The SR-IOV specification provides a mechanism by which a single root function (for example, a single Ethernet port) can appear to be multiple separate physical devices. Intel 82599 10G is an SR-IOV capable device and can be configured (usually by the Hypervisor) to appear in the PCI configuration space as multiple functions (PFs and VFs). The virtual functions (VFs) can be assigned to Nova VMs, causing traffic from the VMs to bypass the Hypervisor and go directly to the fabric interconnect. This feature increases traffic throughput to the VM and reduces CPU load on the UCS Servers.

Capacity, CEPS and Throughput

Sizing a VPC-SI instance requires modeling of the expected call model.

Many service types require more resources than others. Packet size, throughput per session, CEPS (Call Events per Second) rate, IPSec usage (site-to-site, subscriber, LI), contention with other VMs, and the underlying hardware type (CPU speed, number of vCPUs) will further limit the effective number of maximum subscribers. Qualification of a call model on equivalent hardware and hypervisor configuration is required.

Sample VPP Configuration

For 5G-UPF, the FORWARDER_TYPE is "vpp".

The following is a sample output of VPP configuration.

```
show cloud configuration
Thursday January 30 12:18:10 UTC 2020
Card 1:
  Config Disk Params:
  -----------------------
FORWARDER_TYPE=vpp
VNFM_INTERFACE=MAC:fa:11:3e:22:d8:33
MGMT_INTERFACE=MAC:fa:11:3e:44:af:9e
VNFM_IPV4_ENABLE=true
VNFM_IPV4_DHCP_ENABLE=true
SERVICE1_INTERFACE=MAC:fa:11:3e:11:9d:23
SERVICE2_INTERFACE=MAC:fa:11:3e:99:ec:7b
VPP_CPU_WORKER_CNT=8
VPP_DPDK_TX_QUEUES=9
VPP_DPDK_RX_QUEUES=8
CHASSIS_ID=xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx

  Local Params:
  -----------------------
    No local param file available
```

Note   For additional information about VPC-SI build components, boot parameters, configuring VPC-SI boot parameters, VM configuration, vCPU and vRAM options, VPP configuration parameters, and so on, refer the *VPC-SI System Administration Guide*.

UPF Deployment with VPC-SI

For additional information on VPC-SI, supported operating system and hypervisor packages, platform configurations, software download and installation, as well as UPF deployment, contact your Cisco Account representative.

For information on Release Package, refer the corresponding Release Notes included with the build.

UPF Deployment with SMI Cluster Manager

The Ultra Cloud Core Subscriber Microservices Infrastructure (SMI) provides a run time environment for deploying and managing Cisco's cloud native network functions (cNFs), also referred to as applications.

It is built around open source projects like Kubernetes (K8s), Docker, Helm, etcd, confd, and gRPC, and provides a common set of services used by deployed cNFs.

The SMI is a layered stack of cloud technologies that enable the rapid deployment of, and seamless life cycle operations for microservices-based applications.

The SMI stack consists of SMI Cluster Manager that creates the Kubernetes (K8s) cluster and the software repository. The SMI Cluster Manager also provides ongoing Life Cycle Management (LCM) for the cluster including deployment, upgrades, and expansion.

The SMI Cluster Manager leverages the Kernel-based Virtual Machine (KVM)—a virtualization technology—to deploy the User Plane Function (UPF) VMs.

For more information, refer the *UCC SMI Operations Guide*.

Supported Interfaces

This section describes the interfaces supported between the UPF and other network functions in 5GC.

- N3: Interface between the RAN (gNB) and the (initial) UPF; compliant with 3GPP TS 29.281 and 3GPP TS 38.415 (December-2018).

- N4: Interface between the Session Management Function (SMF) and the UPF; compliant with 3GPP TS 29.244 (December-2018).

- N6: Interface between the Data Network (DN) and the UPF; compliant with 3GPP TS 29.561 (December-2018).

License Information

The UPF require specific license(s). Contact your Cisco account representative for more information on how to obtain a license.

Standards Compliance

Cisco UPF complies with the following standards:

- Interface between the Control Plane and the User Plane Nodes: 3GPP TS 29.244 version 15.4.0. (December-2018)

- General Packet Radio System (GPRS) Tunneling Protocol User Plane (GTPv1-U): 3GPP TS 29.281 version 15.5.0 (December-2018).

- NG-RAN; PDU Session User Plane protocol: 3GPP TS 38.415 (December-2018)

- 5G System; Interworking between 5G Network and external Data Networks; Stage 3: 3GPP TS 29.561 (December-2018)

PART I

Features and Functionality

CHAPTER **3**

1:1 Redundancy

Feature Summary and Revision History

Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Revision Details	Release
Support is added for the following functionality: • Zero Accounting Loss in User Plane Function • Early PDU Recovery • Session Prioritization during Recovery • Configuration to change the state of UPF from Pending-Active to Active	2021.02.0
First introduced.	2020.02.0

Feature Description

The 1:1 UPF Redundancy feature, for 5G deployment, supports the detection of a failed User Plane Function (UPF) and seamlessly handles the functions of the failed UPF. Each of the Active UPF has a dedicated Standby UPF. The 1:1 UPF Redundancy architecture is based on the UPF to UPF Interchassis Session Recovery (ICSR) connection.

How it Works

The 5G-UPF deployment leverages the ICSR framework infrastructure for checkpointing and switchover of the UPF node as shown in the following figure. The Active UPF communicates to its dedicated Standby UPF through the Service Redundancy Protocol (SRP) link that is provisioned between the UPFs.

Figure 5: UPF 1:1 Redundancy Using SRP
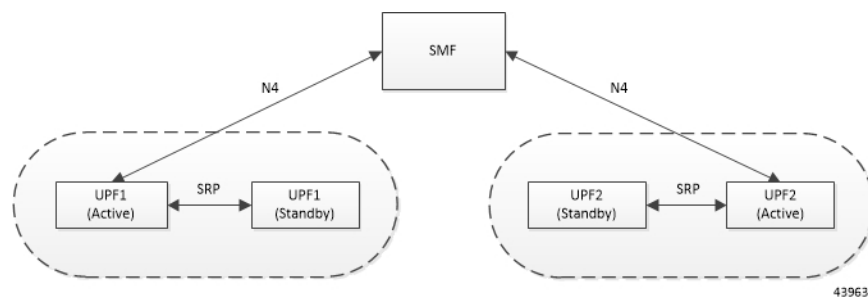
The Session Management Function (SMF) node does not have the Standby UPF information that is available in the UPF group configuration. Therefore, the SMF is not aware of the UPF redundancy configuration and the switchover event among the UPFs.

The Active UPF communicates to the SMF through the N4 interface address configured in the UPF. The Standby UPF takes over the same Sx/N4 address when it transitions to Active during the switchover event. This implies that the Sx/N4 interface is SRP-activated and is in line with the existing configuration method, therefore UPF switchover is transparent to the SMF.

Figure 6: UPF 1:1 Redundancy Switchover

To make redundancy fully compliant, it addresses the following dependencies on the SRP-based ICSR in the 5G environment.

- Configuration Synchronization (or, Replica Configuration on Standby UPF)

- Sx/N4 Association Checkpoint

- Sx/N4 Link Monitoring

Besides the dependencies listed, the UPF implements data collection and checkpoint procedures specific to the UPF node. For example, checkpointing for IP-pool chunks. The UPF integrates these procedures into the existing ICSR checkpointing framework.

Independent Configuration of Standby UPF

After UPF is up with base configuration (for example, services, contexts, interfaces, and so on), the rest of the configuration (for example, ACS and policy-related configuration) is done through Ops-center/Redundancy and Configuration Manager (RCM) POD. This configuration is common for both SMF/UPF policies. For SRP redundancy to work, the Active and Standby UPF has same configuration, except SRP-related configuration with which SRP connections are established between Active and Standby UPF. The RCM configures Active and Standby UPF independently.

BFD Monitor Between Active UP and Standby UP

The Bidirectional Forwarding Detection (BFD) monitors the SRP link between the Active UPF and Standby UPF for a fast failure-detection and switchover. When the Standby UPF detects a BFD failure in this link, it takes over as the Active UPF.

The BFD link can be single-hop or multi-hop.

To configure the BFD monitor, between the Active UP and Standby UP, see *Configuring BFD Monitoring Between Active UPF and Standby UPF*.

Sample Configuration for Multihop BFD Monitoring

Primary UPF:

```
config
  context srp
    bfd-protocol
      bfd multihop-peer 1.1.1.1 interval 50 min_rx 50 multiplier 20
    #exit
    service-redundancy-protocol
      monitor bfd context srp 1.1.1.1 chassis-to-chassis
      peer-ip-address 1.1.1.1
      bind address 1.1.0.1
    #exit
    interface srp
      ip address 1.1.0.1 2.3.4.0
    #exit
    ip route static multihop bfd bfd1 1.1.0.1 1.1.1.1
    ip route 1.1.0.1 2.3.4.0 1.1.0.1 srp
  #exit
end
```

Backup UPF:

```
config
  context srp
    bfd-protocol
      bfd multihop-peer 1.1.0.1 interval 50 min_rx 50 multiplier 20
    #exit
    service-redundancy-protocol
      monitor bfd context srp 1.1.0.1 chassis-to-chassis
      peer-ip-address 1.1.0.1
      bind address 1.1.0.1
    #exit
    interface srp
      ip address 1.1.0.1 255.255.255.0
    #exit
    ip route static multihop bfd bfd1 1.1.1.1 1.1.0.1
    ip route 1.1.0.1 255.255.255.0 1.1.1.1 srp
  #exit
End
```

Router between Primary and Backup UPF:

```
config
  context one
    interface one
      ip address 1.1.0.1 255.255.255.0
    #exit
    interface two
      ip address 1.1.1.1 255.255.255.0
    #exit
  #exit
end
```

Sample Configuration for Single-Hop BFD Monitoring

Primary UPF:

```
config
  context srp
    bfd-protocol
    #exit
    service-redundancy-protocol
      monitor bfd context srp 1.1.0.1 chassis-to-chassis
      peer-ip-address 1.1.0.1
      bind address 1.1.2.1
    #exit
    interface srp
      ip address 1.1.0.1 255.255.255.0
```

```
      bfd interval 50 min_rx 50 multiplier 10
    #exit
    ip route static bfd srp 1.1.2.1
  #exit
end
```

Backup UPF:

```
config
  context srp
    bfd-protocol
    #exit
    service-redundancy-protocol
      monitor bfd context srp 1.1.1.1 chassis-to-chassis
      peer-ip-address 1.1.2.1
      bind address 1.1.3.1
    #exit
    interface srp
      ip address 1.1.2.1 255.255.255.0
      bfd interval 50 min_rx 50 multiplier 10
    #exit
    ip route static bfd srp 1.1.3.1
  #exit
end
```

VPP Monitor

When SRP VPP monitor is configured, the UPF chassis is SRP Active and if the VPP subsystem fails, then SRP initiates switchover to Standby UPF. Currently, VPP health monitoring is limited to heartbeat mechanism between NPUMgr task and VPP process.

To configure the VPP monitor, see *Configuring VPP Monitor on Active UPF and Standby UPF*.

Sx/N4 Association Checkpoint

Whenever an Active UPF initiates an Sx/N4 association to SMF, the Standby UPF checkpoints this data. This maintains the association information even after the UPF switchover.

The Sx/N4 heartbeat messages are sent and the Active UPF responds back even after back-to-back UPF switchovers.

Sx/N4 Monitor

It is critical to monitor the Sx/N4 interface between the UPF and SMF. The SRP monitoring is enabled on Sx/N4 interface and the existing Sx/N4 heartbeat mechanism is leveraged to detect the monitor failure. The Sx/N4 module on Active UPF, on detecting the failure, informs the SRP VPNMgr to trigger UPF switchover event so that the Standby UPF takes over.

Note Sx/N4 monitoring is available only in the UPF.

It is important to ensure that the SMF Sx/N4 heartbeat timeout is higher than the UPF Sx/N4 heartbeat timeout plus UPF ICSR switchover time. This is to ensure that the SMF does not detect the Sx/N4 path failure during a UPF switchover because of the UPF Sx/N4 monitor failure.

The Standby UPF itself has no independent connectivity to the SMF. The Active UPF Sx/N4 context is replicated to the Standby UPF so that it is ready to takeover during SRP switchover. This implies that when the Active UPF has switched over to Standby because of Sx/N4 monitor failure, the new Standby has no way

of knowing if the UPF to SMF link is working. To prevent a switchback of the new Standby to Active state again due to Sx/N4 monitor failure in new Active, use the **disallow-switchover-on-peer-monitor-fail** keyword in the **monitor sx** CLI command.

After a chassis becomes Standby due to Sx/N4 monitoring failure, the Sx/N4 failure status is not reset even if Sx/N4 up checkpoint is received from the new Active UPF. This is to prevent the new Active to cause an unplanned switchback again due to Sx/N4 monitor failure when the previous cause of switchover itself was Sx/N4 monitor failure. This prevents back-to-back switchovers when SMF is down. The Sx/N4 monitor failure status must be manually reset when the operator is convinced that the network connectivity is normal. To reset, use the new **srp reset-sx-fail** CLI command (see *Resetting Sx/N4 Monitor Failure*) in the Standby chassis.

To configure the Sx/N4 monitor, see *Configuring Sx/N4 Monitoring on the Active UPF and Standby UPF*.

Sx/N4 Monitor—Pending-Active

The UPF chassis can turn into Pending-Active state for one of the following reasons:

- When Sx/N4 heartbeat times out during SMF upgrade, the Sx/N4 connection is terminated. So, Sx/N4 monitoring failure triggers ICSR switchover in UPF. This switchover causes the old Standby UPF to transition to Pending-Active state. The UPF in Pending-Active state neither receives any Sx/N4 heartbeats from SMF nor any subscriber traffic. As a result, the UPF remains in Pending-Active state indefinitely and can't be utilized without a manual intervention.

- When appropriate procedure to upgrade UPF is not followed, one of the UPF may end up in Pending-Active state. Also, if SMF goes down during the UPF upgrade or if the UPF switchover takes more time than the SMF heartbeat timeout, then one of the UPF remains in Pending-Active state indefinitely.

- When Sx/N4 session times out between SMF and UPF due to network issues, and if a UPF ICSR switchover happens almost simultaneously (Double fault scenario), the UPF in Pending-Active state doesn't transition to Active state.

Whenever a UPF chassis turns Pending-Active, start a timer with a callback which forcefully transitions the UPF from Pending-Active to Active state. Before forcing the transition, check if the SRP link is up and if the SRP peer is in Standby state. If not, restart the timer. The duration of the timer is configurable using **force-pactv-to-actv-timeout** *value_seconds* CLI command (see *Changing UPF state from Pending-Active to Active* section for configuration details). When this CLI command is not configured, the UPF remains in Pending-Active state indefinitely.

BGP Monitor

Configure BGP peer monitor and peer group monitors for the next-hop routers from UPF (both Gi and Gn side). This is the existing ICSR configuration. BGP may run with BFD assist to detect fast BGP peer failure.

Figure 7: BGP Peer Groups and Routing

To configure BGP monitoring and flag BPG monitoring failure, see *Configuring BGP Status Monitoring Between Each UP and Next-Hop Router*.

UPF Session Checkpoints

The Active chassis sends a collection of UPF data as checkpoints to the peer Standby chassis in the following scenarios:

- New call setup

- For every state change in the call

- Periodically for accounting buckets

On receiving these checkpoints, the Standby chassis acts on the data and updates the necessary information either at the call, node, or instance level.

VPN IP Pool Checkpoints

During Sx/N4 Association, the IP pool allocated to each of the UPF is sent by SMF to the respective UPF. The VPNMgr receives this message in the UPF and checkpoints the same information to the Standby UPF when the SRP is configured.

The IP pool information is also sent during the SRP VPNMgr restart and during the SRP link down and up scenarios.

Validation of the presence of IP pool information in the Standby is vital before switchover. If the IP pool information is not present, then route advertisement is not possible. Therefore, traffic does not reach the UPF.

External Audit and PFD Configuration Audit Interaction

External Audit management is done in Active UPF. The Session Manager gets a start and complete notification of the Configuration Audit. The Session Manager does not start the External Audit if Configuration Audit is in progress. If the Configuration Audit start-notification arrives when the External Audit is already underway, then the Session Manager raises a flag such that the External Audit restarts when it completes. Restarting the External Audit is necessary because it does not achieve its purpose if it occurs when Configuration Audit is already underway.

Zero Accounting Loss for User Plane Function

Zero accounting loss feature is implemented on User Plane Function (UPF) so that accounting-data/billing loss is reduced from 18 seconds, which is the default checkpoint time from Active UPF to Standby UPF, or for the configured accounting checkpoint time.

This change in UPF is to support the Gz, Gy, VoGx, and RADIUS URRs. Only planned switchover is supported for zero accounting loss/URR data counters loss. This feature doesn't impact the current ICSR framework or the way checkpointing is done and recovered.

The Sx/N4 usage report is blocked during the "pending active state" until the chassis becomes Active.

Early PDU Recovery for UPF Session Recovery

Early PDU Recovery feature overcomes the earlier limitation of Session Recovery feature wherein it didn't prioritize the CRRs that were selected for recovery. All the CRRs were fetched from the AAAMgr and then the calls were recovered sequentially. The time taken to fetch all the CRRs was a major factor in the perceived delay during session recovery. When a failure occurred, the delay was sometimes long if there were many sessions in a Session Manager. Also, since the calls were recovered in no particular order, the idle sessions were sometimes recovered before active sessions.

Note   The Early PDU Recovery feature can recover a maximum of 5 percent sessions.

Session Prioritization during Recovery

Without this functionality, the Session Recovery function didn't prioritize the sessions selected for recovery and loops through all the calls in the call recovery list, and are recovered sequentially when the session recovery is triggered.

As part of Session Prioritization during Recovery, a separate skip list is maintained only for priority calls so that these records can be sent from AAAMgr immediately without going through the loop, thus leading to quicker recovery of the priority calls and reducing the data outage time.

There are two types of sessions at User Plane—prioritized sessions and normal sessions. Session is considered to be prioritized session based on message priority flag received from SMF and it's recovered first followed by normal calls. These prioritized sessions also take priority in case of early PDU handling. The early PDU of normal calls initiates recovery only when all prioritized sessions are recovered.

In case of critical flush (GR), checkpoints for prioritized sessions are sent first followed by the normal calls. The data of all the calls (both normal and prioritized) are allowed during switchover.

Note   The SMF is responsible to set the priority flags for all the calls. The UPF uses the priority call details received from the SMF for the Session Prioritization feature.

Configuring 1:1 UPF Redundancy

The following sections provide information about the CLI commands available in support of the feature.

Configuring BFD Monitoring Between Active UPF and Standby UPF

Use the following configuration to configure Bidirectional Forwarding Detection (BFD) monitoring on the Active UPF and Standby UPF. Configure this command in the SRP Configuration Mode.

```
configure
  context context_name
    service-redundancy-protocol
      [ no ] monitor bfd context context_name { ipv4_address | ipv6_address }
{ chassis-to-chassis | chassis-to-router }
      end
```

NOTES:

- **no**: Disables BFD monitoring on the Active and Standby UPF.

- **context** *context_name* : Specifies the context that is used. It refers to the context where the BFD peer is configured (SRP context).

  context_name must be an existing context expressed as an alphanumeric string of 1 through 79 characters.

- *ipv4 _address* | *ipv6_address*: Defines the IP address of the BFD neighbor to be monitored, entered using IPv4 dotted-decimal or IPv6 colon-separated-hexadecimal notation.

  It refers to the IP address of the configured BFD (ICSR) peer.

- **chassis-to-chassis** | **chassis-to-router**:

  chassis-to-chassis: BFD runs between primary and backup chassis on non-SRP links.

  chassis-to-router: BFD runs between chassis and router.

  ⚠️

  | **Caution** | Don't use the **chassis-to-router** keyword for BFD monitoring on the SRP link between the Active UPF and the Standby UPF. |

- This command is disabled by default.

Configuring BGP Status Monitoring Between Each UPF and Next-Hop Router

Use the following commands to configure Border Gateway Protocol (BGP) monitoring between each UPF and next-hop router. The command is configured in the SRP Configuration Mode.

```
configure
  context context_name
    service-redundancy-protocol
      [ no ] monitor bgp context bgp-session-context-name [
nexthop-router-ipv4-address | nexthop-router-ipv6-address ] { vrf
bgp-session-vrf-name } { group group-number  }
      end
```

NOTES:

- **no**: Disables BGP status monitoring on the UPF.

- **bgp context** *bgp-session-context-name*: Specifies the context where BGP peer is configured. *bgp-session-context-name* specifies the context string.

- **nexthop-router-ipv4-address** | **nexthop-router-ipv6-address**: Specifies the configured BGP peer IPv4 or IPv6 address to monitor.

- **vrf** *bgp-session-vrf-name*: Specifies the BGP VPN Routing and Forwarding (VRF) instance. *bgp-session-vrf-name* specifies the VRF name.

- **group** *group-number* : Specifies the BGP peer group where the BGP peer should be included. *group-number* specifies the group number.

  On implementing this keyword, the behavior is as follows:

  - If any BGP peer in that group is up, the BGP peer group is up.

    Omitting group configuration for a BGP monitor includes that monitor in group 0.

    BGP group 0 monitors in a context from an implicit group. Each context forms a separate BGP group 0 implicit monitor group.

    If any BGP peer group is down, BGP monitor is down.

- This command is disabled by default.

Alternate Algorithm to Flag BGP monitoring failure

In this release, an alternate (new) algorithm is introduced to flag BGP monitoring failure.

Use the following commands to flag BGP monitor failure on a single BGP peer (User Plane Function) failure. This command is configured in the SRP Configuration Mode.

```
configure
   context context_name
      service-redundancy-protocol
         [ no ] monitor bgp exclusive-failover
         end
```

NOTES:

- **no**: Disables flagging of BGP monitor failure on a single BGP peer failure.

- On implementing the new **exclusive-failover** keyword, the behavior is as follows:
  - BGP peer group is Up if any BGP peer in that group is Up.
  - Including a BGP peer in group 0 is same as making it non-group (omitting group).
  - BGP monitor is down if any BGP peer group or any non-group BGP peer is down.

- This command is disabled by default.

Configuring Sx/N4 Monitoring on the Active UPF and Standby UPF

Use the following configuration to configure Sx/N4 monitoring on the Active UPF and Standby UPF. This command is configured in the SRP Configuration Mode.

```
configure
  context context_name
    service-redundancy-protocol
      [ no ] monitor sx [ { context context_name | bind-address {ipv4_address
 |ipv6_address } | { peer-address { ipv4_address | ipv6_address } } ]
      end
```

NOTES:

- **no**: Disables Sx/N4 monitoring on the Active and Standby UPF.

- **context** *context_name* : Specifies the context of the Sx/N4 service.

  context_name must be an existing context expressed as an alphanumeric string of 1 through 79 characters.

- **bind-address {** *ipv4 _address* | *ipv6_address* **}**: Defines the service IP address of the Sx/N4 service, entered using IPv4 dotted-decimal or IPv6 colon-separated-hexadecimal notation.

  Note   The IP address family of the **bind-address** and **peer-address** must be same.

- **peer-address {** *ipv4 _address* | *ipv6_address* **}**: Defines the IP address of the Sx/N4 peer, entered using IPv4 dotted-decimal or IPv6 colon-separated-hexadecimal notation.

- **disallow-switchover-on-peer-monitor-fail**:

  Prevents the switchback of the UPF to Active state when the working status of the UPF to SMF link is unknown.

- It's possible to implement this CLI command multiple times for monitoring multiple Sx/N4 connections.

- The Sx/N4 monitor state goes down when any of the monitored Sx/N4 connections are down.

- This command is disabled by default.

Configuring VPP Monitor on Active UPF and Standby UPF

Use the following configuration to configure Vector Packet Processing (VPP) monitor to trigger UPF switchover on the Active UPF if VPP goes down. This command is configured in the SRP Configuration Mode.

```
configure
  context context_name
    service-redundancy-protocol
      monitor system vpp delay-period seconds
      end
```

NOTES:

- If previously configured, use the **no monitor system vpp** CLI command to disable VPP monitoring on the Active and Standby UPF.

- **vpp delay-period** *seconds* : Specifies the delay period in seconds for a switchover, after a VPP failure. *seconds* must be in the range of 0 through 300.

If the delay period is a value greater than zero (0), then the switchover is initiated after the specified delay period when VPP fails. The last VPP status notification within the delay period is the final trigger for switchover action. The default value is 0 seconds, which initiates an immediate switchover.

The need for delay is to address the scenario wherein the VPP is temporarily down and the revival is in process. This implies that a switchover may not be necessary.

• This command is disabled by default.

Preventing User Plane Function Switchback

Use the following configuration to prevent the switchback of the new Standby UPF to Active state again due to Sx/N4 monitor failure in the new Active.

```
configure
  context context_name
    service-redundancy-protocol
      monitor sx disallow-switchover-on-peer-monitor-fail timeout seconds

        end
```

Use either of the following CLIs to allow switchback of the new Standby UPF to Active state.

```
no monitor sx disallow-switchover-on-peer-monitor-fail
```

Or

```
monitor sx disallow-switchover-on-peer-monitor-fail timeout 0
```

NOTES:

• **no**: Disables prevention of switchover.

• **disallow-switchover-on-peer-monitor-fail [timeout** *seconds* **]** : Prevents the switchback of the UPF to Active state when the working status of the UPF to SMF link is unknown.

timeout *seconds*: Timeout after which the switchback is allowed even if the Sx/N4 failure status is not reset in the Standby peer. The valid values range from 0 through 2073600 (24 days).

> **Note**   Assigning 0 seconds as the timeout allows unplanned switchover.

If **timeout**   keyword is not specified, the Active chassis waits indefinitely for the Sx/N4 failure status to be reset in the Standby peer.

• The default configuration is to allow unplanned switchover due to Sx/N4 monitor failure in all conditions.

> **Note**   Manual planned switchover is allowed irrespective of whether this CLI is configured or not.

Preventing Dual Active Error Scenarios

Use the following CLI configuration in CP to prevent dual Active error scenarios for UPF 1:1 redundancy.

```
configure
    user-plane-group group_name
        sx-reassociation disabled
        end
```

NOTE:

- **sx-reassociation disabled**: Disables UP Sx reassociation when the association already exists with the CP.

Resetting Sx/N4 Monitor Failure

Use the following configuration only on the Standby chassis to reset the Service Redundancy Protocol (SRP) Sx/N4 monitor failure information. This command is configured in the Exec Mode.

```
srp reset-sx-fail
```

Changing UPF State from Pending-Active to Active

Use the following configuration to change the UPF chassis state from Pending-Active to Active.

```
configure
  context context_name
    service-redundancy-protocol
      force-pactv-to-actv-timeout value_seconds
```

NOTES:

- *value_seconds*: Specifies the timeout value in seconds and must be in the range of 1-300.

- Use the **show config context** *context_name* CLI command to verify the configuration.

Monitoring and Troubleshooting

This section provides information regarding the CLI command available in support of 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 srp monitor bfd

The output of this CLI command contains the following fields for the 5G UPF 1:1 Redundancy feature:

- Type:

  - (A) - Auth. probe

- (B) - BGP

- (D) - Diameter

- (F) - BFD

- (E) - EGQC

- (C) - Card

- (V) - VPP

- State:

  - (I) - Initializing

  - (U) - Up

  - (D) - Down

- GroupId

- IP Addr

- Port

- Context (VRF Name)

- Last Update

show srp monitor bgp

The output of this CLI command contains the following fields for the 5G UPF 1:1 UPF Redundancy feature:

- Type:

  - (A) - Auth. probe

  - (B) - BGP

  - (D) - Diameter

  - (F) - BFD

  - (E) - EGQC

  - (C) - Card

  - (V) - VPP

  - (S) - Sx

- State:

  - (I) - Initializing

  - (U) - Up

  - (D) - Down

- GroupId

- IP Addr

- Port

- Context (VRF Name)

- Last Update

show srp monitor sx

The output of this CLI command contains the following fields in support of Sx/N4 monitor status:

- Type:

  - (A) - Auth. probe

  - (B) - BGP

  - (D) - Diameter

  - (F) - BFD

  - (E) - EGQC

  - (C) - Card

  - (V) – VPP

  - (S) - SX

- State:

  - (I) - Initializing

  - (U) - Up

  - (D) - Down

- GroupId

- IP Addr

- Port

- Context (VRF Name)

- Last Update

show srp monitor vpp

The output of this CLI command contains the following fields for the 5G UPF 1:1 UPF Redundancy feature:

- Type:

  - (A) - Auth. probe

  - (B) - BGP

- (D) - Diameter

- (F) - BFD

- (E) - EGQC

- (C) - Card

- (V) - VPP

- State:

  - (I) - Initializing

  - (U) - Up

  - (D) - Down

- GroupId

- IP Addr

- Port

- Context (VRF Name)

- Last Update

show srp statistics

The output of this CLI command contains the following fields for the Sx/N4 Monitor—Pending-Active functionality:

- Pending-active timer started

- Pending-active timer stopped

- Pending-active to Active forced

- Pending-active to Active force-failed

- Pending-active to Active force-skipped - peer-not-sby

- Pending-active to Active force-skipped - not-PActv

CHAPTER **4**

APN ACL Support

This chapter covers the following topics:

Feature Summary and Revision History

Summary Data

Table 5: Summary Data

Applicable Product (s) or Functional Area	5G-UPF
Applicable Platforms	VPC-SI
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

Table 6: Revision History

Revision Details	Release
A configurable mechanism has been introduced to apply traffic classification and policy enforcement on selective subscriber sessions.	2021.01.0
First introduced.	2020.02.0

Feature Description

IP Access Lists, commonly known as Access Control Lists (ACLs), control the flow of packets into and out of the system. The configuration is per-context basis and consists of "rules" (ACL rules) or filters that control the action applicable for packets that match the filter criteria. Once configured, an ACL can be applied to an individual subscriber. Separate ACLs can be created for IPv4 and IPv6 access routes.

The following are the two main aspects of ACLs:

- Rule(s)

- Rule Order

Rule(s)

A single ACL consists of one or more ACL rules. Each rule is a filter configured to take a specific action when packets match a specific criteria.

Each rule specifies the action to take when a packet matches the specifies criteria. This section discusses the rule actions and criteria supported by the system.

Actions

ACLs specify that one of the following actions can be taken on a packet that matches the specified criteria:

- **Permit**: The packet is accepted and processed for classification and policy enforcement.

- **Deny**: The packet is rejected.

- **Redirect CSS**: The behaviour is same as Permit action.

NOTES:

- In UPF, it's recommended to use Permit option instead of Redirect CSS. Functionally, both the options are equivalent in UPF. Support for Redirect CSS option is only for backward compatibility and should be used only in such scenarios.

- Configured ACLs consisting of no rules imply a "deny any" rule. This is the default behavior for an empty ACL.

- In UPF, if ACLs aren't associated with an APN, then call is up. By default, traffic is processed for classification and policy enforcement. For non-UPF architecture, call fails as Redirect CSS is mandatory.

- If only Deny option is given in the ACL for certain traffic, then to pass the rest of the traffic, Permit option must be given explicitly.

- If only permit option is given in the ACL for certain traffic, then to pass the rest of the traffic, permit must be given explicitly for that traffic.

- Router Advertisement/Router Solicitation (RA/RS) packets are candidate for ACL. So, take caution in putting the IPv6 ACL.

- Configuration change in ACL is applied for a new call and not on the existing call.

Criteria

Each ACL consists of one or more rules specifying the criteria that packets will be compared against.

The following criteria are supported:

- **Any**: Filters all packets

- **Host**: Filters packets based on the source host IP address

- **ICMP**: Filters Internet Control Message Protocol (ICMP) packets

- **IP**: Filters Internet Protocol (IP) packets

- **Source IP Address**: Filter packets based on one or more source IP addresses

- **TCP**: Filters Transport Control Protocol (TCP) packets

- **UDP**: Filters User Datagram Protocol (UDP) packets

Each of the above criteria is described in detail in the sections that follow.

- **Any**: The rule applies to all packets.

- **Host**: The rule applies to a specific host as determined by its IP address.

- **ICMP**: The rule applies to specific Internet Control Message Protocol (ICMP) packets, Types, or Codes. ICMP type and code definitions can be found at *www.iana.org* (RFC 3232).

- **IP**: The rule applies to specific Internet Protocol (IP) packets or fragments.

- **Source IP Address**: The rule applies to specific packets originating from a specific source address or a group of source addresses.

- **TCP**: The rule applies to any Transport Control Protocol (TCP) traffic and could be filtered on any combination of source/destination IP addresses, a specific port number, or a group of port numbers. TCP port numbers definitions can be found at *www.iana.org*.

- **UDP**: The rule applies to any User Datagram Protocol (UDP) traffic and could be filtered on any combination of source/destination IP addresses, a specific port number, or a group of port numbers. UDP port numbers definitions can be found at *www.iana.org*.

Rule Order

A single ACL can consist of multiple rules. Each packet is compared against each of the ACL rules, in the order in which they were entered, until a match is found. Once a match is identified, all subsequent rules are ignored.

Limitations

Following are the known limitations of APN ACL feature in UPF:

- Readdress option in ACL is not supported.

- Redirect ACL for context and next-hop is not supported.

- Log option is not supported in ACLs.

• APN-level bulkstats for ACL drops (only IPv4) are supported.

Configuring ACL

To apply the ACL to individual subscriber through via APN, use the following configuration:

```
configure
   context dest_context_name [ -noconfirm ]
      { ip | ipv6 } access-list acl_list_name
         { permit | deny | redirect }acl
         end
configure
   apn apn_name
      { ip | ipv6 } access-group acl_list_name [ in | out ]
      end
```

Notes:

• The ACL to be applied must be in the destination context of the APN (which can be different from the context where the APN is configured).

• If neither the **in** nor the **out** keyword is specified, the ACL will be applied to all inbound and outbound packets.

• Four access-groups can be applied for each APN, for example:

ip access-group acl_list_name_1 in

ip access-group acl_list_name_2 out

ipv6 access-group acl_list_name_3 in

ipv6 access-group acl_list_name_4 out

Verifying ACL Configuration

Use the following CLI commands in Exec mode to check if your ACL lists were applied properly, and also for packet drops due to ACL:

• **show subscriber user-plane-only full all**

• **show subscribers user-plane-only full callid** *call_id*

• **show user-plane-service pdn-instance statistics** *name*

IP Source Violation

Source validation requires the source address of incoming packets to match the IP address of the subscriber during the session. This allows operators to configure the network to prevent problems when a user gets handed back and forth between two gateways several times during a handoff scenario.

When the UPF receives a subscriber packet with a source IP address violation, the system increments the IP source violation drop-limit counter and starts the timer for the IP source violation period. Every subsequent packet received with a bad source address during the IP source violation period causes the drop-limit counter

to increment. For example, if you set the drop limit to 10, after 10 source violations, the call is dropped. The detection period timer continues to count throughout this process.

The following must be configured in the User Planes APN configuration:

```
ip source-violation { ignore | check [ drop-limit limit ] } [
exclude-from-accounting ]
```

NOTE: For information on IP source violation CLI commands, refer to the StarOS *Command Line Interface Reference*.

Gating Control

Gating Control in the UPF enables or disables the forwarding of IP packets belonging to a service data flow or detected application's traffic to pass through to the desired endpoint. See 3GPP TS 23.203, subclause 4.3.2.

The SMF controls the gating in the UPF by creating PDRs for the service data flow(s) or application's traffic to be detected, and by associating a QER, including the Gate Status IE, to the PDRs.

The Gate Status IE indicates whether the service data flow or detected application traffic is allowed to be forwarded (the gate is open) or to be discarded (the gate is closed) in the uplink and/or in downlink directions.

The UPF identifies the UL and DL flows by the Source Interface IE in the PDI of the PDRs or the destination Interface IE in the FARs. The UPF applies UL and DL gating accordingly.

The SMF requests the UPF to discard the packets that are received for the PDR by setting the gate fields in the Gate Status IE of QERs to CLOSED.

APN AMBR Traffic Policing

Feature Summary and Revision History

Summary Data

Table 7: Summary Data

Applicable Product (s) or Functional Area	5G-UPF
Applicable Platforms	VPC-SI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Table 8: Revision History

Revision Details	Release
First introduced.	2021.01.0

Feature Description

The APN-AMBR is a subscription parameter stored per APN in the HSS. S-GW provides APN-AMBR during default bearer establishment procedure. APN-AMBR limits the aggregate bit rate that can be expected to be

provided across all non-GBR bearers and across all PDN connections of the same APN. Each of those non-GBR bearers can potentially utilize the entire APN-AMBR, for example, when the other non-GBR bearers does not carry any traffic. The P-GW enforces the APN-AMBR in downlink and uplink direction.

Limitations

The **token-replenishment-interval** and **violate-action shape** CLI commands are not supported.

Configuring the APN AMBR Traffic Policing Feature

This section describes how to configure the APN-AMBR Traffic Policing feature.

```
configure
        context context_name
          apn apn_name
      apn-ambr rate-limit direction { downlink | uplink } [ burst-size
{ auto-readjust duration { milliseconds msecs | seconds } | violate-action
 { drop | lower-ip-precedence | transmit }
                end
```

NOTES:

- **rate-limit direction { downlink | uplink }**: Specifies that the rate limit is to be applied to either the downlink (network to subscriber) traffic or the uplink (subscriber to network) traffic.

- **burst-size { auto-readjust duration milliseconds msecs | seconds }**: This parameter is used by policing algorithms to permit short bursts of traffic not to exceed the allowed data rates. It's the maximum size of the token bucket.

  - **auto-readjust duration**_seconds_: The duration (in seconds) used in this burst size calculation: burst size = peak data rate/8 * auto-readjust duration.

    - Seconds must be an integer value from 1-30. Default is 1 second.

  - **milliseconds**: _msecs_ must be an integer value from 100-900, in increments of 100 milliseconds. For example, 100, 200, or 300, and so on.

- **violate-action { drop | lower-ip-precedence | transmit }**: The action that the P-GW takes when the data rate of the bearer context exceeds the AMBR.

  - **drop**: Drops violating packets.

  - **lower-ip-precedence**: Sets the DSCP value to zero ("best effort") for violating packets.

  - **transmit**: Transmits violating packets. This is the default behavior of the feature.

- Prior to this feature, the default behavior was to drop the violating packets.

Monitoring and Troubleshooting

This section provides information about the commands available to monitor and/or troubleshoot the APN-AMBR Traffic Policing feature.

Show Commands and/or Outputs

This section provides information about the show commands available for monitoring and/or troubleshooting the APN-AMBR Traffic Policing feature.

- **show user-plane-service pdn-instance name <apn_name>**

Use this show command in UPF to see if the rate limit is enabled/disabled, burst size, and other such parameters for downlink/uplink traffic:

  - APN-AMBR

    - Downlink Apn Ambr: Indicates if the rate limit is enabled or disabled for downlink traffic.

      - Burst Size: Indicates the burst size of the downlink traffic.

      - Auto Readjust: Indicates if the auto-readjust is enabled or disabled for downlink burst size.

      - Auto Readjust Duration: Indicates the duration used in downlink burst size calculation.

      - Burst Size(bytes): Indicates the burst size in bytes.

      - Violate Action: Indicates the action that the P-GW takes when the data rate of the bearer context exceeds the AMBR for downlink traffic.

    - Uplink Apn Ambr: Indicates if the rate limit is enabled or disabled for uplink traffic.

      - Burst Size: Indicates the burst size of the uplink traffic.

      - Auto Readjust: Indicates if the auto-readjust is enabled or disabled for uplink burst size.

      - Auto Readjust Duration: Indicates the duration used in uplink burst size calculation.

      - Burst Size(bytes): Indicates the burst size in bytes.

      - Violate Action: Indicates the action that the P-GW takes when the data rate of the bearer context exceeds the AMBR for uplink traffic.

    - Token Replenishment Interval: Indicates the token replenishment interval duration.

- **show sub user-plane-only full all**

Use this show command in UPF to see the count of packets that are dropped, and IP precedence lowered due to APN-AMBR policer. The following fields are introduced in support of this feature:

  - APN AMBR Uplink Pkts Drop: Indicates the number of APN-AMBR packets that are dropped for uplink traffic.

- APN AMBR Uplink Bytes Drop: Indicates the number of APN-AMBR bytes that are dropped for uplink traffic.

- APN AMBR Uplink Pkts IP pref lowered: Indicates the number of APN-AMBR uplink packets for which IP precedence is lowered.

- APN AMBR Uplink Bytes IP pref lowered: Indicates the number of APN-AMBR uplink bytes for which IP precedence is lowered.

- APN AMBR Downlink Pkts Drop: Indicates the number of APN-AMBR packets that are dropped for downlink traffic.

- APN AMBR Downlink Bytes Drop: Indicates the number of APN-AMBR bytes that are dropped for downlink traffic.

- APN AMBR Downlink Pkts IP pref lowered: Indicates the number of APN-AMBR downlink packets for which IP precedence is lowered.

- APN AMBR Downlink Bytes IP pref lowered: Indicates the number of APN-AMBR downlink bytes for which IP precedence is lowered.

Bulk Statistics Support

Feature Summary and Revision History

Summary Data

Table 9: Summary Data

Applicable Product (s) or Functional Area	5G-UPF
Applicable Platforms	VPC-SI
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

Table 10: Revision History

Revision Details	Release
The following new bulk statistic schemas are now supported: • P2P Schema • Sx Schema • System Schema • Userplane Schema	2021.01.0

Revision Details	Release
First introduced.	2020.02.0

Feature Description

This chapter identifies bulk statistic schemas for the Cisco Ultra Cloud 5G User Plane Function (UPF) software release.

Bulk statistics is a collection of software features and framework that collects and exports the important performance and health-related statistics of the packet core node to an external node. These statistics provide an effective way for the operators to perform the following functions:

- Monitor the overall health and performance of the nodes.

- Help take corrective actions.

- Optimize the packet core network for better utilization.

- Reduce the overall operation expenses.

The individual statistics are configured to be collected in a group called 'schema.'

The system-supported bulk statistics allows operators to choose statistics that are of importance to them and configure the presentation format. This simplifies the post-processing of statistical data because it allows data formatting that facilitates external, backend processors to parse it.

Statistics or bulk statistics reporting is important on a Mobile Packet Core node. For a product to be deployed in the network, it has to support statistics that meets Carrier Grade requisites.

Operators use bulk statistics for the following:

- Performance KPI monitoring

- Network Fault analysis and debugging

- Network Optimization

- Traffic pattern analysis

- Node health analysis

When used along with an element management system (EMS), the data can be parsed, archived, and graphed.

In the 5G environment, the system can be configured to collect for the following network functions:

- Access and Mobility Management Function (AMF)

- Network Repository Functions (NRF)

- Network Slice Selection Functions (NSSF)

- Policy Control Function (PCF)

- Session Management Function (SMF)

- User Plane Function (UPF)

The system supports the configuration of up to four sets (primary and secondary) of receivers. Each set is configured to collect specific sets of statistics from the supported list of schemas. Statistics can be pulled manually from the system or sent at configured intervals. The bulk statistics are stored on the receivers in files.

A user can configure the format of the bulk statistic data files. Users can specify the following:

- Format of the filename

- File headers and footers to include information such as the date, system hostname, and system uptime

- IP address of the system generating the statistics (available for only for headers and footers)

- Time that the file was generated

An EMS is capable of further processing the statistics data through XML parsing, archiving, and graphing. The Bulk Statistics Server component of an EMS parses collected statistics and stores the information in its PostgreSQL database. It can also generate XML output and can send it to a Northbound NMS or an alternate bulk statistics server for further processing.

Additionally, the Bulk Statistics server can archive files to an alternate directory on the server. The directory can be on a local file system or on an NFS-mounted file system on an EMS server.

The implementation of bulk statistics in 5G is as follows:

- The NFs collect and export the statistics separately to an aggregator node in the 5G architecture.

- The receiver correlates the statistics from the NFs using the node-names or any other information that is configured as part of the bulk statistics configuration. Any EMS tool can render this data similar to how it is rendered from a standalone system.

Supported Schemas

This release supports the following schemas in the 5G architecture.

APN Schema

The APN schema provides Access Point Name (APN) statistics.

Card Schema

The Card schema provides card-level statistics.

ECS Schema

The ECS schema provides Enhanced Charging Service statistics.

GTP-U Schema

The GTP-U schema provides GPRS Tunneling Protocol- User message statistics.

P2P Schema

The P2P schema provides P2P statistics.

P-GW Schema

The P-GW schema provides user-plane service statistics.

Port Schema

The Port schema provides port-level statistics.

Rulebase

The Rulebase schema provides rule base statistics.

Sx Schema

The Sx schema provides N4 related message statistics.

System Schema

The System schema provides system-level statistics.

Userplane Schema

The Userplane schema provides User Plane statistics.

☞

Important For more information on bulk statistic configuration, refer to the *Bulk Statistics* chapter in the *ASR 5500 System Administration Guide*.

CHAPTER 7

Charging Support

This chapter covers the following topics:

Feature Summary and Revision History

Summary Data

Table 11: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI
Feature Default Setting	Disabled - Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Table 12: Revision History

Revision Details	Release
The feature is enhanced to support following functionality: • PTT no-quota Limited Pass • PTT quota exhaust Limited Pass • Tariff Time • TCP Maximum Segment Size	2021.02.0
Usage reporting with Rating-Group and Service ID is introduced.	2020.02.5
First introduced.	2020.02.0

Feature Description

The usage measurement and reporting function in User Plane Function (UPF) is controlled by the Session Management Function (SMF). The SMF controls these functions by:

- Creating the necessary PDRs to represent the service data flow, application, bearer or session (if they are not existing already).

- Creating the URRs for each Charging Key and combination of Charging Key and Service ID. Also, creating URRs for a combination of Charging Key, Sponsor ID, and Application Service Provider Id.

- Associating the URRs to the relevant PDRs defined for the PFCP session, for usage reporting at SDF, Session or Application level.

- For online charging, the SMF provisions Volume and Time quota, if it receives it from the Online Charging Server (OCS).

Offline Charging Events Reporting over N4

The User Plane Function (UPF) supports session-based offline charging, PDU session level reporting triggers in URR (volume and time threshold), PFCP session report procedure, and usage report IE support in the PFCP modification response for the Session-AMBR change, QoS, and User Location triggers.

Online Charging Support over N4

The UPF supports flow-based online charging support, which includes URR enhancements for Volume and Time quota and Usage reporting IE in PFCP modify response. In addition, the UPF supports online charging triggers, which include a PFCP session report request support with usage reporting IE.

How it Works

Call Flows

The following sections describe the call flows between SMF and UPF for PFCP Session Management.

PFCP Session Establishment Procedure

The PFCP Session Establishment procedure establishes a PFCP session between SMF and UPF. It also configures rules in UPF for handling incoming packets. In addition, the SMF sends Create URR IE, which comprises of triggers and thresholds that are intended for reporting.

The following call flow depicts the PFCP Session Establishment procedure.

Step	Description
1	The SMF receives the trigger to establish a new PDU Session or change the UPF for an established PDU Session.
2	The SMF sends the PFCP Session Establishment Request message to the UPF. This message contains the structured control information, which defines the UPFs behavior.
3	The SMF provisions URR with Create URR IE. The Create URR associates with PDRs by adding URR-ID IE in Create PDR IE. It includes various triggers and thresholds for usage reporting.
4	When the same URR is associated with multiple PDRs, URRs are linked with another URR. Therefore, if a report for an URR is sent, its linked URR is also reported.
5	The UPF responds with the PFCP Session Establishment Response message to the SMF. For instance, Created PDR IE, in which UPF Flow-TEID is sent to gNB for GTP-u encap for data traffic.
6	The SMF interacts with the network function, which triggered this procedure. For instance, AMF or PCF.

PFCP Session Modification Procedure

The SMF uses the PFCP Session Modification procedure to modify an existing PFCP session on the UPF. For instance, configuring a new rule, modifying an existing rule, or deleting an existing rule, and so on. The SMF sends the Create URR IE, Update URR IE (to update the trigger or threshold) and Remove URR IE (to remove an existing URR created earlier by SMF during Session Establishment Procedure) in the same message.

The following call flow depicts the PFCP Session Modification procedure.

Step	Description
1	The SMF receives the trigger to modify the existing PDU Session.
2	The SMF sends an N4 session modification request message to the UPF. This message contains the structured control information, which defines the UPFs behavior.
3	The UPF identifies the PFCP session context for the Session ID to modify. It updates the parameters of this session context according to the list of parameters sent by the SMF. It then responds with a PFCP Session Modification Response message. The message contains the information, which the UPF must provide to the SMF (in response to the control information received).
4	If the SMF sends the QAURR flag set in PFCPSMReq-Flag IE or URR ID (s) with Query URR IE (e), then UPF sends the usage report IE for the corresponding URR with the PFCP Session Modification response.
5	The UPF provisions and acts based on the Create URR, Update URR or Remove URR IE sent by the SMF.
6	The SMF interacts with the network function, which triggered this procedure. For instance, AMF or PCF.

PFCP Session Reporting Procedure

The UPF uses PFCP Session Reporting procedure to report information that is related to the PFCP session to the SMF (usage report IE). Once the threshold hits the volume, time or event measurement and sets the corresponding trigger for reporting, the message is sent to the SMF by the UPF.

The following call flow depicts the PFCP Session Reporting procedure.
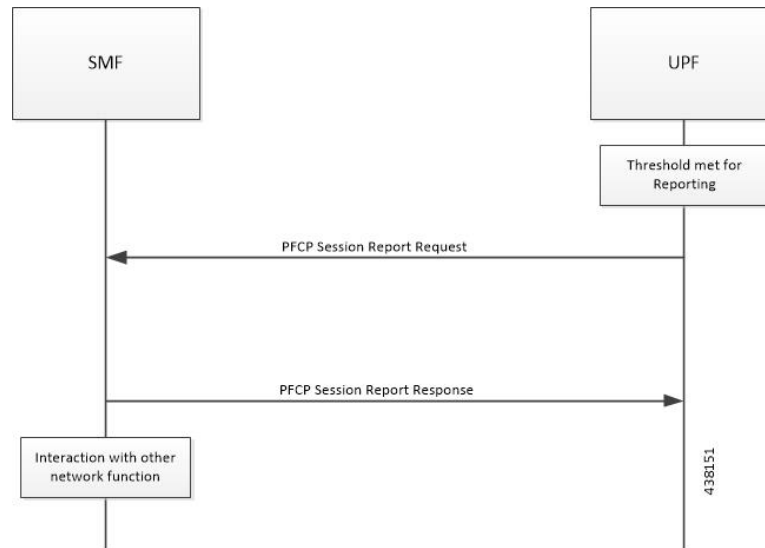
Step	Description
1	Once the provisioned threshold is met (for time, volume or event, and trigger is set for reporting), the UPF sends PFCP Session Report Request with usage report IE and usage details for volume, time, or threshold.
2	The SMF responds with PFCP Session Modification Response with success or failure message. No failure handling is needed on the UPF.
3	The SMF interacts with the network function, which triggered this procedure. For instance, AMF or PCF.

PFCP Session Deletion Procedure

The PFCP Session Deletion procedure deletes an existing PFCP session between the SMF and UPF. The SMF initiates a PFCP Session Deletion procedure toward the UPF to delete an existing PFCP session. The UPF sends the Session Deletion Response including the Usage Report for all URRs provisioned earlier.

The following call flow depicts the PFCP Session Deletion procedure.

Step	Description
1	The SMF receives the trigger to remove the PFCP session context for the PDU Session.
2	It sends the PFCP Session Delete Request message to the UPF.
3	The UPF identifies the PFCP session context for the Session ID to remove. It then removes the whole session context. In addition, the UPF responds with a PFCP Session Delete Response message that contains any information the UPF provides to the SMF. For instance, the UPF sends usage report for all the URR provisioned for this session.
4	The SMF interacts with the network function, which triggered this procedure. For instance, AMF or PCF.

IEs Supported for Offline Charging Reporting

The following trigger Information Elements (IEs) support offline charging Reporting over N4:

- Periodic Reporting – When this trigger is set, the UPF sends resource usage report periodically to Session Management Function (SMF). The intervals that are required for periodic reporting are sent with the measurement period IE.

- Volume Threshold (when the volume threshold reaches UL, DL, and Total) – This trigger is set when the volume-based measurement is required. The SMF sends the traffic volume value along with the volume threshold IE, while the UPF sends the traffic usage report when the traffic volume is reached for the specific Usage Reporting Rule (URR).

- Time Threshold (when the time threshold is reached) – This trigger is set when the time-based measurement is set. The SMF sends the time threshold value along with the time threshold IE, while the UPF sends resource usage report when the time threshold is reached for the specific URR.

- Linked Usage Reporting – The UPF sends the usage report of this specific URR when this trigger is set. In addition, the usage report is sent to any of the URRs linked to UPF when this trigger is set. The UPF sends the linked URR-Id along with the linked URR-Id IE.

- Packet Forwarding Control Protocol (PFCP) Session Deletion – A usage report generates (in a PFCP Deletion Response) for a URR due to the termination of the PFCP session. Similarly, a usage report generates (in a PFCP modification response) for a URR due to the removal of a specific URR.

- Update URR – This trigger is set when update URR request is received.

IEs Supported for Online Charging Reporting

The following IEs support online charging:

- Volume Quota – The SMF requests the UPF to stop forwarding packets or allow forwarding some limited user plane traffic (based on the operator policy in UPF) with this IE. If no Volume Threshold is provisioned – to generate a usage report – and when the measured traffic reaches the quota, this IE is used.

- Time Quota - The SMF requests the UPF to stop forwarding packets or allow forwarding some limited user plane traffic (based on operator policy in UPF) with this IE. If no Volume Threshold is provisioned – to generate a usage report – and when the measured traffic reaches the quota, this IE is used.

- Monitoring Time – This IE is used by the SMF to send the time (UTC format) at which the UPF can re-apply the volume or time threshold. Also, the SMF sends any one of the Subsequent Volume, Time, Volume Quota, Time Quota, and Quota IEs, which is re-applied at the Monitoring Timestamp.

- FAR (Forwarding Action Rule) ID for Quota Action – This IE is used by the SMF to identify the substitute FAR the UPF applies – for the traffic that is associated to the URR – when any of the Volume, Time or quota is exhausted. This FAR requires the UPF to drop the packets or redirect the traffic toward a redirect destination.

- Subsequent Volume Threshold – When volume-based measurement is used and Monitoring Time IE is available, this IE is also present. The presence of this IE indicates the existence of the traffic volume value (the network resources usage reported by the UPF to the SMF) for this specific URR and the period after the Monitoring Time.

- Subsequent Time Threshold - When time-based measurement is used and Monitoring Time IE is available, this IE is also present. The presence of this IE indicates the existence of the time usage (the network resources usage reported by the UP function to the CP function) for this specific URR and the period after the Monitoring Time.

- Linked URR ID – When the linked usage reporting is required, this IE is used. It is possible to link multiple URR-IDs with an URR. Also, linked usage reporting is also sent in the Reporting Trigger IE.

- Measurement Method – The SMF specifies the measurement method of the network usage with the presence of this IE. The measurement method is based on volume and duration.

- Measurement Period – This IE is present to modify the measurement period.

- Periodic Reporting - When this trigger is set, the UPF sends resource usage report periodically to the SMF. The intervals that are required for periodic reporting are sent with the measurement period IE. When the trigger is set to 1, a request for periodic reporting is sent.

- Volume Threshold – This trigger is set when volume-based measurement is required. The SMF sends the traffic volume value along with the volume threshold IE, while the UPF sends the traffic usage report when the traffic volume is reached for the specific Usage Reporting Rule (URR). When the trigger is set to 1, a request for reporting – when the data volume usage reaches a volume threshold – is sent.

- Time Threshold - This trigger is set when time-based measurement is set. The SMF sends the time threshold value along with the time threshold IE, while the UPF sends resource usage report when the

time threshold is reached for the specific URR. When the trigger is set to 1, a request for reporting – when the time usage reaches a time threshold - is sent.

- Start of Traffic – The UPF sends the Usage Report once the traffic starts for an application, when this trigger is set.

- Linked Usage Reporting - The UPF sends the usage report of this specific URR when this trigger is set. In addition, the usage report is sent to any of the URRs linked to UPF when this trigger is set. The UPF sends the linked URR-Id along with the linked URR-Id IE. When the trigger is set to 1, a request for linked usage reporting is sent.

Usage Reporting in PFCP Modification Response

The UPF sends session modification response after receiving session modification request based on the IEs received in the request message. The UPF includes usage report IE in the session modification response for the following scenarios:

- Query URR Handling—The URR-Id IE is included when the SMF requests immediate usage reports from the UPF in the session modification response (for the URR-Id present in this specific IE).

- Query All URRs (QAURR) Handling—The UPF sends the usage report with session modification response for all the URRs provisioned prior by the SMF for this PFCP session once it receives the QUARR flag set in PFCPSMReq-Flags IE from SMF.

- Update URR—The SMF updates the new value of the existing IE with the old value during the session modification procedure.

- Remove URR—During the session modification procedure, the SMF removes the IE, which is not received but was available earlier.

Usage Reporting for Online and Offline Charging

Usage Reporting for Online and Offline Charging is supported in the following ways:

- URR for online charging based on Rating-Group level even if the Service ID is present under Charging-Action. This behavior is seen when diameter ignore-service-id is configured under Credit Control Group.

- URR for offline charging based on a combination of Rating-Group level and Service ID, for static and predefined rules, as configured in the Charging-Action.

Both URRs are linked by the SMF. These URRs are linked such that when an online URR is reported, an offline URR is also reported.

Usage Reporting with Rating-Group and Service ID

The functionality enables usage reporting to the SMF with the Rating-Group (RG) and/or Service ID (SI) populated in the Usage Report IE within the Session Report Request.

The RG and SI are populated using proprietary PFCP IEs and are applicable for usage reporting of URRs associated only with Static and Predefined configured rules. The values are derived from the configured

charging-action associated with the ruledefs, resulting in creation of the URRs during predefined activation or traffic hit for static rules.

Any change in RG/SI properties of the charging-action is reflected only in new URRs. The existing URRs associated with such charging-actions continue to report usage with the earlier RG+SI values.

UPF does not differentiate between usage reporting for Online and Offline URRs, and reports the RG+SI/RG/SI values configured in the charging-action, resulting in creation of the URRs.

NOTE: To know how SMF handles this functionality, refer *Dynamic Configuration Change Support* section in the *SMF Charging* chapter of *UCC 5G SMF Configuration and Administration Guide*.

Implementing the QAURR Flag

The SMF sets the QAURR flag of PFCPSMReq-Flags IE to request immediate usage reports for all the URRs previously provisioned earlier. Alternatively, SMF queries report for selected URR by sending URR-ID with Query URR IE. The UPF sends the usage report IE for corresponding URR with PFCP session modification response when the SMF sends the QAURR flag set in PFCPSMReq-Flag IE or URR-Id with Query URR IE.

Supported Functionality and Limitations

Basic call flow with Volume-Quota mechanism is supported with the following limitations:

- Dynamic Rules with Online Enabled is supported; both at Session-Setup and Mid-Session.

- Predefined Rules (dynamic-only) is supported; both at Session-Setup and Mid-Session. No restriction on configuring the "preemptively request".

- Static-rules with Online Charging are supported.

- Ignore-service-id is supported.

- Volume-Quota/Volume-Threshold mechanisms are supported.

- Event-Triggers (through which the Query URR occurs), and sending of usage information to the OCS is supported.

- The "updateURR" procedure, through the Sx/N4-Session-Modification procedure where the OCS grants a fresh Quota, is supported.

- Pending-Traffic-Treatment (PTT) Drop/Pass is supported with following limitations:

    - The scenarios supported for now are no-quota and quota-exhausted.

    - The trigger/re-authorization scenarios are not supported.

    - The PTT action (Forward/Drop) is considered after the quota-get is exhausted.

- Wall-Clock time-quota mechanism is supported.

- Other Time Quota Mechanisms (Discrete Time Period and Continuous-Time-Period) are not supported.

- Final-Unit-Indication Terminate mechanism is supported.

- FUI-Restrict is not supported.

- Server-Unreachable (SU) mechanism is now supported with minor change in behavior compared to non-5G UPF P-GW.

  - When an URR needs quota at UPF, the usage-report is generated to SMF and until the SMF responds with the linked SU_URR, the packets matching this URR are treated with Pending-Traffic-Treatment configuration.

  - When the SU Time Quota is used and it's reported to SMF for the Quota Exhaust, and if the session goes into Server-Unreachable state again, the time elapsed from the last Usage-Report is accounted in the usage.

- Pending-Traffic-Treatment Buffer mechanism is not supported.

- Quota-Hold-Time is supported.

- Quota-Consumption-Time mechanism is not supported.

- Quota-Validity-Time is supported.

- Configuring different "rating-group" value other than the "content-id" is supported.

  - The RG 0 is not supported.

- Trigger to PCF for the Out-of-Credit, Reallocation-of-Credit events are not qualified.

> ☞
>
> **Important**   Event-trigger Out-of-Credit towards PCF is validated with a limitation of having only one time Grant-Quota (Keeping Total Volume and Granted Volume at same value).

- Service-Specific-Units are not supported.

- Tariff-Time change is supported as per 3GPP specification.

- FUI-Redirect is supported with following limitations:

  - Redirection for HTTPs is not supported.

  - The FUI-Redirect with Filter-IDs/Filter-Rules are not supported.

  - The WSP Protocol is not supported.

  - The **redirect-require-user-agent** CLI command is not supported; the redirection continues to work even if the user-agent is not present.

  - Appending the original URL is not supported.

  - Token based mechanism, to come out of Redirection, is not supported. To end the redirection in 5G UPF, OCS sends Redirect Validity-Time or RAR.

  - FUI-Redirection is supported only for the URL, similar to the behavior in non-5G UPF architecture.

  - Check pointing of FUI Redirection URL is not supported.

PTT no-quota Limited Pass

This feature allows the subscriber to use the network while waiting for the response from OCS. The Limited-Pass configuration allows to specify the Volume which the subscriber can consume while waiting for the quota-response from OCS. The usage is accounted in the respective charging bucket and is adjusted against the next-quota allocation.

Use the following CLI commands to enable the feature:

```
configure
   active-charging service service_name
      credit-control
         pending-traffic-treatment noquota limited-pass volume volume
         end
```

Limited Pass Volume is used only for **noquota** case (Rating Group (RG) seeking quota for the first time) and not for **quota-exhausted**. Limited Pass Volume isn't used for subsequent credit requests.

The traffic is allowed to pass until the Limited-Pass Volume gets exhausted. The usage is counted in the respected charging-bucket and adjusted against the "Quota" granted. If the "Quota" allocation is less than the actual usage, immediate reporting towards OCS with the usage-report occurs requesting for more quota allocation. The subsequent incoming packets are handled as per the "quota-exhausted" PTT configuration.

If the Limited Pass Volume is NOT exhausted before the OCS responds with denial of quota, traffic is blocked after the OCS response. The gateway reports usage on Limited-Pass Volume in next SX_SESSION_REPORT_REQUEST.

If the Limited Pass Volume is exhausted before the OCS responds, then the subsequent incoming packets for the session are dropped until quota is granted from OCS.

The default pending-traffic-treatment for **noquota** is Drop. The **default pending-traffic-treatment noquota** command removes any Limited Pass Volume size configured.

PTT quota exhaust Limited Pass

Quota Exhausted Limited pass is proposed as an alternative to the Quota Exhausted Buffer due to the practical issues of the later in the high-speed network. Buffering requires packet buffering for large number of packets at the gateway. The large number of packets can result in risking to run out of memory affecting the bandwidth speed. So, Limited Pass is an alternate to the Buffer option. Limited Pass allows the traffic to pass through until the configured limit on the Quota-Exhaust scenarios.

Use the following CLI command to enable the feature:

```
configure
   active-charging service service_name
      credit-control
        pending-traffic-treatment quota-exhausted limited-pass volumevolume

        end
```

Note   The above CLI is only applicable for the 5G UPF architecture.

After the Limited-Pass volume exhausts, the further packets drop until the quota allocation comes.

Limited Pass allows the traffic until the Limited-Pass volume exhausts. The Limited Pass counts and adjusts the usage in the respective charging-bucket against the "Quota" granted. If the "Quota" allocation is less than the actual usage, there's immediate reporting towards OCS with the usage-report and asking for more quota allocation.

If the limited pass volume doesn't exhaust before the OCS responds with denial of quota, there's traffic blockage after the OCS response. Gateway reports the usage in SX_SESSION_REPORT_REQUEST.

If the limited pass volume exhausts before the OCS responds, then further incoming packets for the session drop until grant quota from OCS.

The default pending-traffic-treatment for quota-exhausted is drop. The default pending-traffic-treatment quota-exhausted command removes any Limited Pass Volume size configured.

Tariff Time Support

The Tariff switch time functionality is applied when a subscriber switch from one tariff plan to another.

The Tariff-Time-Change AVP is used to determine the tariff switch time, and the Monitoring-Time IE is used to support the Tariff Time support functionality.

After a tariff timer expiry, the Gateway accumulates the usage separately in a charging bucket and continues to consume from the original quota value. At the time of next reporting (Quota exhausted or another control events), the Gateway reports both usages (before and after tariff time change) for the same Charging Bucket.

The first reporting of this charging-bucket will have the Reporting-Reason: Monitoring Time and the second bucket will contain the last reporting reason, and the quota usage after the tariff-timer expiry.

The data traffic usage can be split into resource usage before a tariff switch and resources used after a tariff switch. The Tariff-Change-Usage AVP is used within the Used-Service-Units AVP to distinguish reported usage before and after the tariff time change.

Limitations

Following are the known limitations of this feature:

- Only one tariff time per RG/Service ID combination is supported.

- Allocation of different quota before and after tariff time change isn't supported. This functionality isn't in compliance with the 3GPP standards.

TCP Maximum Segment Size

TCP IP Stack always inserts Maximum Segment Size (MSS) field in the header. This causes difference in MSS insertion behavior with and without TCP Proxy.

Using **tcp mss** configurations, TCP MSS can be limited if already present in the TCP SYN packets. If there are no errors detected in IP header/TCP mandatory header, and there are no memory allocation failures, TCP optional header is parsed. If TCP MSS is present in the optional header and its value is greater than the configured MSS value, the value present in the TCP packet is replaced with the one that is configured.

If the TCP optional header is not present in the SYN packet and there are no errors in already-present TCP header, the configured TCP MSS value is inserted while sending out the current packet.

Configuring Credit Control for Usage Reporting

This configuration enables to accept/ignore service ID in the Service-Identifier AVP defined in the Diameter dictionaries.

```
configure
   require active charging
   active-charging service service_name
      credit-control group group_name
         diameter ignore-service-id
         end
```

- **diameter ignore-service-id** : This command can be used to disable the usage of the Service-Identifier AVP for Gy interface implementations even if any of the Diameter dictionaries support the Service-Identifier AVP, and if this AVP should not be used for Gy interactions but must be present in GCDRs/eGCDRs.

Configuring ACS Rulebase for Usage Reporting

This section describes how to create, configure, or delete an ACS rulebase. A rulebase is a collection of protocol rules to match a flow and associated actions to be taken for matching flow. The default rulebase is used when a subscriber/APN is not configured with a specific rulebase to use.

Rulebase configuration is the one that combines all the specified configurations together to construct the static and predefined PCC rules.

```
configure
   active-charging service service_name
      rulebase rulebase_name
         action priority action_priority { [ dynamic-only ] |
static-and-dynamic | timedef timedef_name ] { group-of-ruledefs
ruledefs_group_name | ruledef ruledef_name } charging-action charging_action_name [
 monitoring-key monitoring_key ] [ description description ] }
         cca quota { holding-time holding_time content-id content_id |
retry-time retry_time [ max-retries retries ] }
         credit-control-group cc_group_name
         dynamic-rule order { always-first | first-if-tied }
         egcdr threshold { interval interval [ regardless-of-other-triggers
 ] | volume { downlink | total | uplink } bytes }
         route priority route_priority ruledef ruledef_name analyzer { dns |
file-transfer | ftp-control | ftp-data | h323 | http | imap | mipv6 | mms
 | pop3 | pptp | radius | rtcp | rtp | rtsp | sdp | secure-http | sip [
advanced | basic-and-advanced ] | smtp | tftp | wsp-connection-less |
wsp-connection-oriented } [ description description ]
         tcp check-window-size
         tcp mss tcp_mss { add-if-not-present | limit-if-present |
limit-if-present add-if-not-present }
         tcp packets-out-of-order { timeout timeout_duration | transmit [
```

```
after-reordering | immediately ] }
        end
```

NOTES:

- **rulebase** *rulebase_name*: Specifies the name of the ACS rulebase. *rulebase_name* must be an alphanumeric string of 1 to 63 characters.

- **action priority** *action_priority* **{ [dynamic-only] | static-and-dynamic | timedef** *timedef_name* **] { group-of-ruledefs** *ruledefs_group_name* **| ruledef** *ruledef_name* **} charging-action** *charging_action_name* **[monitoring-key** *monitoring_key* **] [description** *description* **] }**: Configures the priority order in which ruledefs are matched and the associated charging action.

    - *priority* must be an integer value in the range of 1-65535.

    - *monitoring_key* must be an integer value in the range of 100000-4000000000.

- **cca quota { holding-time** *holding_time* **content-id** *content_id* **| retry-time** *retry_time* **[max-retries** *retries* **] }**: Configures the quota for the online charging.

    - *holding_time*: must be an integer value in the range of 1-4000000000

    - *content_id*: must be an integer value in the range of 1-2147483647

    - *retry_time*: must be an integer value in the range of 0-86400

    - *retries*: must be an integer value in the range of 1-65535

- **credit-control-group** *cc_group_name*: Configures the online charging parameters used by this rulebase. *cc_group_name* must be an alphanumeric string of 1 to 63 characters.

- **dynamic-rule order**: Configures the order of dynamic rule matching vs the static rules in a rulebase.

- **egcdr threshold { interval** *interval* **[regardless-of-other-triggers] | volume { downlink | total | uplink } bytes }**: Configures the threshold for offline charging.

    - **interval**: must be an integer value in the range of 60-40000000.

    - **downlink**: must be an integer value in the range of 100000-4000000000. Default: 4000000000.

    - **uplink**: must be an integer value in the range of 100000-4000000000. Default: 4000000000.

    - **total**: must be an integer value in the range of 100000-4000000000.

- **route priority** *route_priority* **ruledef** *ruledef_name* **analyzer { dns | file-transfer | ftp-control | ftp-data | h323 | http | imap | mipv6 | mms | pop3 | pptp | radius | rtcp | rtp | rtsp | sdp | secure-http | sip [advanced | basic-and-advanced] | smtp | tftp | wsp-connection-less | wsp-connection-oriented } [description** *description* **]**: This command is used only on UPF.

    - *route_priority* must be an integer value in the range of 0-65535.

    - *ruledef_name* must be an alphanumeric string of 1 to 63 characters.

- **tcp check-window-size**: This command is used only on UPF.

- **tcp mss** *tcp_mss*: This command is used only on UPF. *tcp_mss* must be an integer value in the range of 496-65535.

    - **add-if-not-present** : Specifies to add the TCP MSS if not present in the packet.

- **limit-if-present** : Specifies to limit the TCP MSS if present in the packet.

- **limit-if-present add-if-not-present** : Specifies to limit the TCP MSS if present, else, adds it to the packets.

> ✎
>
> **Note**   The **tcp mss** *tcp_mss***limit-if-present add-if-not-present** CLI command is available in 2021.02.0 and later releases.

- **tcp packets-out-of-order { timeout** *timeout_duration* **| transmit [after-reordering | immediately] }**: This command is used only on UPF.

  - *timeout_duration* must be an integer value in the range of 100-30000. Default value is 5000.

Sample Configuration

```
active-charging service acs
    ruledef ip-any-rule
      ip any-match = TRUE
    #exit
    urr-list upf
      rating-group 10 ser 10 urr-id 10
      rating-group 10 urr-id 50
    #exit
    charging-action starent
      content-id 10
      service-identifier 10
      billing-action egcdr
      cca charging credit rating-group 10
    exit
    credit-control group CCG
      diameter ignore-service-id
    #exit
    rulebase starent
      billing-records egcdr
      action priority 30 ruledef ip-any-rule charging-action starent
      egcdr threshold interval 3600
      egcdr threshold volume total 200000
      egcdr threshold volume downlink 100000 uplink 100000
      dynamic-rule order first-if-tied
      credit-control-group CCG
    #exit
#exit

context ISP
    apn starent.com
      accounting-mode gtpp
      gtpp group my_grp accounting-context ISP
      ip context-name ISP
    #exit
    gtpp group my_grp
      gtpp egcdr service-data-flow threshold interval 1200
      gtpp egcdr service-data-flow threshold volume downlink 13000
      gtpp egcdr service-data-flow threshold volume uplink 17000
      gtpp egcdr service-data-flow threshold volume total 22222
    #exit
end
```

Monitoring and Troubleshooting

Show Commands and/or Outputs

This section provides information about the show CLI commands that are available in support of the feature.

show-user-plane-service statistics rulebase name <name>

Use this CLI command to see the following fields that are available in support of TCP Maximum Segment Size (MSS) feature:

- TCP MSS Inserted Pkts: Displays the total number of MSS Inserted packets.

- TCP MSS Limited Pkts: Displays the total number of TCP MSS Limited packets.

Collection and Reporting of Usage Data over N4 Interface

This chapter covers the following topics:

Feature Summary and Revision History

Summary Data

Table 13: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Table 14: Revision History

Revision Details	Release
First Introduced	2020.02.0

Feature Description

With this release, the User Plane Function (UPF) supports offline charging and reporting of usage data over the N4 interface.

Here, the SMF controls the collection and reporting of usage data by creating necessary PDRs and URRs, and associates the URRs with its relevant PDRs defined for a PFCP session. It also controls data usage reporting at an IP-CAN bearer level, IP-CAN session, TDF session, SDF, or at an application level.

The URR consists of the usage measurement method, reporting triggers, threshold, and quota values.

☞

Important   In this release, only URR creation is supported during PFCP session establishment.

How it Works

This section describes how UPF supports offline charging of usage data.

To implement offline charging, the charging information is sent to the SMF only during PFCP session deletion.

Time and volume-based reporting is supported in the offline charging implementation. The following call flow illustrates offline charging in UPF.

Figure 8: Offline Charging in UPF

During the PFCP session deletion, UPF transfers the following charging information to the SMF:

- Timestamp of the first and last data packet

- Duration measurement – This IE specifies the time difference between URR creation and usage-reporting

- Volume measurement – This IE specifies the uplink data, downlink data and the total bytes transferred from the UPF to gNodeB.

Standards Compliance

UPF support for collection and reporting of data is compliant with the following standards:

- 3GPP TS 29.244 - LTE; Interface between the Control Plane and the User Plane of EPC Nodes

- 3GPP TS 23.501 - 5G; System Architecture for the 5G System

- 3GPP TS 23.502v - 5G; Procedures for the 5G System

Configuration to Collect and Report Volume Measurement over N4 Interface

This section describes the configuration required to collect and report volume measurement (usage data). However, to achieve this, SMF-based configurations for volume measurement needs to be configured.

The following SMF-based configuration is required to send volume measurement data in the URR by the UPF.

Configuring Charging Action for a Required Billing Action

Use the following configuration to configure charging-action for a required billing-action:

```
configure
   require active-charging
   active-charging service service_name
      charging-action charging_action_name
         billing-action interface_name
         end
```

NOTES:

- **billing-action**: Enables the specified billing type. The supported interfaces are:

  - **egcdr**: Enables the GGSN charging data record.

Associating a Charging Action with a Rulebase

Use the following configuration to associate a charging action with a rulebase:

```
configure
```

```
require active-charging
active-charging service service_name
   rulebase rulebase_name
       billing-records interface_name
       action priority priority_value ruledef ruledef_name charging-action
charging_action_name
           end
```

NOTES:

- **rulebase**: Enables the Active Charging Service Rulebase configuration.

- **billing-records**: Enables the generation of billing records. The supported interface is **egcdr**

- **action**: Decides the action to be taken on the ruledef.

- **priority**: Assigns priority to a ruledef in the rulebase. Priority must be a unique integer value ranging from 1 to 65535.

- **ruledef**: Specifies the ruledef.

- **charging-action**: Specifies the charging action.

CHAPTER 9

Control Plane-Initiated N4 Association Support

This chapter covers the following topics:

Feature Summary and Revision History

Summary Data

Table 15: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Table 16: Revision History

Revision Details	Release
First Introduced.	2020.02.0

Feature Description

SMF initiated N4 Association Setup Procedure

The N4 association set up procedure sets up an N4 association between the Session Management Function (SMF) and User Plane Function (UPF). It enables the SMF to use the UPF resources to establish the N4 sessions. The SMF and UPF exchange the supported functionalities on each side during this procedure.

How it Works

The N4 association setup is initiated by the SMF. The setting of N4 association setup procedure is controlled through **sx-association initiated-by-cp** CLI command in the Control Plane Group Configuration mode. By default, the configuration is set to support the UPF-initiated N4 association setup procedure.

Call Flows

Session Management Function Initiated N4 Association Setup Procedure

The following call flow depicts the SMF-initiated N4 Association Setup procedure.

Step	Description
1	The SMF initiates the N4 Association Setup procedure to request the setup of an N4 association towards a UPF prior to establishing a first N4 session on this UPF.
2	After receiving an N4 Association Setup Request, the UPF sends an N4 Association Setup Response.

Configuring the CP-Initiated N4 Association Setup Feature

This section describes how to configure the CP-Initiated N4 Association Setup feature.

Configuring this feature involves using the **"sx-association initiated-by-cp"** CLI command in the Control Plane Group Configuration mode. The default configuration is UPF-initiated N4 association setup procedure.

Use the following configuration to configure the N4 association setup feature.

```
configure
   context
      control-plane-group group_name
         peer-node-id ipv4-address ip_address interface n4
         sx-association { initiated-by-cp | initiated-by-up }
         end
```

NOTES:

- **initiated-by-cp:** This keyword is used to initiate the Sx association request through control plane.

- **initiated-by-up:** This keyword is used to initiate the Sx association request through user plane.

- By default, the UPF-initiated N4 association setup procedure is configured.

- To revert to the default setting, use the **no sx-association** command.

CP-Initiated N4 Association Setup Feature OAM Support

This section describes operations, administration, and maintenance information for this feature.

Show Command Support

Use the following show command to verify the CP-initiated N4 Association Setup feature configuration.

show control-plane-group all

The following is a sample output of the show command.

```
show control-plane-group all
Control Plane Group
-------------------------------
Name          : default
Sx-Association : initiated-by-up
Name          : default
Sx-Association : initiated-by-up
Node-Id       : 1.1.2.2
Interface     : N4
```

Show Command Support

Converged Datapath

- Feature Summary and Revision History, on page 73
- Feature Description, on page 74
- How it Works, on page 74
- Configuring Converged Datapath, on page 87
- Monitoring and Troubleshooting, on page 88

Feature Summary and Revision History

Summary Data

Table 17: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

Table 18: Revision History

Revision Details	Release
Support added for WiFi to LTE handover, and configuration to enable Converged Datapath feature at UPF.	2021.02.0
First introduced.	2021.01.0

Feature Description

The Converged Datapath feature allows interconnection of the same UE's session at UPF instance with SGW-C/cnSGW and Session Management Function (SMF)/Inter-Working Function (IWF) to build converged/collapsed datapath and achieve higher throughput. With this feature:

- The UP/UPF selection logic is enhanced to aide same node selection on SGW-C/cnSGW and SMF.

- The SxDemux selects the same Session Manager (SessMgr) instance based on existing session of N4 or Sxa respectively.

- The Sxa session and N4 session correlation is done at SessMgr.

- The datapath is allowed to be collapsed in the forwarding plane.

- Extra hop in subscriber's datapath is eliminated, resulting in reduced latency and improved user experience.

Architecture

As part of this feature, there are two sessions on the same UPF instance established by SGW-C/cnSGW and SMF. Once they are established, the software logic determines the peer session so that the converged/collapsed datapath for packet processing is possible at the UPF node.

How it Works

This section describes how the feature works.

SxDemuxMgr

In distributed architecture of UP/UPF, sessions (Sxa or N4) run on different SessMgr instances. To support collapsing/converging the sessions to a single SessMgr, the SessMgr instance is selected by both sessions during establishment.

At SxDemux, when Sx Establishment Request (Sxa or N4) is received for selecting the SessMgr instance, it's parsed for finding the SessMgr instance from remote F-TEID, where corresponding sessions (N4 or Sxa respectively) are established. The F-TEID, that contains the Tunnel Identifier embedded with SessMgr instance, is extracted.

UPF also maintains IMSI entry at SxDemux. IMSI entry has information of SessMgr ID where the current session is hosted. When Converged Datapath feature is enabled, on receiving Sx Session Establishment Request, SxDemux first tries to find the SessMgr ID using F-TEID. If F-TEID is not present, IMSI lookup is done. If both F-TEID and IMSI is not present, then SxDemux won't select the same SessMgr ID for Sxa/N4 leg.

SessMgr

There are two sessions, Sxa and N4, that exist on the same SessMgr instance. To converge them, the following logic is used to identify the session:

- For Uplink Packet: Egress FARs F-TEID matches with Ingress PDRs F-TEID.

- For Downlink Packet: Ingress FAR's F-TEID should match with Egress PDR's F-TEID.

The F-TEID includes both Tunnel Identifier and the endpoint IP address. After the session is identified, the required information is used in datapath to build the converged datapath.

Datapath

After convergence of session occurs at SessMgr, the SessMgr removes the existing Bearer stream (3 tuple) from Fast Path that is installed for Sxa session. It's established only when flow-level stream (6 tuple), based on received packet, is analyzed.

The uplink packet is received by S-GW ingress PDR endpoint. The downlink packet is forwarded using S-GW Ingress FAR-based outer header.

Charging

Charging of SMF leg (N4 leg) is supported.

Note S-GW charging is not supported.

Call Flows

This section describes the call flows associated with Converged Datapath feature.

Initial Attach with SGW-C/cnSGW and SMF/IWF

The following illustration describes the initial attach call flow with collapsed UPF.
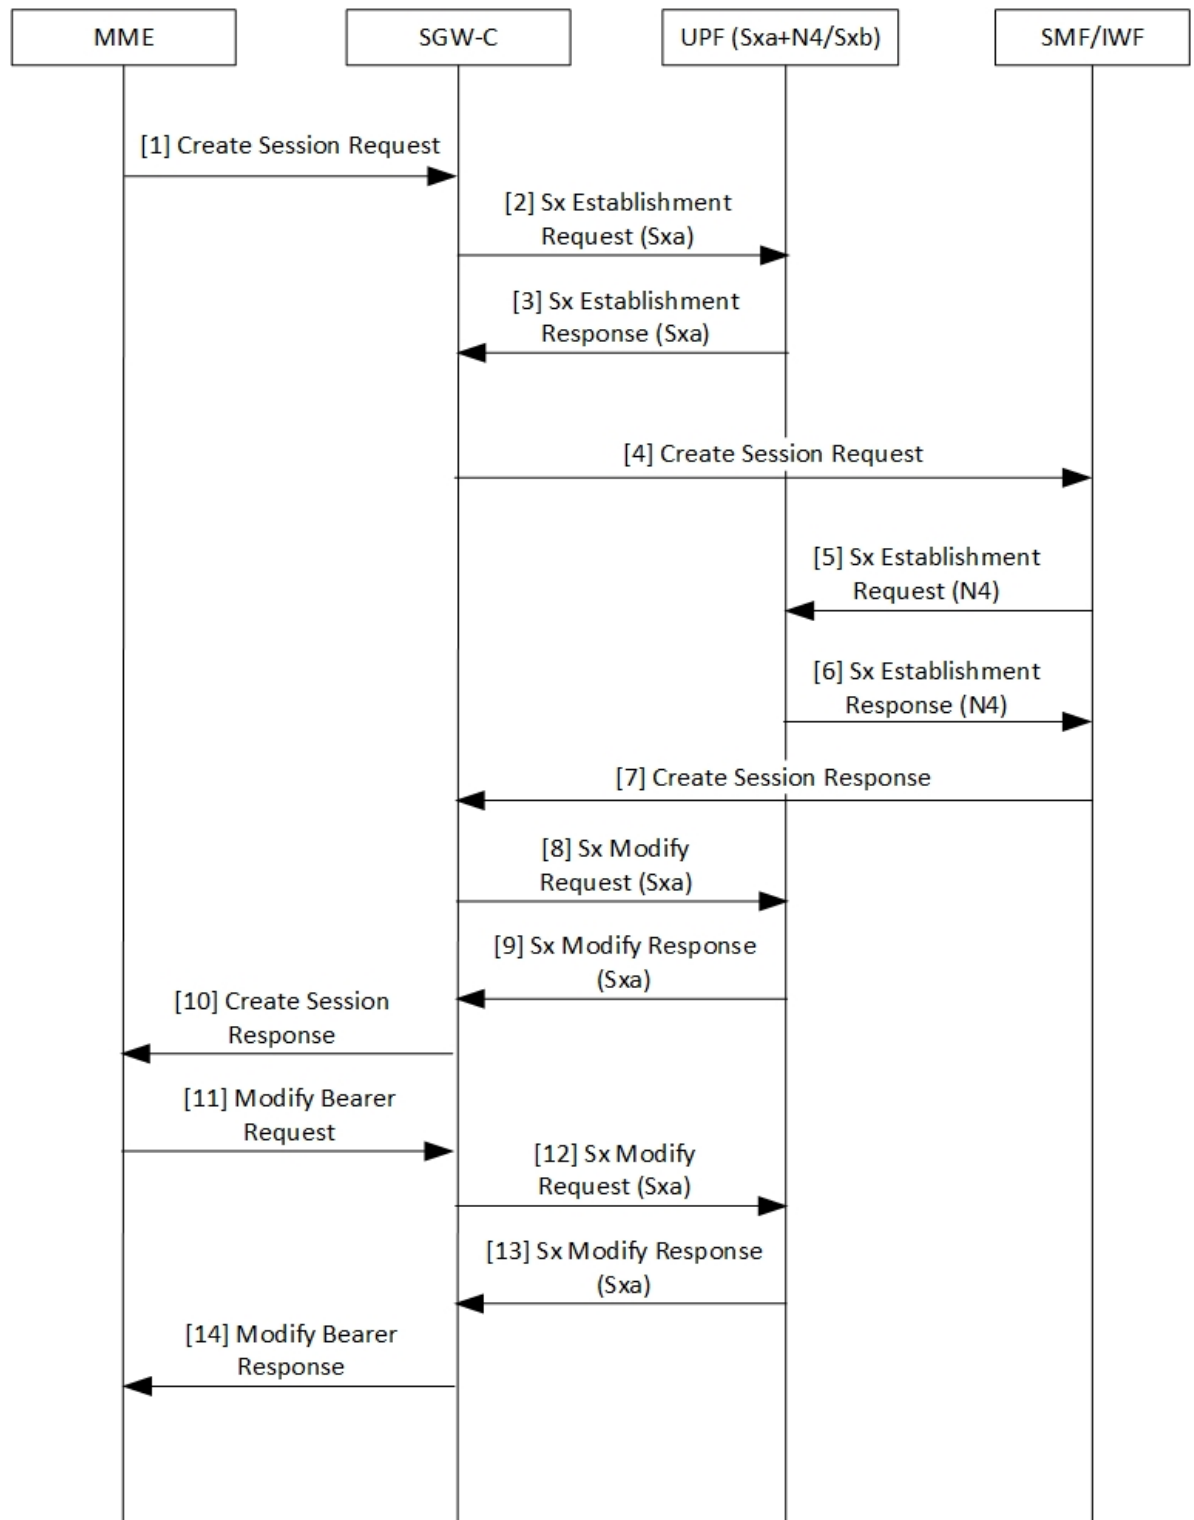
Table 19: Initial Attach with SGW-C/cnSGW and SMF/IWF Call Flow Description

Step	Description
1	Create Session Request (CSReq) is received by SGW-C/cnSGW and it selects the UPF.
2	The SGW-C/cnSGW sends Sx Establishment Request (Sxa) to the UPF. The UPF: • Allocates Sxa session • Allocates S-GW Ingress and Egress local F-TEID
3	The UPF sends Sx Establishment Response (Sxa) back to SGW-C/cnSGW.
4	The SGW-C/cnSGW sends CSReq to SMF/IWF. The SMF/IWF selects the same UPF that is selected by SGW-C/cnSGW.
5	The SMF/IWF sends Sx Establishment Request (N4) to the UPF.
6	The UPF sends Sx Establishment Response (N4) to the SMF/IWF.
7	The SMF/IWF sends Create Session Response to the SGW-C/cnSGW.
8	The SGW-C/cnSGW sends Sx Modify Request (Sxa) to the UPF. The UPF updates P-GW data F-TEID as part of Egress FAR. The UPF also interconnects Sxa and N4 session using internal logic and removes already-created Bearer Stream (3 tuple).
9	The UPF sends Sx Modify Response (Sxa) to the SGW-C/cnSGW.
10	The SGW-C/cnSGW sends Create Session Response to the MME.
11	The MME sends Modify Bearer Request to the SGW-C/cnSGW.
12	The SGW-C/cnSGW sends Sx Modify Request (Sxa) to the UPF. The UPF updates eNodeB F-TEID as part of Ingress FAR.
13	The UPF sends Sx Modify Response to the SGW-C/cnSGW.
14	The SGW-C/cnSGW sends Modify Bearer Response to the MME.

5G to 4G Handover with Collapsed UPF

The following illustration describes the 5G to 4G handover call flow with collapsed UPF.

Table 20: 5G to 4G Handover with Collapsed UPF Call Flow Description

Step	Description
1	As part of UE initial attach, N4 session is already established with SMF and UPF.
	The MME sends Create Session Request (CSReq) to SGW-C/cnSGW. If it's a handover request, the SGW-C/cnSGW selects the same UPF that is selected by the SMF.
2	The SGW-C/cnSGW sends Sx Establishment Request (Sxa) to the UPF. At UPF:
	• SxDemux selects the same SessMgr instance extracted from the P-GW F-TEID that is received in FAR. Both Sxa and N4 session are on the same SessMgr.
	• Allocates Sxa session
	• Allocates S-GW Ingress and Egress local F-TEID
	• Interconnects Sxa and N4 session using internal logic and doesn't install Bearer Stream (3 tuple)

Step	Description
3	The UPF sends Sx Establishment Response (Sxa) back to SGW-C/cnSGW.
4	The SGW-C/cnSGW sends Create Session Response to the MME.
5	The MME sends Modify Bearer Request to the SGW-C/cnSGW.
6	The SGW-C/cnSGW sends Sx Modify Request (Sxa) to the UPF. The UPF updates eNodeB F-TEID for downlink data.
7	The UPF sends Sx Modify Response (Sxa) to the SGW-C/cnSGW.
8	The SGW-C/cnSGW sends Modify Bearer Request to the SMF/IWF.
9	The SMF/IWF sends Sx Modify Bearer Request (N4) to the UPF. The UPF: • Updates N4 session FAR towards S-GW with F-TEID • Updates TEP entries at VPP with new F-TEID
10	The UPF sends Sx Modify Response (N4) to the SMF/IWF.
11	The SMF/IWF sends Sx Modify Response to the SGW-C/cnSGW.
12	The SGW-C/cnSGW sends Modify Bearer Response to the MME.

Intra S-GW Handover with Collapsed UPF

The following illustration describes the intra-SGW handover call flow with collapsed UPF.
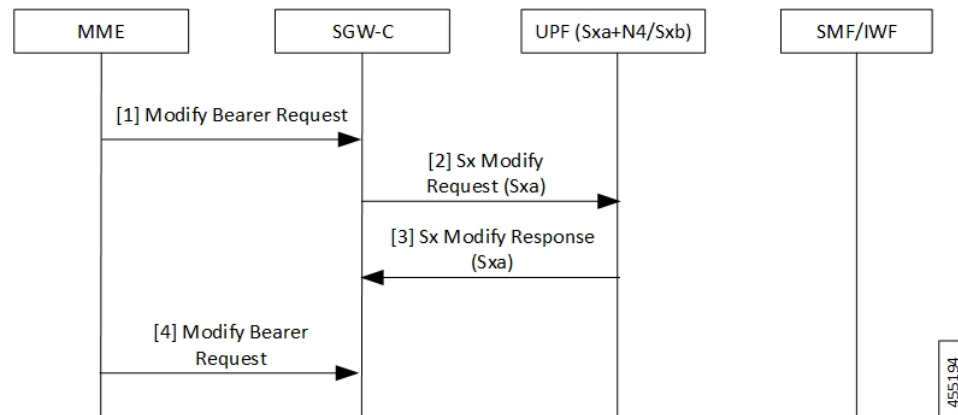
Table 21: Intra S-GW Handover with Collapsed UPF Call Flow Description

Step	Description
1	As part of UE initial attach, N4 and Sxa session is already established with SMF and UPF. At UPF, N4+Sxa session exists and are interconnected. The MME sends Modify Bearer Request to the SGW-C/cnSGW. The SGW-C/cnSGW updates eNodeB F-TEID in FAR.

Step	Description
2	The SGW-C/cnSGW sends Sx Modify Request (Sxa) to the UPF. The UPF: • Updates eNodeB F-TEID for downlink data • Updates TEP entries at VPP with new Remove TEID
3	The UPF sends Sx Modify Response (Sxa) to the SGW-C/cnSGW.
4	The SGW-C/cnSGW sends Modify Bearer Response to the MME.

Idle/Active DDN Handling with Collapsed UPF

The following illustration describes the Idle/Active DDN handling with collapsed UPF.

Table 22: Idle/Active DDN Handling with Collapsed UPF Call Flow Description

Step	Description
1	As part of UE initial attach, N4 and Sxa session is already established with SMF and UPF. At UPF, N4+Sxa session exists and are interconnected. At MME, the UE goes from Active to Ide state. The MME sends Release Access Bearer Request to the SGW-C/cnSGW. The SGW-C/cnSGW informs UPF to buffer packets.
2	The SGW-C/cnSGW sends Sx Modify Request (Sxa) to the UPF. The UPF updates FAR action towards eNodeB to buffer state. 6 tuple flows are onloaded for buffering.
3	The UPF sends Sx Modify Response (Sxa) back to SGW-C/cnSGW.
4	The MME sends Release Access Bearer Response to the SGW-C/cnSGW.
5	The ISP/Gi sends the received downlink packet to the UPF. The packet received by UPF at N4 session is passed to interconnect Sxa session for buffering.
6	The UPF sends Usage Report Request (Sxa) to the SGW-C/cnSGW.
7	The SGW-C/cnSGW sends Usage Report Response (Sxa) to the UPF.
8	The SGW-C/cnSGW sends DDN notification to the MME.
9	The MME sends DDN Acknowledgment back to the SGW-C/cnSGW.
10	At MME, the UE moves from Idle to Active. The MME sends Modify Bearer Request to the SGW-C/cnSGW. The SGW-C/cnSGW: • Updates eNodeB F-TEID in FAR • Updates FAR action to Forward
11	The SGW-C/cnSGW sends Sx Modify Request (Sxa) to the UPF. The UPF: • Updates eNodeB F-TEID for downlink data • Updates TEP entries at VPP with new Remove TEID • Releases buffered packets by finding respective 6 tuple streams
12	The UPF sends Sx Modify Response (Sxa) to the SGW-C/cnSGW.
13	The SGW-C/cnSGW sends Modify Bearer Response to the MME.

IDFT Handling during S1 Handover

The following illustration describes the IDFT handling during S1 handover with collapsed UPF.

Table 23: IDFT Handling during S1 Handover Call Flow Description

Step	Description
1	As part of initial attach, N4 and Sxa session is already established with SMF and UPF. S1 handoff is triggered. The MME sends Indirect Tunnel Request Message to the SGW-C/cnSGW.
2	The SGW-C/cnSGW sends Sx Modify Request (Sxa) to the UPF. At UPF: • IDFT PDR is detected at SessMgr • New F-TEID is allocated for IDFT tunnel • Converged datapath is not required and traffic goes through Slowpath
3	The UPF sends Sx Modify Response (Sxa) to the SGW-C/cnSGW.
4	The SGW-C/cnSGW sends Create Indirect Tunnel Response Message to the MME.
5	The Source eNodeB sends IDFT data packet to the UPF. If there's no matching 3 tuple stream at UPF, the packet is forwarded to SessMgr.
6	The UPF sends IDFT data packet to the destination eNodeB.

S-GW Relocation with Same SGW-U

The following illustration describes the S-GW relocation with destination S-GW selecting the same UPF.
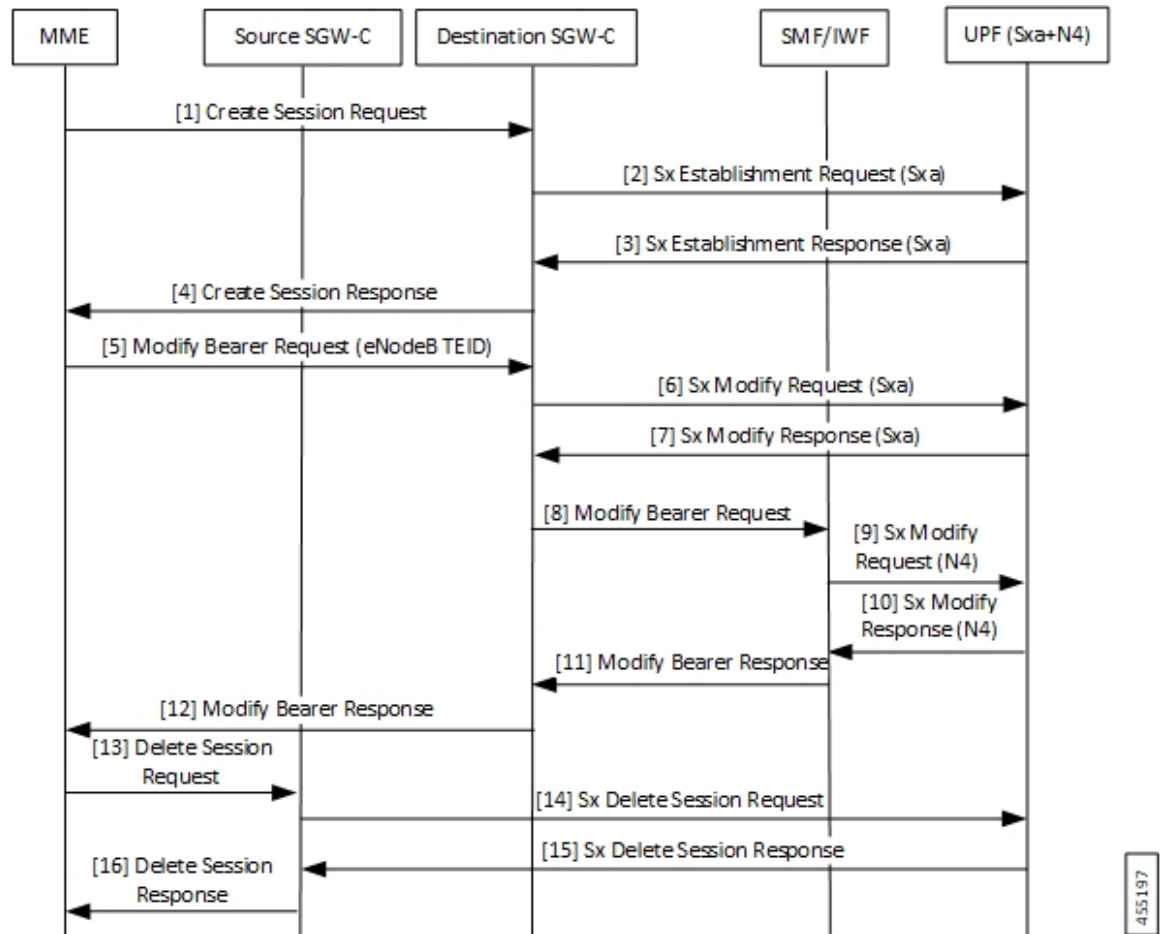
Table 24: S-GW Relocation with Same SGW-U Call Flow Description

Step	Description
1	As part of initial attach, N4 and Sxa session is already established with Source SGW-C/cnSGW, SMF/IWF, and UPF. The MME sends Create Session Request to the destination SGW-C/cnSGW.
2	The destination SGW-C/cnSGW sends Sx Establishment Request (Sxa) to the UPF. The UPF links new Sxa session with the N4 session for uplink packets (Slowpath).
3	The UPF sends Sx Establishment Response (Sxa) to the destination SGW-C/cnSGW.
4	The destination SGW-C/cnSGW sends Create Session Response to the MME.
5	The MME sends Modify Bearer Request (eNodeB F-TEID) to the destination SGW-C/cnSGW.
6	The destination SGW-C/cnSGW sends Sx Modify Request (Sxa) to the UPF.
7	The UPF sends Sx Modify Response (Sxa) to the destination SGW-C/cnSGW.
8	The destination SGW-C/cnSGW sends Modify Bearer Request to the SMF/IWF.

Step	Description
9	The SMF/IWF sends Sx Modify Request (N4) to the UPF. The UPF switches downlink tunnel and links the N4 session with the new Sxa session.
10	The UPF sends Sx Modify Response (N4) to the SMF/IWF.
11	The SMF/IWF sends Modify Bearer Response to the destination SGW-C/cnSGW.
12	The destination SGW-C/cnSGW sends Modify Bearer Response to the MME.
13	The MME sends Delete Session Request to the source SGW-C/cnSGW.
14	The source SGW-C/cnSGW sends Sx Delete Session Request to the UPF.
15	The UPF sends Sx Delete Session Response to the source SGW-C/cnSGW.
16	The source SGW-C/cnSGW sends Delete Session Request to the MME.

WiFi to LTE Handover

The following illustration describes the WiFi to LTE handover call flow.

Table 25: WiFi to LTE Handover Call Flow Description

Step	Description
1	N4 Session is established for WiFi between SMF/IWF and UPF. The MME sends Create Session Request (relocating to 4G) to the SGW-C/cnSGW.
2	The SGW-C/cnSGW sends Sx Session Establishment Request (No F-TEID) to the UPF.
3	IMSI-based lookup is done at SxDemux for SessMgr ID. The UPF sends Sx Establishment Response to the SGW-C/cnSGW.
4	The SGW-C/cnSGW sends Create Session Request to the SMF/IWF.
5	The SMF/IWF sends Sx Modify Request with new Uplink RB PDR to the UPF.
6	The UPF sends Sx Modify Response to the SMF/IWF with new TEID for Uplink data.
7	The SMF/IWF sends Create Session Response to the SGW-C/cnSGW.
8	The SGW-C/cnSGW sends Sx Modify Request to the UPF with updated Uplink FAR.
9	Uplink path is established. The UPF sends Sx Modify Response to the SGW-C/cnSGW.
10	The SGW-C/cnSGW sends Create Session Response to the MME.
11	The MME sends Modify Bearer Request to the SGW-C/cnSGW.
12	The SGW-C/cnSGW sends Sx Modify Request to the UPF with eNodeB TEIDs.
13	The UPF sends Sx Modify Response back to the SGW-C/cnSGW.
14	The SGW-C/cnSGW sends Modify Bearer Request to the SMF/IWF.
15	The SMF/IWF sends tunnel switch Sx Modify Request to the UPF.
16	Collapsed datapath with Sxa session is established at UPF. The UPF sends Sx Modify Response to the SMF/IWF.
17	The SMF/IWF sends Modify Bearer Response to the SGW-C/cnSGW.
18	The SGW-C/cnSGW sends Modify Bearer Response to the MME.

Limitations

The following are the known limitations of the feature:

- If Sxa leg is of one user-plane-service and N4 leg is of another user-plane-service, then datapath won't be collapsed.

- If Sxa leg is under one context and N4 leg is in another context, then datapath can't be collapsed.

- The S-GW part of the **show subscribers user-plane-only full all** CLI output doesn't display ToS-marked Uplink and Downlink packets.

- For the S-GW part of the **show user-plane-service statistics rat all** CLI output, the session statistics for Unknown is incremented, however, the data statistics aren't incremented under RAT-type Unknown.

- Lawful Intercept for S-GW isn't supported.

- S-GW charging isn't supported.

- S-GW bearer inactivity timeout isn't honored, as it's determined by S-GW URR for which processing isn't done for collapsed datapath.

- If S-GW leg of call is locally purged for converged session, then the UPF continues to send data toward eNodeB.

Configuring Converged Datapath

This section describes the CLI commands available in support of the feature.

Enabling Converged Datapath at UPF

Use the following configuration to enable Converged Datapath at UPF.

```
configure
  user-plane converged-mode
  end
```

NOTES:

- **user-plane**: Specifies the UPF related to the configuration.

- **converged-mode**: Specifies the collapse datapath of Sxa and N4 calls.

- By default, the CLI is disabled.

- It is recommended to add the CLI in boot configuration.

Configuring Remote Peers for Sxa and N4

Use the following CLI commands to configure remote peers for Sxa and N4 interface.

```
configure
  control-plane-group group_name
    peer-node-id ipv4-address ipv4_address interface n4
    peer-node-id ipv4-address ipv4_address
    end
```

Configuring User Plane Service for Sxa and N4

Use the following CLI commands to configure UP Service for Sxa and N4 interface.

```
configure
  user-plane-service service_name
    associate gtpu-service gtpu_service upf-ingress
```

```
associate gtpu-service gtpu_service sgw-ingress
associate gtpu-service gtpu_service sgw-egress
associate gtpu-service gtpu_service cp-tunnel
associate sx-service sx_service
associate control-plane-group group_service
end
```

Monitoring and Troubleshooting

Show Commands and/or Outputs

This section provides information about the show CLI commands and/or outputs available in support of the Converged Datapath feature.

show subscribers user-plane-only all

The output of this CLI command is enhanced to display the following fields to indicate if the session is converged or non-converged:

- Converged

- Non Converged

show subscribers user-plane-only full all

The output of this CLI command is enhanced to display the following fields:

- Converged Session

- Converged Peer Callid

show user-plane-service statistics all

The output of this CLI command is enhanced to display the following fields:

- Converged Data Session PDNs:

  - Active

  - Setup

  - Released

CHAPTER 11

Deep Packet Inspection and Inline Services

Feature Summary and Revision History

Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

Revision Details	Release
Support has been added for the following functionality: • IP Readdressing • RTP Dynamic Flow Detection • Rule-matching for Bearer-specific Filters • QUIC IETF implementation	2021.02.0

Revision Details	Release
New L7 protocols have been introduced as part of Deep Packet Inspection (DPI).	2021.01.0
The following EDR attributes have been added for TCP: • SYN and SYN-ACK packet • SYN-ACK and ACK packet	2021.01.0
New DNS attributes have been introduced in EDRs.	2021.01.0
First introduced.	2020.02.0

Feature Description

One of the key product capability of Cisco 5G-UPF is integrated Deep Packet Inspection (DPI) based services. DPI is the examination of layer 7 (L7), which contains Uniform Resource Identifier (URI) information. In some cases, layer 3 (L3) and layer 4 (L4) analyzers that identify a trigger condition are insufficient for billing purposes, so layer 7 (L7) examination is used.

DPI performs packet inspection beyond L4 inspection and is typically deployed for detection of URI information at L7 (for example, DNS, HTTP, HTTPS, RTP, and RTSP URLs).

How it Works

This section describes the following functionality of DPI:

• DSCP Marking of downlink and uplink packets.

• Traffic Readdressing or Redirecting.

DSCP Marking for Downlink and Uplink Packets

Transport-level marking is the process of marking traffic at the UPF with a Differentiated Services Code Point (DSCP) value. The transport-level marking, executed on per-QoS flow, is based on the mapping from the 5QI and optional Allocation and Retention Policy (ARP) configuration from the SMF.

The SMF controls the transport-level marking by providing the DSCP in the ToS (IPv4) or Traffic Class (IPv6) within the "Transport Level Marking" IE in the FAR, that is associated to the PDR matching the traffic to be marked. The UPF performs the transport level marking for the detected traffic and sends the marked packet to the peer entity. The SMF can change the transport-level marking by changing the "Transport Level Marking" IE in the related FAR.

The UPF also supports the inner packet marking in which it marks the tunnel packets. As the 3GPP specification does not determine any specific IE, the UPF uses a private IE named "Inner Packet Marking".

In addition, there is also a provision to copy the DSCP of inner packet to the outer IP header. As the 3GPP specification does not determine any specific IE, the UPF uses a private IE named "Transport Level Marking Options".

Transport Level Marking IE

The "Transport Level Marking" IE type is encoded as shown in the following figure. It indicates the DSCP value for the downlink transport-level marking.
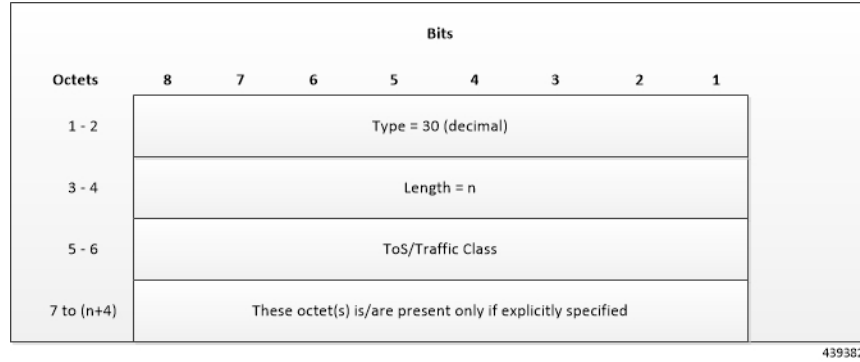
The encoding for Type-of-Service (ToS) or Traffic Class takes place in the form of two octets as an OctetString. The first octet contains the DSCP value in the IPv4 Type-of-Service or the IPv6 Traffic-Class field and the second octet contains the ToS/Traffic Class mask field, which is set to *0xFC*.

Transport Level Marking Options IE

The "Transport Level Marking Options" IE type is encoded as shown in the following figure. The DSCP for downlink transport-level marking is copied from the inner packet.

The Copy-Inner and Copy-Outer flags are present in bit-0 and bit-1 of octet 5. Copy-Outer flag is not used for downlink packets because there is no outer header present in packets coming from ISP. If a Copy-Inner flag is present, then the UPF uses DSCP value from the inner packet to mark the transport-level IP header.

Inner Packet Marking IE

The "Inner Packet Marking" IE type is encoded as shown in the following figure. It indicates the DSCP value for the downlink inner packet marking.

The encoding for ToS/Traffic Class takes place in the form of two octets as an OctetString. The first octet contains the DSCP value in the IPv4 ToS or the IPv6 Traffic Class field and the second octet contains the ToS/Traffic Class mask field, which is set to *0xFC*.

NOTES:

- The original ECN bits in the IP header of User Plane packets do not change after applying transport-level marking or inner packet marking.

- If "Transport Level Marking" IE, "Inner Packet Marking" IE, or both the IEs are associated with uplink FAR, then the following rule applies for uplink packet marking:

  - If "Transport Level Marking" or "Inner Packet Marking" IE is present, its DSCP value is used.

  - If both "Transport Level Marking" and "Inner Packet Marking IE" are present, then the value from "Transport Level Marking" IE is used for uplink packet marking.

Traffic Readdressing or Redirecting

Traffic Redirection is the process of redirecting uplink application traffic to a redirect destination. For example, redirect some HTTP flows to service provisioning page. The redirect destination is provided by the PCF or it is preconfigured in the SMF or in the UPF.

The traffic redirection enforcement is applicable for the SMF or in the UPF if the traffic that the UPF supports subjects to traffic redirection. The UPF reports to the SMF whether it supports traffic redirection enforcement in the UPF through the "UP-Function Features" IE.

To enforce the traffic redirection in the UPF, the SMF takes the following actions:

- Creates the necessary PDRs, if it does not exist, to represent the traffic redirection.

- Creates a FAR with:

  - The "Redirect Information" IE, that includes the redirect destination, if the traffic needs redirection towards a redirect destination that is provided by the SMF. The redirect destination from the SMF prevails over a redirect destination that is preconfigured in the UPF.

    - For HTTP traffic redirection, the Redirection Address Type is set to "URL" and the SMF sets the "Destination Interface" IE in the FAR to "Access" (to forward the HTTP Response message with a status-code indicating redirect). For other types of traffic redirection, the "Destination Interface" IE in the FAR is set to "Core".

• Associates the FAR to the above PDRs of the PFCP session.

Redirect Information IE

Redirect information is encoded as follows:

"Redirect Address Type" indicates the type of the redirect address:

Redirect Address Type	Value (Decimal)
IPv4 address	0
IPv6 address	1
URL	2
SIP URI	3
Spare, for future use.	4–15

The "Redirect Server Address Length" indicates the length of the "Redirect Server Address". The "Redirect Server Address" encoding is in UTF8String format and contains the address of the redirect server (for example, HTTP redirect server, SIP server) with which the end user connects.

☞

Important   In this release, only Redirect Address Type URL is supported for dynamic rule when FAR is associated with URR where quota expires.

Supported Inline Services

Application Detection and Control

The ADC in-line service is mainly used to detect Peer-to-Peer protocols by analyzing traffic. Other popular applications that generate the bulk of Internet traffic like Social Networking and Gaming applications can be detected.

The in-line service known as ADC is also referred as "P2P". Peer to Peer (P2P) is a term used in two slightly different contexts. At a functional level, it means protocols that interact in a peering manner, in contrast to client-server manner. There is no clear differentiation between the function of one node or another. Any node

can function as a client, a server, or both — a protocol may not clearly differentiate between the two. For example, peering exchanges may simultaneously include client and server functionality, sending and receiving information.

Note The ADC support is a licensed feature. Contact your Cisco Account or Support representative for information on how to obtain a license.

QUIC IETF Implementation

In the current framework, Deep Packet Inspection (DPI) is done for every packet in a flow when it reaches the plugin. The DPI is done by analyzing the packets and extracting deterministic patterns. The DPI is done in-order to detect the application and to classify its subtype. Plugin excludes the flow after the DPI. The flow is offloaded after the detection. As part of QUIC IETF, the initial QUIC handshake packets (Client/Server Hello) are encrypted over the network. Hence, there are no deterministic patterns available for detection of the application. Support is added in p2p plugin to decrypt and obtain the SNI (Server Name Indication) for detection.

Configuring QUIC IETF

Use the following configuration to enable or disable the QUIC IETF decryption.

```
configure
   active-charging service acs_service_name
      p2p-detection debug-param protocol-param p2p_quic_ietf_decrypt 1
      end
```

Note
- By default, the CLI is disabled and there's minimal impact on the performance due to TLS decryption.

- Runtime change of configuration doesn't impact the existing flow. Change is applicable only for new flow.

Statistics

show user-plane-service statistics analyzer name p2p

Use this show CLI command to determine the packets that are analyzed for QUIC and application.

Content Filtering

Content Filtering is an in-line service available for 3GPP and 3GPP2 networks to filter HTTP requests from mobile subscribers based on the URLs in the requests. This enables operators to filter and control the content that an individual subscriber can access, so that subscribers are inadvertently not exposed to universally unacceptable content and/or content inappropriate as per the subscribers' preferences.

Content Filtering Configuration

Use the following additional configuration to enable the content filtering:

```
configure
   require user-plane content-filtering
      content-filtering category database directory path path_address
      content-filtering category database max-version version_number
      end
```

✎

Note   The above configuration must be configured on the UPF, during boot time, to enable Content Filtering. Defining the above configuration post the User Plane configuration will lead to errors and inconsistencies.

Show Commands Input and/or Outputs

show subscribers user-plane-only callid *call_id* **full all**

SMF provides Content Filtering Policy ID in the Session Establishment/Modification Request. The following fields are displayed in support of this feature:

- SUBSCRIBER PARAMS

    - Content Filtering Policy ID

DNS Snooping

Charging

The charging of DNS Snooping takes place at SM-P.

Rule Definitions

Use the following CLI commands for specifying the rule definition hostnames (domain-names) and part of the host names.

ruledef *<ruledef_name>*
       ip [server-domain-name {contains|=|ends-with|starts-with} *<url_string>*]

       ip [server-domain-name {contains|=|ends-with|starts-with} *<url_string>*]

         multi-line-OR enabled

Use the no version of this CLI to delete the ruleline for ip server- domain-name.

ruledef *<ruledef_name>*
  no ip [server-domain-name {contains|=|ends-with|starts-with} *<url_string>*]

   exit

Use the following CLI for configurable timer of DNS entries at ECS level.

configure
       active-charging service *service_name*
          ip dns-resolved-entries timeout *<value_secs>*
          end

Whenever the ruledef containing the ip server-domain-name keyword is defined and used in rulebase, the ip-table is created per rulebase per instance.

Rule Matching

The functionality remains the same as the non-CUPS architecture.

Show CLIs

Use the following CLIs to check the table for DNS IP entries:**show user-plane-service [statistics dns-learnt-ip-addresses {summary | sessmgr instance <id> |all [verbose] }]**

Bulkstats

The following bulkstats are available in support of DNS Snooping feature:

- ecs-dns-learnt-ipv4-entries

- ecs-dns-flushed-ipv4-entries

- ecs-dns-replaced-ipv4-entries

- ecs-dns-overflown-ipv4-entries

- ecs-dns-learnt-ipv6-entries

- ecs-dns-flushed-ipv6-entries

- ecs-dns-replaced-ipv6-entries

- ecs-dns-overflown-ipv6-entries

The above bulkstats are added in the ECS schema same as in the non-CUPS architecture.

Note   The SNMP Trap generation commands are not supported in CUPS DNS snooping feature.

Event Data Records

Feature Description

Event Data Records (EDR) are usage records with support to configure content information, format, and generation triggers by the system administrative user.

When a flow is terminated, the UPF generates EDRs with detail information of the terminated flow.

How It Works

EDRs are generated from User Plane on flow termination. During call setup and call modification, all call-specific attributes required for EDR generation is sent from SMF to UPF as part of the "Subscriber Params" IE within the Sx Establishment/Modification request messages.

On flow termination, the charging counters are fetched from VPP. All configured call-level attributes in the EDR format configuration along with the charging/volume counter attributes is sent to the CDRMOD proclet. This proclet writes these records to a file/disk, which is transferred to a configured external server.

TCP Fast Open

TCP Fast Open (TFO) is an extension to speed up the opening of successive TCP connections between two endpoints. It works by using a TFO cookie (a TCP option), which is a cryptographic cookie stored on the client and set upon the initial connection with the server. When the client reconnects, it sends the initial SYN packet along with the TFO cookie data to authenticate itself. If successful, the server starts sending data to the client even before the reception of the final ACK packet of the three-way handshake. Due to this, the difference between following packets are recorded to calculate and record time difference between control packets of TCP flow in EDR:

- SYN and SYN-ACK packet

- SYN-ACK and ACK packet

For information about rule variables that are added to capture the information in EDRs, refer *Configuring Additional TCP Fields* section.

Transaction Complete EDR

Transaction Complete EDRs are generated for HTTP EDRs when an HTTP transaction is completed. On completion, the charging counters are fetched from VPP. All configured call-level attributes in the EDR format configuration along with the charging/volume counter attributes is sent to the CDRMOD proclet. This proclet writes these records to a file/disk, which is transferred to a configured external server.

The following EDR attributes are supported:

- attribute sn-start-time

- attribute sn-end-time

- attribute sn-start-time format MM/DD/YYYY-HH:MM:SS:sss

- attribute sn-end-time format MM/DD/YYYY-HH:MM:SS:sss

- attribute radius-calling-station-id

- attribute radius-called-station-id

- rule-variable bearer 3gpp imsi

- rule-variable bearer 3gpp imei

- rule-variable bearer 3gpp rat-type

- rule-variable bearer 3gpp user-location-information

- rule-variable ip subscriber-ip-address

- rule-variable ip dst-address

- attribute sn-ruledef-name

- attribute sn-subscriber-port

- attribute sn-server-port

- attribute sn-app-protocol

- attribute sn-volume-amt ip bytes uplink

- attribute sn-volume-amt ip bytes downlink

- attribute sn-flow-start-time format seconds

- attribute sn-flow-end-time format seconds

- attribute sn-volume-amt ip pkts uplink

- attribute sn-volume-amt ip pkts downlink

- attribute sn-direction

- rule-variable traffic-type

- rule-variable p2p protocol

- rule-variable p2p app-identifier tls-cname

- rule-variable p2p app-identifier tls-sni

- rule-variable p2p app-identifier quic-sni

- rule-variable bearer 3gpp sgsn-address

- attribute sn-rulebase

- attribute sn-charging-action

- rule-variable flow tethered-ip-ttl

- rule-variable flow ttl

- rule-variable flow ip-control-param

- rule-variable bearer qci

- rule-variable tcp flag

- rule-variable ip server-ip-address

- attribute sn-flow-id

- attribute sn-closure-reason

- attribute sn-duration

- rule-variable ip src-address

- rule-variable ip protocol

- attribute sn-charge-volume ip bytes uplink

- attribute sn-charge-volume ip bytes downlink

The following HTTP EDR attributes are supported:

- rule-variable http url length 2000

- rule-variable http request method

- rule-variable http content type

- rule-variable http user-agent length 255

- rule-variable http reply code

- rule-variable http referer

- rule-variable http host

- rule-variable http cookie

- rule-variable http header-length

- attribute transaction-uplink-bytes

- attribute transaction-downlink-bytes

The following DNS EDR attributes are supported:

- rule-variable dns answer-ip-list

- rule-variable dns answer-name

- rule-variable dns previous-state

- rule-variable dns query-name

- rule-variable dns query-type

- rule-variable dns return-code

- rule-variable dns state

- rule-variable dns tid

Limitations

The EDR feature in UPF has the following limitations:

- EDR will be generated only for flow end condition: Idle timeout, HAGR, normal flow termination, and during the end of a session.

- Charging-Action based EDR configuration is not supported.

- Reporting EDRs are not supported.

Configuring Event Data Records

Configuring EDRs on UPF

Use the following configuration to configure EDRs on UPF:

```
active-charging service service_name
  rulebase rulebase_name
    flow end-condition { timeout | normal-end-signaling | session-end |
 hagr } charging-edr charging_edr_format_name
    edr transaction-complete { http | dns } charging-edr
charging_edr_format_name
```

```
      exit
      edr-format format_name
         attribute attribute_name
            end
```

NOTES:

- **flow end-condition**: This command allows you to configure the end condition of the session flows related to a user session and triggers EDR generation.

- **timeout**: Creates an EDR with the specified EDR format whenever a flow ends due to a timeout condition.

- **normal-end-signaling**: Creates an EDR with the specified EDR format whenever flow end is signaled normally. For example, detecting FIN and ACK for a TCP flow, and create an EDR for the flow using the specified EDR format.

- **session-end**: Creates an EDR with the specified EDR format whenever a subscriber session ends. By this option session manager creates an EDR with the specified format name for every flow that has had any activity since last EDR was created for the flow on session end.

- **charging-edr** *charging_edr_format_name*: Specifies the charging EDR format.

- **hagr**: Creates an EDR with the specified EDR format whenever a flow is terminated due to Inter-chassis Session Recovery action.

- **http**: Specifies HTTP protocol related configuration.

- **dns**: Specifies DNS protocol related configuration.

Configuration to Enable EDR Module

Use the following configuration to enable EDR module.

```
configure
   context context_name
      edr-module active-charging-service
         end
```

Configuring Additional TCP Fields

Prior to using the following CLI commands to configure additional TCP fields in the EDR, ensure that all the other EDR configurations are present.

```
configure
   active-charging service service_name
      edr-format edr_format_name
         rule-variable tcp syn-synack-rtt priority priority_value
         rule-variable tcp synack-ack-rtt priority priority_value
            end
```

Monitoring and Troubleshooting

show user-plane-service statistics rulebase name *rulebase_name*

The following fields are displayed in support of this feature:

- Rulebase Name

- EDRs

- Charge Volume

  - Uplink Pkts

  - Uplink Bytes

  - Downlink Pkts

  - Downlink Bytes

- Charging EDRs

  - Total Charging EDRs generated

  - EDRs generated for handoff

  - EDRs generated for timeout

  - EDRs generated for normal-end-signaling

  - EDRs generated for session end

  - EDRs generated for rule match

  - EDRs generated for hagr

  - EDRs generated for flow-end content-filtering

  - EDRs generated for flow-end url-blacklisting

  - EDRs generated for content-filtering

  - EDRs generated for url-blacklisting

  - EDRs generated for any-error packets

  - EDRs generated for firewall deny rule match

  - EDRs generated for transaction completion

  - EDRs generated for voip call end

  - EDRs generated for dcca failure handling

  - EDRs generated for TCP optimization on

  - EDRs generated for tethering signature change

  - EDRs generated for interim interval

  - Total Flow-Overflow EDRs

  - Total zero-byte EDRs suppressed

show user-plane-service edr-format all

The following fields are displayed in support of Additional TCP Fields in EDR feature:

- Service Name

> • EDR Format Name
>
> > • rule-variable tcp syn-synack-rtt priority 1
> >
> > • rule-variable tcp synack-ack-rtt priority 2

Flow Idle Timeout Randomization

Every two seconds, the Session Manager polls the time of the latest packet from Session Manager instance, or the fastpath stream to determine idle flows. Short length flows become idle quickly as they are short due to the lesser number of packets and are short lived, within 5–10 seconds. As a result, large number of idle flows must be deleted due to the timeout at the given polling cycle of two seconds. Deletion of idle flows is CPU intensive as it involves statistics reconciliation, EDR generation, and fast path stream deletion. You can accomodate more flows with this feature as the short lived flows get cleared aggressively.

Configuring Flow Idle Timeout Randomization in ACS

Use the following configuration to randomize the idle timeout flow.

```
configure
active-charging service service_name
  idle-timeout randomize-range range
    { default | no } idle-timeout randomize-range
      end
```

NOTES:

- **idle timeout**: Specifies the maximum duration that a flow can remain idle for, in seconds. Seconds must be an integer from 5 through 30. The flow will then be terminated based on the random value.

- **randomize-range**: Specifies the range of a period of time in seconds. The idle timeout applied, will be different for each flow.

  For example,

  ```
  idle-timeout randomize-range 20
  ```

  An integer random number is generated from 0 through 20. This number is added to the configured idle timeout value to check if the flow has become idle in the two second timer processing. If the idle timeout configured is 60 seconds, the actual timeout that is applied to each flow will be random in the range between *60 + 20* seconds causing staggered flow deletion.

- **no**: Disables the idle timeout randomization. This command is disabled by default.

- **default**: Configures the idle timeout randomization command with its default setting in seconds. Seconds must be an integer from 0 through 30. Default range is 0–30 seconds.

  For example, **default idle-timeout randomize-range** is equal to **idle-timeout randomize-range 30**.

HTTP URL Filtering

The HTTP URL Filtering feature simplifies rule definitions used for URL detection.

The HTTP request packet can have a proxy (prefixed) URL and an actual URL. If a proxy URL is found in the HTTP request packet, the HTTP URL Filtering feature truncates this URL from the parsed information and only the actual URL is used for rule matching and Event Data Records (EDR) generation.

Configuring the HTTP URL Filtering Feature

This section describes how to configure the HTTP URL Filtering feature.

Configuring Group of Prefixed URLs

To configure the group of prefixed URLs, use the following CLI commands:

```
configure
   active-charging service ecs_service_name
      group-of-prefixed-urls prefixed_urls_group_name
      end
```

Configuring URLs in the Group of Prefixed URLs

To configure URLs to be filtered in the group of prefixed URLs, use the following CLI commands:

```
configure
   active-charging service ecs_service_name
      group-of-prefixed-urls prefixed_urls_group_name
         prefixed-url url_1
         ...
         prefixed-url url_10
         end
```

Enabling the Group of Prefixed URLs in Rulebase

To enable the group of prefixed URLs in rulebase for processing prefixed URLs, use the following CLI commands:

```
configure
   active-charging service ecs_service_name
      rulebase rulebase_name
         url-preprocessing bypass group-of-prefixed-urls
prefixed_urls_group_name_1
         ...
         url-preprocessing bypass group-of-prefixed-urls
prefixed_urls_group_name_64
         end
```

This configuration on the control plane chassis will be pushed to the user plane with a PFD message for "group-of-prefixed-urls" and "rulebase-url-preprocessing"separately.

The group of prefixed URLs has the list of proxy URLs, which must be truncated. The rulebase contains multiple group of prefixed urls, which must be filtered. Charging ruledefs contain rules for actual URLs that must be searched after truncating URLs in the group of prefixed URLs.

✏️

Note
- Each group of prefixed URLs can have a maximum of ten prefixed URLs.

- A maximum of 64 group of prefixed URLs can be created and configured.

Show Commands

show user-plane-service group-of-prefixed-urls all | name *group_name*

This show command can be used on the user plane to verify whether the group of prefixed URLs are pushed or not. The output of this command is as follows:

- Name of the group of prefixed URLs

- Prefixed URLs

- Total number of prefixed URLs found

show user-plane-service rulebase name *rbase_name*

This show command can be used on the user plane to check whether the group of prefixed URLs is configured in rulebase or not. The output of this command is as follows:

- Name of rulebase

- Name of the groups of prefixed Urls for URL pre-processing

show user-plane-service statistics analyzer name http

The output of this command is as follows:

- Total HTTP Sessions

- Current HTTP Sessions

- Total Uplink Bytes

- Total Downlink Bytes

- Total Uplink Pkts

- Total Downlink Pkts

- Uplink Bytes Retrans

- Downlink Bytes Retrans

- Uplink Pkts Retrans

- Downlink Pkts Retrans

- Total Request Succeed

- Total Request Failed

- GET Requests

- POST Requests

- CONNECT Requests

- PUT requests

- HEAD requests

- Websocket Flows

- Invalid packets

- Wrong FSM packets

- Unknown request method

- Pipeline overflow requests

- Corrupt request packets

- Corrupt response packets

- Unhandled request packets

- Unhandled response packets

- Partial HTTP Header Anomaly prevented

- New requests on closed connection

- Memory allocation failures

- Packets after permanent failure

- Prefixed Urls Bypassed

- FastPath Statistics

- Total FP Flows

- Uplink (Total FP Pkts)

- Downlink (Total FP Pkts)

- Uplink (Total FP Bytes)

- Downlink (Total FP Bytes)

Note   Prefixed URLs Bypassed counter has been added in http analyzer stats as a performance measurement to show the number of truncated prefixed URLs.

IP Readdressing

The IP Readdressing feature enables redirecting unknown gateway traffic based on the destination IP address of the packets to known/trusted gateways.

IP Readdressing is configured in the flow action defined in a charging action. IP readdressing works for traffic that matches particular ruledef, and hence the charging action. IP readdressing is applicable to both uplink and downlink traffic. In the Enhanced Charging Subsystem, uplink packets are modified after packet inspection, rule matching, and so on, where the destination IP/port is determined, and replaced with the readdress IP/port just before they are sent out. Downlink packets (containing the readdressed IP/port) are modified as soon as they are received, before the packet inspection, where the source IP/port is replaced with the original server IP/port number.

For one flow from an MS, if one packet is re-addressed, then all the packets in that flow is re-addressed to the same server. Features like DPI and rule-matching remain unaffected. Each IP address and port combination are defined as a ruledef.

In case of IP fragmentation, packets with successful IP re-assembly are re-addressed. However, IP fragmentation failure packets are not re-addressed.

There are two different approaches for the readdress server selection, in case, server-list is configured under charging-action.

- **round-robin**: In round-robin approach, server selection happens in round-robin manner for every new flow. In round-robin, only active servers in the list is considered for selection.

- **hierarchy based approach**: In hierarchy based approach, servers are tagged as primary, secondary, tertiary, and so on depending on the order they are defined in the readdress server-list. All flows are readdressed to the primary server till it is up and running. If Primary server goes down, then flows are readdressed to secondary server and the same logic goes on. Once primary server is active then flows will switch back to primary server for readdressing.

An additional CLI is provided that enables user to select from hierarchy or round-robin approach for server selection. Both round-robin and hierarchy based server selection approaches are applicable for both IPv4 and IPv6 based servers.

Configuring IP Readdressing

Readdressing of packets based on the destination IP address of the packets enables redirecting unknown gateway traffic to known/trusted gateways. This is implemented by configuring the re-address server in the charging action.

To configure the IP Readdressing feature, use the following configuration:

```
configure
   active-charging service acs_service_name
      charging-action charging_action_name
         flow action readdress server ipv4_address/ipv6_address [
discard-on-failure ]  [ dns-proxy-bypass ]  [ port port_number [
discard-on-failure ] [ dns-proxy-bypass ] ]
         end
```

To configure the IP Readdressing feature when the readdress server-list is defined under charging-action, use the following configuration:

```
configure
   active-charging service acs_service_name
      charging-action charging_action_name
         flow action readdress server-list server_list_name [ hierarchy ] [
round-robin ] [ dns-proxy-bypass ] [ discard-on-failure ]
         end
```

Following is the sample server-list configuration:

```
readdress-server-list DRE
consecutive-failures 1
response-timeout 10000
server 1.1.1.1 port 53
server 1.1.1.2 port 53
server 2001:420:54fe::1019 port 53
server 2001:420:54fe::1039 port 53
server 2001:420:54fe::1049 port 53
server 2001:420:54fe::1059 port 53
server 192.1.1.1 port 8080
#exit
```

![Note icon]

| Note | A maximum of 10 servers can be configured in the list and a maximum of 10 lists can be configured in active-charging service. |

Show Commands

This section provides information about the show CLI commands available in support of IP Readdressing feature.

CLI Command	Description
show user-plane-service readdress-server-list statistics all **show user-plane-service readdress-server-list statistics instance** *instance* **show user-plane-service readdress-server-list statistics name** *name*	Use these show CLI commands to display the readdress server list statistics.
show user-plane-service readdress-server-list statistics	

Use the **clear user-plane-service readdress-server-list statistics all** CLI to clear the readdress server list statistics.

Use the **show user-plane-service statistics charging-action all** CLI to check the "flows readdressed" counter.

L7 Protocol

The following L7 protocols are supported as part of DPI:

- DNS
- FTP
- HTTP
- HTTPS
- RTP/RTSP
- SIP

DNS

The UPF supports DNS protocol as part of L7 Analyzer.

FTP

The UPF supports FTP proocol as part of L7 Analyzer.

HTTP

On completion of HTTP Request/Response, the uplink/downlink data packets are offloaded to VPP in the following cases:

- Content-Length – Volume-based offloading is supported for methods like GET and POST. The HTTP flow with chunk-encoding data transfer mechanism does not get offloaded irrespective of the method defined in HTTP. If the stream is offloaded based on content-length, then the stream on the other end will also get offloaded until the former is not onloaded.

- CONNECT Method– The method where both uplink and downlink streams are offloaded after flow is upgraded to CONNECT.

- WebSocket Method– After the flow is classified as WebSocket protocol, both uplink and downlink streams are offloaded.

- The streams are onloaded back in either of the following cases:

  - FIN packet received.

  - Content-length is breached.

  - PDN update.

Header Parsing

Only the header fields defined in ruledefs, which are included in rulebase, are parsed. Or, in case of features like x-header, redirection is configured which has dependencies on some of the HTTP header fields.

HTTP Charging

- Complete packets are charged.

- Partial packets are charged on completion. Packet completing the partial packet is also charged.

- Concatenated packets are charged.

- Delay Charging is enabled – Control packets are charged against application-based rule, depending on delay charging CLI configuration.

- Response-based charging is enabled – After HTTP request's response is received, then the HTTP request is charged against response rule's CA.

X-Header Parsing and Rule-Matching

Ruledefs with x-header rule-lines are parsed and matched.

WebSocket

Involves charging of subsequent packets of the flow after HTTP GET request as per the HTTP request, if the HTTP flow is upgraded to be a websocket flow.

Response-Based TRM

Transactional Rule Matching is engaged after HTTP response packet is received.

URL-Based Redirection

For flow action redirect-url, encrypt is not supported. Currently, the following dynamic fields are supported:

- #HTTP.URI#
- #HTTP.HOST#
- #HTTP.URL#
- #ACSMGR_BEARER_CALLED_STATION_ID#
- #RULEBASE#
- #RTSP.URI#

X-Header Insertion

X-header Insertion is supported in HTTP Requests. Note that:

- Flows, for which X-header is inserted in a packet, are not offloaded.
- With X-header configuration, all TCP OOO packets irrespective of transmit order CLI, will be buffered and sent out after reordering.

Limitation

- X-Header Spoofing is not supported.
- X-Header Insertion in Response packet is not supported.
- X-Header Encryption with RSA and RC4MD5 is supported but not supported with AES.
- Monitor protocol for X-Header is not supported.
- Following X-Header fields insertion is not supported in a packet:
  - QoS
  - UIDH
  - Customer ID
  - Hash Value
  - Time of the Day
  - RADIUS String
  - Session-Id
  - Congestion Level
  - User-Profile

HTTPS

The UPF supports HTTPS protocol as part of L7 Analyzer.

RTP/RTSP

The UPF supports RTP and RTSP protocols as part of L7 Analyzer.

SIP

Session Initiation Protocol is an IETF-defined signaling protocol widely used for controlling communication sessions such as voice and video calls over Internet Protocol (IP). The protocol can be used for creating, modifying and terminating two-party (unicast) or multiparty (multicast) sessions. Sessions may consist of one or several media streams.

The UPF supports SIP as part of L7 Analyzer.

Monitoring and Troubleshooting

This section provides information regarding the CLI command available in support of 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.

DNS

Use the following CLI command to get statistics related to DNS:

show user-plane-service statistics analyzer name dns

FTP

Use the following CLI command to get statistics related to FTP:

show user-plane-service statistics analyzer name ftp

HTTP

Use the following CLI command to get statistics related to HTTP:

show user-plane-service statistics analyzer name http

HTTPS

Use the following CLI command to get statistics related to HTTPS:

show user-plane-service statistics analyzer name secure-http

RTP

Use the following CLI command to get statistics related to RTP:

show user-plane-service statistics analyzer name rtp

RTSP

Use the following CLI commands to get statistics related to RTSP:

- **show user-plane-service statistics analyzer name rtsp**

> • `show user-plane-service statistics analyzer name rtsp verbose`

SIP

Use the following CLI command to get statistics related to SIP:

`show user-plane-service statistics analyzer name sip`

Tethering Detection

Feature Description

Tethering refers to the use of a mobile smartphone as a USB dongle/modem to provide Internet connectivity to PC devices (laptops, PDAs, tablets, and so on) running on the smartphone's data plan. Typically, for smartphone users, most operators have in place an unlimited data plan, the usage of which is intended to be from the smartphone as a mobile device. However, some subscribers use the low cost / unlimited usage data plan to provide Internet connectivity to their laptops in places where normal Internet connection through broadband/WiFi may be costly, unavailable, or insecure.

The Tethering Detection feature enables detection of subscriber data traffic originating from PC devices tethered to mobile smartphones, and also provides effective reporting to enable service providers take business decisions on how to manage such usage and to bill subscribers accordingly. Tethering Detection is supported for IPv4 (TCP)and IPv6 traffic flows.

In this release, IP-TTL based tethering is supported. This feature is configurable at the rulebase level and is applicable on all flows for all subscribers having IP-TTL configuration within the rulebase.

Configuring Tethering Support

This section describes how to configure the Tethering Support feature.

Configuring the Tethering Support feature involves the following steps:

> • Rulebase Configuration for Tethering

> • Ruledef Configuration for Tethering

> • EDR Configuration for Tethering

Rulebase Configuration for Tethering

Use the following commands to configure the rulebase parameters for tethering.

```
configure
   active-charging service service_name
      rulebase rulebase_name
         tethering-detection ip-ttl valuettl_value
         end
```

NOTES:

> • **tethering-detection**: This command allows you to enable/disable the Tethering Detection feature for the current rulebase, and specifies the database to use.

- **ip-ttl value** *ttl_value*: Specifies to perform tethering detection using IP-TTL configuration. *ttl_value* must be an integer from 1 through 255 to configure TTL values for tethered flows.

Ruledef Configuration for Tethering

Use the following commands to configure ruldef parameters for tethering.

```
configure
  active-charging service service_name
    ruledef ruledef_name
      ip any-match operator_condition
      tethering-detection ip-ttl flow-tethered
      end
```

NOTES:

- **ip any-match** *operator_condition*: This command allows you to define rule expressions to match all IPv4/IPv6 packets.

- **ip-ttl**: Specifies to select flows that were tethered or non-tethered as per IP-TTL values.

- **flow-tethered**: Specifies to match if tethering is detected on flow.

EDR Configuration for Tethering

Use the following commands to configure EDR for tethering:

```
configure
  active-charging service service_name
    edr-format format_name
      rule-variable flow tethered-ip-ttl priority priority_value
      rule-variable flow ttl priority priority_value
      end
```

NOTES:

- **edr-format** *format_name*: configures EDR formats.

- **flow**: Configures the flow related fileds in an EDR.

- **tethered-ip-ttl**: IP-TTL based tethering detected on flow.

- **ttl**: Time To Live/Max hops value received in the first packet of the flow.

Monitoring and Troubleshooting

This section provides information regarding the CLI command available in support of 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 user-plane-service statistics tethering-detection

The following fields are displayed in support of this feature:

• Current Tethered Subscribers

• Total Tethered Subscribers

• Total flows scanned

• Total Tethered flows detected

• Total Tethered flows recovered

• Total flows bypassed for scanning

• Tethering Detection Statistics (ip-ttl)

  • Total flows scanned

  • Tethered flows detected

  • Tethered uplink packets

  • Tethered downlink packets

show user-plane-service statistics rulebase name rulebase_name

The following fields are displayed in support of this feature:

• Tethering Detection (ip-ttl)

  • Total flows scanned

  • Tethered flows detected

  • Tethered uplink packets

  • Tethered downlink packets

RTP Dynamic Flow Detection

The **rtp dynamic-flow-detection** CLI command, under the ACS Rulebase Configuration mode, enables the Real Time Streaming Protocol (RTSP) and Session Description Protocol (SDP) analyzers to detect the child RTP and RTCP flows. If you configure the RTSP/SIP and SDP analyzers, and **rtp dynamic-flow-detection** CLI is present, then there's no need for configuring RTP/RTCP explicitly. With the **rtp dynamic-flow-detection** CLI command, the child RTP or RTCP flows get corelated to their parent RTSP/SIP-SDP flows.

Once the parent flow (RTSP/SIP-SDP) gets cleared, the child RTP/RTCP flows also gets cleared. In the absence of this CLI, the L7 layer analysis for RTP and RTCP needs a separate analyzer configuration. There's no correlation of RTP/RTCP flows to RTSP/SIP-SDP flow.

Rule-matching for Bearer-specific Filters

The Rule-matching for Bearer-specific Filters functionality includes:

• IMSI-based rules are matched as per the subscribers IMSI.

• APN-based rules allow you to define rule expressions to match Access Point Name (APN) of the bearer flow.

• RAT-Type allows you to define rule expressions to match Radio Access Technology (RAT) in the bearer flow.

Configuring IMSI Pool

Use the following CLI commands to configure the IMSI pool.

```
configure
   active-charging service service_name
      imsi-pool pool_name
         imsi { imsi_number | range start_imsi to end_imsi }
```

The imsi-pool can contain either IMSI value or range of IMSI.

Configuring Rule-line ACS Ruledef

Use the following CLI commands to configure rule-line in ACS Ruledef Configuration mode.

```
configure
   active-charging service service_name
      ruledef ruledef_name
         bearer 3gpp imsi { = imsi_value } | { range imsi-pool pool_name }
         bearer 3gpp apn operator apn_name
         bearer 3gpp rat-type operator rat_type
```

IMSI range can be configured in a rule with the help of IMSI pool.

bearer 3gpp rat-type *operator rat_type*:

- *operator* must be one of the following:

    - **!=** : Does not equal

    - **=** : Equals

- **NOTE**: In this release, **wlan** is the qualified *rat_type*.

Configuring HTTP Content-Type

Use the following CLI commands to define rule expressions to match value in HTTP Content-Type entity-header field.

```
configure
   active-charging service service_name
      ruledef ruledef_name
         http content type [ case-sensitive ] operator content_type
```

Show CLIs

Use the following CLI on UPF to see information about IMSI pool that is configured in a service:**show user-plane-service imsi-pool name** *pool_name*

URL Blacklisting

Feature Description

The URL blacklisting feature regulates the subscriber's access to view or download content from websites whose URL or URI has been blacklisted. It uses a database that records a list of URLs that indicates if the detected URL is categorized to be blocked or not.

How it Works

To enable the URL blacklisting feature on UPF, URL blacklisting database should be present with a name "optblk.bin" under flash, or SFTP or under its sub-directory. This database directory path needs to be configured on user-plane, after user-plane services are brought up.

HTTP Analyzer must be enabled for URL blacklisting. The HTTP analyzer extracts URL information from the incoming HTTP request data packet. Extracted URL content is compared with the URL Blacklisting database. When the URL of incoming HTTP data packet matches with the database URL entry, that URL is treated as blacklisted URL and one of the following actions takes place on that HTTP packet:

- Termination of flow

- Packet is discarded

The URL backlisting configurations must be configured under Rulebase configuration in Active Charging Service. Additionally, two URL blacklisting methods – Exact and Generic, are supported at Active Charging Service-level configuration.

☞

Important   Blacklisting database(s) are provided by – Internet Watch Foundation (IWF) and National Center for Missing and Exploited Children (NCMEC). The UPF always receives the blacklisting database in Optimized Format.

URL Blacklisting Database Upgrade

URL database upgrade is supported in following two ways:

- Timer-based upgrade or Auto upgrade

- CLI-based upgrade or Manual upgrade

Timer-based or Auto-upgrade

After the database is loaded on the chassis for the first time, a timer, for a duration of 5 minutes, is started. This process is started to auto upgrade the database.

If at the expiry of the timer, a valid database with higher version is available at the directory path, then database upgrade procedure is initiated, and a newer version of the database is loaded on the UPF.

To upgrade a URL blacklisting database, a higher version of valid URL Blacklisting database with name "optblk_f.bin" should be present at same directory as that of current database "optblk.bin".

After the database is upgraded successfully, the earlier "optblk.bin" file gets renamed as "optblk_0.bin" and "optblk_f.bin" file gets renamed as "optblk.bin". Here, "optblk_0.bin" file is treated as a backup file of older database.

If an additional upgrade is performed, then "optblk_0.bin" file will be renamed as "optblk_1.bin" file and current "optblk.bin" will get renamed as "optblk_0.bin", and so on.

See the *Loading URL Blacklisting Database on UPF* section to configure the number of backup files to be stored in the database.

CLI-based or Manual Upgrade

See the *Upgrading the URL Blacklisting Database* section to upgrade the current database to a newer version.

Configuring URL Blacklisting

Loading URL Blacklisting Database on UPF

Use the following configuration to load URL blacklisting database on UPF.

```
configure
   url-blacklisting database directory path database_directory_path
   url-blacklisting database max-versions max_version_value
   end
```

NOTES:

- **database directory path**: Configures the database directory path.

  The *database_directory_path* is a string of size 1 to 255.

- **max-versions**: Configures the maximum database upgrade versions.

  The *max_version_value* is an integer from 0 to 3.

Upgrading the URL Blacklisting Database

Use this configuration to manually upgrade the URL blacklisting database.

```
upgrade url-blacklisting database
end
```

Configuration to Enable URL Blacklisting

Use the following configuration to enable URL blacklisting feature on UPF.

```
configure
   active-charging service service_name
      url-blacklisting match-method [ exact | generic ]
      rulebase rulebase_name
         url-blacklisting action [ discard | terminate-flow ]
         end
```

NOTES:

- **match-method [exact | generic]**: Specifies the match method used for URL blacklisting.

  - **exact**: URL Blacklisting perform an exact-match of URL.

  - **generic**: URL Blacklisting perform generic-match of URL.

- **url-blacklisting action [discard | terminate-flow]**:

- **discard**: Discards the HTTP packet received.

- **terminate-flow**: Terminates the flow of the HTTP packet received.

Monitoring and Troubleshooting

This section provides information regarding the CLI command available in support of 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 user-plane-service url-blacklisting database

The following fields are displayed in support of this feature:

- URL Blacklisting Static Rating Databases:
  - Last Upgrade Status
  - Path
    - Database Status
    - Number of URLs in DB
    - Type
    - Version
    - Creation Time
    - Hostname
    - Comment
    - Last Access Time
    - Last Modification Time
    - Last Status Change Time

show user-plane-service url-blacklisting database url database_directory_path

The following fields are displayed in support of this feature:

- URL Blacklisting Static Rating Databases:
  - Last Upgrade Status
  - Path
    - Database Status
    - Number of URLs in DB
    - Type

- Version

- Creation Time

- Hostname

- Comment

- Last Access Time

- Last Modification Time

- Last Status Change Time

show user-plane-service url-blacklisting database facility sessmgr all

The following fields are displayed in support of this feature:

- URL-Blacklisting SessMgr Instance Based Database Configuration

  - SessMgr Instance

  - BL DB Load Status

  - BL DB Version

  - Number of URLs

  - Checksum

show user-plane-service rulebase name rulebase_name

The following fields are displayed in support of this feature:

- URL-Blacklisting Action

- URL-Blacklisting Content ID

show user-plane-service inline-services info

The following fields are displayed in support of this feature:

- URL-Blacklisting: Enabled

  - URL-Blacklisting Match-method: Generic

show user-plane-service inline-services url-blacklisting statistics

The following are displayed in support of this feature:

- Cumulative URL-Blacklisting Statistics

  - Blacklisted URL hits

  - Blacklisted URL misses

  - Total rulebases matched

show user-plane-service inline-services url-blacklisting statistics rulebase name rulebase_name

The following fields are displayed in support of this feature:

- Rulebase Name
    - URL-Blacklisting Statistics
    - Blacklisted URL hits
    - Blacklisted URL misses
- Total rulebases matched

Configuring the Static and Pre-Defined Rules

This section describes how to configure the static and pre-defined rules under the charging action configuration.

```
configure
    active-charging service service_name
        charging-action charging_action_name
            flow action { discard [ downlink | uplink ] | redirect-url
redirect_url | terminate-flow }
            end
```

NOTES:

- **flow action**: Specifies the action to take on packets that match rule definitions.
    - **discard [downlink | uplink]**: Specifies to discard downlink or uplink packets.
    - **redirect-url** *redirect_url*: Specifies the URL to be redirected. For example, http://search.com/subtarg=#HTTP.URL#
    - **terminate-flow**: Specifies to terminate the flow.
- For redirect-url, configure HTTP analyzer under rulebase. Example:

```
route priority 70 ruledef http-port analyzer HTTP
   ruledef http-port
   tcp either-port = 80
   rule-application routing
   exit
```

Configuring ACS Ruledef for L7 Protocols for DPI

A ruledef represents a set of matching conditions across multiple L3 – L7 protocol based on protocol fields and state information. Each ruledef can be used across multiple rulebases within the active charging service.

Note　In UPF, if rule-line addition or deletion inside a ruledef is done during active calls and data flows, then this configuration change is not applied for current flows. However, the configuration change applies to new calls and new flows on same calls.

The following is a sample configuration that describes how to create, configure, or delete ACS rule definitions.

```
configure
   active-charging service service_name
      ruledef ruledef_name
         dns { any-match value | query-type query_type | query-name query_name
 }
         ip any-match [ = | != ] [ TRUE | FALSE ]
         ip dst-address { operator { { ipv4_address | ipv6_address } | {
ipv4_address/mask | ipv6_address/mask} | address-group ipv6_address } | { !range |
 range } host-pool host_pool_name }
         ip server-ip-address { operator { { ipv4_address | ipv6_address } | {
ipv4_address/mask | ipv6_address/mask} | address-group ipv6_address } | { !range |
 range } host-pool host_pool_name }
         multi-line-or all-lines
         rule-application { charging | post-processing | routing }
         { tcp | udp } { either-port port_number }
         end
```

NOTES:

- **ruledef** *ruledef_name*: Specifies the ruledef to add, configure, or delete. *ruledef_name* must be the name of an ACS ruledef, and must be an alphanumeric string of 1 to 63 characters, and can contain punctuation characters. Each ruledef must have a unique name. Host pool, port map, IMSI pool, and firewall, routing, and charging ruledefs must have unique names.

- If the named ruledef does not exist, it is created, and the CLI mode changes to the ACS Ruledef Configuration Mode wherein the ruledef can be configured.

- If the named ruledef already exists, the CLI mode changes to the ACS Ruledef Configuration Mode for that ruledef. The ACS Ruledef Configuration Mode is used to create and manage rule expressions in individual rule definitions (ruledefs).

- **ip any-match [= | !=] [TRUE | FALSE]:** This command defines the rule expressions to match IPv4/IPv6 packets. The *operator* and *condition* in the command specifies the following:

    - *operator*

        - !=: Does not equal

        - < =: Equals

    - *condition*

        - FALSE

        - TRUE

- **ip dst-address { ** *operator* **{ { ** *ipv4_address* | *ipv6_address* **} | { ** *ipv4_address/mask* |*ipv6_address/mask* **} | address-group** *ipv6_address* **} | { !range | range } host-pool** *host_pool_name* **}**: This command allows defining rule expressions to match IP destination address field within IP headers.

  - *ipv4_address* | *ipv6_address*: Specifies the IP address of the destination node for outgoing traffic. *ipv4_address* | *ipv6_address* must be an IP address in IPv4 dotted-decimal or IPv6 colon-separated-hexadecimal notation.

  - *ipv4_address/mask* | *ipv6_address/mask*: Specifies the IP address of the destination node for outgoing traffic. *ipv4_address/mask* | *ipv6_address/mask* must be an IP address in IPv4 dotted-decimal or IPv6 colon-separated-hexadecimal notation with subnet mask bit. The mask bit is a numeric value which corresponds to the number of bits in the subnet mask.

  - *address-group ipv6_address*: Specifies a group of IPv6 addresses configured with wildcard input and/or specialized range input. Multiple wildcard characters can be accepted as input and only one 2 byte range input will be accepted. Both wildcard character input and 2-byte range input can be configured together within a given IPv6 address.

  - **host-pool** *host_pool_name*: Specifies the name of the host pool. *host_pool_name* must be an alphanumeric string of 1 to 63 characters.

  - The *operator* in the command specifies the following:

    - !=: Does not equal

    - <: Lesser than or equals

    - =: Equals

    - >=: Greater than or equals

- **multi-line-or all-lines**: This command allows a single ruledef to specify multiple URL expressions. When a ruledef is evaluated, if the multi-line-or all-lines command is configured, the logical OR operator is applied to all the rule expressions in the ruledef to decide if the ruledef matches or not. If the multi-line-or all-lines command is not configured, the logical AND operator is applied to all the rule expressions.

- **rule-application { charging | post-processing | routing }**: This command specifies the rule application for a rule definition**.**

  - **charging**: Specifies that the current ruledef is for charging purposes.

  - **post-processing**: Specifies that the current ruledef is for post-processing purposes. This enables processing of packets even if the rule matching for them has been disabled.

  - **routing**: Specifies that the current ruledef is for routing purposes. Up to 256 rule definitions can be defined for routing in an Active Charging Service. Default: Disabled.

- **dns { any-match** *value* **| query-type** *query_type* **| query-name** *query_name* **}**: This command allows you to define rule expressions to match all DNS packets, or packets based on the query type or query name.

  ip server-ip-address *ip_address_value*: This command allows you to define rule expressions to match the IP address of the destination end of the connection.

- **{ tcp | udp } { either-port** *port_number* **}**: This command allows you to define rule expressions to match either a destination or source port number in UDP/TCP headers.

Charging Action Configuration for L7 Protocols for DPI

This section describes how to configure charging action. The charging action represents actions to be taken when a configured rule is matched. Actions could range from generating an accounting record (for example, an EDR) to dropping the IP packet, and so on. The charging action will also determine the metering principle—whether to count retransmitted packets and which protocol field to use for billing (L3, L4, L7, and so on).

The charging action configuration is used to define the QoS and charging related parameters associated with ruledefs.

```
configure
   active-charging service service_name
     charging-action charging_action
       allocation-retention-priority priority [ pci pci_value | pvi pvi_value

       billing-action egcdr
       cca charging credit [ rating-group coupon_id ] [ preemptively-request
 ]
       content-id content_id
       flow action { discard [ downlink | uplink ] | redirect-url
redirect_url | terminate-flow }
       flow limit-for-bandwidth { { direction { downlink | uplink }
peak-data-rate bps peak-burst-size bytes violate-action { discard |
lower-ip-precedence } [ committed-data-rate bps committed-burst-size bytes
 [ exceed-action { discard | lower-ip-precedence } ] ] } | { id id } }
       nexthop-forwarding-address ipv4_address/ipv6_address
       qos-class-identifier qos_class_identifier
       service-identifier service_id
       tft packet-filter packet_filter_name
       tft-notify-ue
       tos { af11 | af12 | af13 | af21 | af22 | af23 | af31 | af32 |
af33 | af41 | af42 | af43 | be | ef | lower-bits tos_value } [ downlink |
uplink ]
```

NOTES:

- **charging-action** *charging_action_name*: Specifies the name of a charging action. *charging_action_name* must be an alphanumeric string of 1 to 63 characters and can contain punctuation characters. Each charging action must have a unique name.

- If the named charging action does not exist, it is created, and the CLI mode changes to the ACS Charging Action Configuration Mode wherein the charging action can be configured.

- If the named charging action already exists, the CLI mode changes to the ACS Charging Action Configuration Mode for that charging action.

- **allocation-retention-priority** *priority* [**pci** *pci_value* | **pvi** *pvi_value*]: Configures the Allocation Retention Priority (ARP). *priority* must be an integer value in the range of 1-15.

   - **pci** *pci_value* : **Specifies the Preemption Capability Indication (PCI) value. The options are:**

      - MAY_PREEMPT - Flow can be preempted. This is the default value.

> • NOT_PREEMPT - Flow cannot be preempted

> • **pvi** *pvi_value*: Specifies the Preemption Vulnerability Indication (PVI) value. The options are:

>> • NOT_PREEMPTABLE - Flow cannot be preempted. This is the default value.

>> • PREEMPTABLE - Flow can be preempted

- **billing-action**: Configures the billing action for packets that match specific rule definitions.

- **cca charging credit**: Enables or disables credit control charging credit behaviour.

- **content-id**: Configures the rating group.

- **flow action**: Specifies the action to take on packets that match rule definitions.

  - **discard [downlink | uplink]**: Specifies to discard downlink or uplink packets.

  - **redirect-url** *redirect_url*: Specifies the URL to be redirected.

  - **terminate-flow**: Specifies to terminate the flow.

- **flow limit-for-bandwidth**: Configures the QoS parameters such as MBR, GBR, and so on.

  - peakdatarate(MBR): Default is 3000 bps

  - peakburstsize: Default is 3000 bytes

  - committedDataRate(GBR): Default is 144000 bps

  - committedBurstSize: Default is 3000 bytes

- **nexthop-forwarding-address** *ipv4_address/ipv6_address* ,: Configures the nexthop forwarding address.

- **qos-class-identifier** *qos_class_identifier* : Configures the QCI for a charging action. *qos_class_identifier* must be an integer value in the range of 1-9 or from 128-254 (Operator specific).

- **service_identifier** *service_id*: Configures the service identifier to use in generated billing records.*service_id* must be an integer value in the range of 1-2147483647.

- **tft packet-filter** *packet_filter_name*: Specifies the packet filter to add or remove from the current charging action. *packet_filter_name* must be the name of a packet filter, and must be an alphanumeric string of 1 to 63 characters.

- **tft-notify-ue**: Control the TFT updates towards the UE based on certain trigger conditions.

- **tos**: Configures the Type of Service (ToS) octets.

CHAPTER 12

Device ID in EDNS0 Records

Feature Summary and Revision History

Summary Data

Table 26: Summary Data

Applicable Product (s) or Functional Area	5G-UPF
Applicable Platforms	VPC-SI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

Table 27: Revision History

Revision Details	Release
TCP support is added.	2021.02.1
First introduced.	2021.01.2

Feature Description

The Device ID in EDNS0 offers each enterprise with a customized domain blocking through Umbrella. To enable this functionality:

- The UPF must reformat a subscriber DNS request into an EDNS0 request, and

- The UPF must include an Umbrella "Device ID" in the EDNS0 packet so that the Umbrella DNS resolver can use the Device ID to apply the domain filter associated/configured with the Device ID in the EDNS0 packet.

Presently, the Session Management Function (SMF) receives the domain filtering policy ID from PCRF/PCF. The SMF passes the domain filtering policy ID to the User Plane Function (UPF) in the Subscriber Parameters. The UPF uses the domain filtering policy ID to apply domain filtering functionality to the subscriber.

How it Works

New CLIs are introduced to configure and trigger the EDNS0 functionality.

The EDNS0 packet receives the 64-bit device ID as OPT RR data. The first 32 bits of all device IDs is a fixed value configured in the UPF. The last 32 bits of a subscriber device ID is the content filter ID value received from the PCRF/PCF. The UPF concatenates the two 32-bit values to build a subscriber full 64-bit Device ID for populating in the subscriber EDNS0 queries. New CLI helps to configure the first 32 bit of static device-id value. If you don't configure the 32-bit static prefix CLI, the outgoing packet shows the device-id = 32-bit CF PolicyID.

64 bit Umbrella Device ID

32 bits

32 bits

New static configurable parameter in the P-GW/UPF

Content filter policy ID passed down from the PCRF over the Gx interface for the subscriber

458078

The Device ID number in the EDNS0 record allows the Umbrella DNS system to apply a custom set of domain filters for the EDNS0 queries.

Process Flow

The following process flow describes about the Content Filtering enhancement to insert Device ID in EDNS0 records:

EDNS0 Packet Format

The enterprise policy ID (CF_POLICY_ID) from PCRF helps to create the Device ID. The SMF sends the device ID to the UPF. Adding the Device ID to the DNS packet helps in creating the EDNS0 packet. The format of EDNS0 packets is specified by RFC2671. The following are few specifics:

- Following is the structure for the fixed part of an OPT RR:

```
Field Name    Field Type      Description
----------------------------------------------------
NAME          domain name     empty (root domain)
TYPE          u_int16_t       OPT
CLASS         u_int16_t       sender's UDP payload size
TTL           u_int32_t       extended RCODE and flags
RDLEN         u_int16_t       describes RDATA
RDATA         octet stream    {attribute, value} pairs
```

- Following is the variable part of an OPT RR encoded in its RDATA:

```
     +0 (MSB)                          +1 (LSB)
     +---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
  0: |                          OPTION-CODE                         |
     +---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
  2: |                          OPTION-LENGTH                       |
```

```
            +---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
         4: |                                                               |
            /                        OPTION-DATA                            /
            /                                                               /
            +---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

- OPTION-CODE: Assigned by IANA

- OPTION-LENGTH: Size (in octets) of OPTION-DATA

- OPTION-DATA- Varies per OPTION-CODE

Example: If received policy-id from PCF/PCRF is "1234" and static prefix configured on UPF is "5678".

64-bits Device-ID will be "0000162e000004d2".

- 0000162e -- 5678 (Decimal)

- 000004d2 -- 1234 (Decimal)

RDATA 69 42 00 0f 4f 70 65 6e 44 4e 53 00 00 16 2e 00 00 04 d2

- 6942 -- option-code

- 000f -- option-length

- 4f70656e444e53 -- OpenDNS (String)

- 0000162e -- 5678 (MSB)

- 000004d2 -- 1234 (LSB)

EDNS0 with IP Readdressing

The new CLI is configured within trigger action to readdress the DNS traffic to the Umbrella DNS. This CLI uses the existing readdress server list configuration from the ACS service. Readdressing of packets based on the destination IP address of the packets enables redirecting gateway traffic to configured server/port in the readdressed server list.

Behavior and Restrictions

Following are the behavior and restrictions applicable for this feature:

- Trigger Condition is evaluated at flow creation time. Any change in trigger condition in between the flow doesn't affect the existing flow but affects the new flows.

- Any change to trigger action is applicable on the same flow.

- Neither CF nor EDNS is enforced when the CF Policy ID range is defined but Service-schema is not defined, or the Trigger condition pertaining to EDNS is not configured.

- If no CF Policy ID is received from Gx, range check is not performed, and content filtering works as defined in rule base.

- Cases where the 'security-profile' CLI is not associated with the 'EDNS format' CLI in Trigger Action, the device-id in the outgoing EDNS packet is sent with only 32-bit CF Policy ID.

- DNS queries with type other than A, AAAA, CNAME, NS, PTR, SRV, TXT, NULL are not to be EDNS converted.

- CF Policy ID change over Gx in between inflow are not applicable for the current flows. The current flows continue to insert the CF Policy ID present at the time of flow creation.

Limitation

Following are the limitations for this feature:

- The feature doesn't support the EDNS response packet reformat.

- The UPF must be able to include the IMSI MSISDN tag value in the EDNS0 queries. This feature doesn't support the encrypted IMSI in EDNS0 packet. This feature also doesn't support the following configuration on the EDNS fields currently.

```
configure
   active-charging-service  service_name
      edns
         fields  fields_name
            tag default device-id
            tag 101 imsi encrypt
            tag 102 pgw-address
            end
```

Configuring EDNS Format and Trigger Action

Use the following configuration to enable DNS filtering:

```
configure
   active-charging-service service_name
      content-filtering range start_min_val to end_max_val
```

If the range parameter is set from 10 through 1000, any subscriber profile with a content filter policy ID from 10 through 1000 uses the standard content filtering functionality. Any subscriber profile with a content filter policy ID higher than 1000 or lower than 10 triggers the new EDNS0 functionality.

Use the following configuration to disable DNS filtering:

```
configure
   active-charging-service service_name
      no content-filtering range
```

When DNS filtering is disabled, the standard content filtering policies resume as configured or as received from the PCF.

Use the following configuration to configure the EDNS packet action and format under the active-charging service:

```
configure
   active-charging-service service_name
      trigger-condition trigger_condition_name
         external-content-filtering
```

```
            app-proto = dns
          end
```

NOTES:

- **external-content-filtering**: Enables EDNS0 feature. When this flag is true along with the range criteria, EDNS0 feature is enabled. By default, this flag is disabled.

- **app-proto** = **dns**: Avoids the IP readdressing of the non-DNS traffic. If this CLI is enabled with multiline-or cli, then all DNS traffic is EDNS encoded.

The following configuration leads the trigger action to define the EDNS format to be inserted in the EDNS packet:

```
configure
  active-charging-service service_name
    trigger-action trigger_action_name
      edns-format format_name
        security-profile profile_name
          flow action readdress server-list server_list_name [ hierarchy ]
 [ round-robin ] [ discard-on-failure ]
            end
```

NOTES:

- **trigger-action** *trigger_action_name*: Enables you to configure the flow action CLIs in the trigger action.

- **edns-format** *format_name*: Use the EDNS format when EDNS is applied.

- **security-profile** *profile_name*: Defines the security profile configuration in the EDNS to add mapping with the Device-ID.

  > **Note** Device ID in EDNS0 Records feature supports multiple security profiles.

- **flow action readdress server-list** *server_list_name* **[hierarchy] [round-robin] [discard-on-failure]**: Associates the EDNS with IP readdressing. Use IP readdressing to readdress the packets to the configured server IPs. This CLI in trigger action supports only server list configuration. It doesn't support single-server IP or port configuration like charging action.

Use the following configuration to insert the CF policy ID in the EDNS:

```
configure
  active-charging-service service_name
    edns
      fields fields_name
        tag { val { imsi | msisdn | cf-policy-id }}
        end
```

To configure the 32 MS bit, static value is provided at the EDNS level with the security profile.

```
security-profile security_profile cf-policy-id-static-prefix value
```

Use the following configuration to insert a new tag specifying the payload length:

```
tag default payload-length [ tcp | udp ] value
```

Value range: 576–4096

Sample Configuration

Following is the sample configuration for configuring the EDNS packets:

```
configure
active-charging service ACS
content-filtering range 10 to 100
ruledef dns-port
udp either-port = 53
tcp either-port = 53
multi-line-or all-lines
rule-application routing
#exit
 readdress-server-list re_adr_list_ta
   server 100.100.100.14
   server 2001::14
   server 100.100.100.15
   server 2001::15
#exit
rulebase starent
route priority 20 ruledef dns-port analyzer dns
#exit
edns
security-profile sec_profile cf-policy-id-static-prefix 123456
fields test_fields
tag 26946 cf-policy-id
#exit
format test_format
fields test_fields encode
#exit
#exit
trigger-action TA1
edns format test_format security-profile sec_profile
flow action readdress server-list re_adr_list_ta hierarchy
#exit
trigger-condition TC1
external-content-filtering
app-proto = dns
#exit
service-scheme SS1
trigger flow-create
priority 1 trigger-condition TC1 trigger-action TA1
#exit
subs-class SC1
rulebase = starent
multi-line-or all-lines
#exit
subscriber-base SB1
priority 1 subs-class SC1 bind service-scheme SS1
exit
end
```

Monitoring and Troubleshooting

Following are the show commands and outputs in support of enhance content filtering support to Insert device ID in EDNS0 records.

Show Commands and Outputs

Following are the show commands and outputs that are modified in support of the enhance content filtering support to Insert device ID in EDNS0 records.

- **show user-plane-service inline-services info**

```
CF Range: Enabled <<<<
   Start Value:  1 <<<<
   End Value:    1000 <<<
```

- **show user-plane-service statistics analyzer name dns**: output is modified to include the "EDNS Response Received" in both "EDNS Over UDP" and "EDNS Over TCP" sections.

- **show subscribers user-plane-only full callid**: output is modified to include the following parameters in the EDNS statistics per subscriber.

  - DNS-to-EDNS Uplink Pkts

  - DNS-to-EDNS Uplink Bytes

  - EDNS Response Received

- **show user-plane-service edns all**

```
Fields:
    Fields Name: fields_1
    tag 26946 cf-policy-id

    Fields Name: fields_2
    tag 2001 imsi
    tag 2002 msisdn
    tag 26946 cf-policy-id

    Format:
    Format Name: format_1
    fields fields_1 encode

    Format Name: format_2
    fields fields_2 encode

    Security-profile Name: high
    CF Prefix Policy ID: 1234
```

Use the following show commands to view the Trigger Action statistics:

- **show user-plane-service statistics trigger-action all**

- **show user-plane-service statistics trigger-action name** *trigger_action_name*

- **show user-plane-service trigger-condition all**

```
Trigger-Condition: TC1
      External-content-filtering : Enabled
      App-proto : dns
      Multi-line-OR All lines : Disabled
```

- **show user-plane-service trigger-action all**

```
Trigger-Action: TA1
      HTTP Response Based TRM      : none
      HTTP Response Based Charging : none
      Throttle Suppress            : Disabled
```

```
        Flow Recovery            : Disabled
        Traffic Optimization     : Disabled
        Step Up GBR              : Disabled
        Step Down GBR            : Disabled
        TCP Acceleration         : Disabled
        TCP Acceleration Threshold : Disabled
        Service-Chain            : none
        UP-Service-Chain         : none
        EDNS-Encode              : Enabled
        Flow-IP-Readdressing     : Enabled
```

Bulk Statistics

The following bulk statistics are available in support of the Device ID in EDNS0 Records feature:

SCHEMA: ECS	
Statistics	**Description**
ecs-dns-udp-edns-encode-succeed	The count of DNS to EDNS converted packets over UDP
ecs-dns-udp-edns-encode-failed	The count of failed DNS to EDNS conversions over UDP
ecs-dns-udp-edns-encode-response	The count of responses received for EDNS query over UDP
ecs-dns-tcp-edns-encode-succeed	The count of DNS to EDNS converted packets over TCP
ecs-dns-tcp-edns-encode-failed	The count of failed DNS to EDNS conversions over TCP
ecs-dns-tcp-edns-encode-response	The count of responses received for EDNS query over TCP

CHAPTER 13

Downlink Data Notification

Feature Summary and Revision History

Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

Revision Details	Release
First introduced.	2021.02.0

Feature Description

The Downlink Data Notification (DDN) messages with support for DDN Delay and DDN Throttling, and buffering in SAEGW when UE is in Idle State, is supported in 5G-UPF.

How It Works

This section provides an overview of how this feature works.

- Buffering is supported at SAEGW-U.

- Support of buffering starts when UE moves to IDLE state due to Release Access Bearer.

- ACTIVE to IDLE transition:

  - When the UE moves to ECM-IDLE state, since the SAEGW supports buffering capability and decides to activate buffering in SAEGW-U for the session, the SAEGW-C informs the SAEGW-U through an Sx session modification.

  - After the buffering starts, when the first downlink packet arrives on any bearer, the SAEGW-U informs the SAEGW-C. The SAEGW-U sends an Sx reporting message to the SAEGW-C, unless specified otherwise, and identifies the S5/S8 bearer on which the downlink packet is received.

  - On receiving the reporting message, the SAEGW-C decides whether to send a DDN message to the MME, as defined in 3GPP TS 23.401 [2]. The DDN notification is sent with the Sx-Usage-Report.

- IDLE to ACTIVE transition:

  - At the UE transition to ECM-CONNECTED state, the SAEGW-C updates the SAEGW-U through Sxa interface with the F-TEIDu of the eNodeB/RNC/SGSN. The buffered data packets, if any, are then forwarded to the eNodeB/RNC/SGSN by the SAEGW-U.

- If the Apply Action is BUFFER, and SGW-U recovers, the SGW-U initiates Sx Report (with DLDR Report Type) on arrival of the downlink data packet.

- In SGW-U, a timer is implemented that starts after each Sx Report (with DLDR report Type) is sent. If the Apply Action is not changed, then on timer expiry, Sx Report (with DLDR Report Type) gets initiated again.

- ARP of the bearer is included in the DDN message.

- In a multi-PDN session, if the DDN is initiated for one PDN and then data is received on another PDN, wherein the bearer has higher priority, then the DDN is initiated again with the higher priority ARP value.

Downlink Data Notification – Delay (DDN-D) Support

Under certain conditions, when UE triggers a service request, uplink and downlink data is triggered and is received at the SGW-C even before the Modify Bearer Request (MBR) is received causing unnecessary Downlink Packet Notification messages sent that increases the load in MME.

In such cases, the MME monitors the rate at which these events occur. If the rate becomes significant (as configured by the operator) and the MME's load exceeds an operator configured value, the MME indicates "Delay Downlink Packet Notification Request" with parameter D to the Serving Gateway, where D is the requested delay given as an integer with multiples of 50 milliseconds, or zero. The S-GW then uses this delay in between receiving downlink data and sending the Downlink Data Notification message.

The Downlink Data Notifications are supported for both Collapsed and Pure-S calls.

Due to the distributed nature of the system, sessions from a particular MME are offloaded on different session managers. Therefore, all session managers are notified when a session is offloaded. Also, the functionality is designed to not allow all session managers to message the DEMUX manager.

- In DDN Delay feature, DDN delay timer support is at Session Management Function.

- When first data packet arrives, Sx/N4 Report message is initiated but DDN message is initiated from Session Management Function after the expiry of Delay timer.

- DDN Delay feature is a peer level feature and so, it is applied for all the session on that peer from where the DDN Delay value is received.

- In case a previous delay value was received from a peer and it is absent in the current message, the delay value will be considered as 0.

Session Recovery and ICSR is supported for DDNs.

5G SMF Calls

Downlink Data Notification - 5G UE

When UE turns to Idle state in the 5G call mode, SMF sends Sx_Modify_Request with FAR Apply Action set to BUFFER value. For every QFI (default/dedicated), the FAR Apply Action is set to BUFFER value.

Once the Downlink packet is received from the server, UPF sends the Sx_Report_Request with Downlink Data Notification with Rule Id (PDR Id)/QFI as per packet rulematch. UPF continues buffering packets until the packet limit reaches 5 for every FAR.

When UE is active, Sx_Modify_Request reaches UPF with FAR Apply Action set to FWD (Forward) and TEID (Tunnel Endpoint Identifier). UPF debuffers the packets and sends them to UE as FIFO. For each packet, rule match will take place after the debuffering process.

DDN Throttling Support

Too many DDN requests towards MME from SGW-C could lead to processing overload at MME. To reduce this load, MME dynamically requests SGW-C to reduce a certain percentage of DDN messages sent towards it for a given period time.

For DDN throttling, S-GW is required to drop a given percentage of DDNs over a given period of time. S-GW implements this functionality using a probabilistic algorithm at each session manager.

Whereas, the conventional implementation of DDN throttling requires each session manager to share its list of pending DDNs for low priority bearers with a central entity that would then calculate the net load of pending DDNs and then decide how many DDNs each session manager would have to drop. This implementation would require buffering of DDN messages at session manager. Also, due to distributed processing nature of

software subsystem in chassis, it would require considerable amount of messaging between the session managers and the central entity (demuxmgr in case of Boxer) at regular intervals.

Implementing a probabilistic algorithm removes the need for buffering at session manager and also messaging with demuxmgr. Accuracy of probabilistic algorithm increase with increasing low ARP priority paging load at session manager. Even with lower paging load, accuracy would be fairly close to the throttling factor provided.

For non-release 10 compliant MME, SGW_C provides option to enable throttling through the CLI.

Threshold ARP values for low priority bearer must be configured through S-GW Service Configuration. For example, if configured ARP value is 9, any bearer with ARP > 9 is considered low priority bearer. DDN throttling is enabled through this configuration. If DDN throttling is enabled through SGW service configuration, each DDN message towards MME would contain the ARP IE.

No User Connect Timer Support

- Timer is introduced when a Modify Bearer Request is not received after positive Downlink Data Notification acknowledgment.

- It is initiated at SGW-C when DDN acknowledgment is received.

- On arrival of Modify Bearer Request, SGW-C stops this timer.

- On timer expiry SGW-C informs SGW-U to drop buffered packets.

DDN Call Flows

DDN Success Scenario

1. MME sends Release Access Bearer request to SGW-C to release downlink remote TEIDs of all the bearers for that UE.

2. On arrival of Release Access Bearer request, SGW-C informs the same to SGW-U by updating FAR with Apply Action as BUFFER in Sx Modification Request for all the PDNs.

3. SGW-U send Sx Modification response after applying Buffering in SGW-U for corresponding PDN.

4. SGW-C sends Release Access Bearer response to MME.

5. First Downlink data arriving in SGW-U triggers Sx Report Request (with Report Type as Downlink Data Report) towards SGW-C.

6. On arrival of Sx Report Request message, the SGW-C initiates Downlink Data Notification request message towards MME.

7. SGW-C sends Sx Report Response message towards SGW-U.

8. If MME is able to send a paging request towards UE, it sets the cause as "Request Accepted" in Downlink Data Notification Acknowledgment Message and sends it to SGW-C.

9. On successful paging, MME sends a Modify Bearer request to the S-GW with eNodeB TEIDs that sets up the S1-U connection at the SGW.

10. SGW-C sends Sx Modification request with updated FAR for new TEID information to SGW-U. SGW-U can now forward all the buffered data to UE through eNodeB.

11. SGW-U sends Sx Modification response to SGW-C.

DDN Failure Scenario

1. MME sends Release Access Bearer request to SGW-C to release downlink remote TEIDs of all the bearers for that UE.

2. On arrival of Release Access Bearer request, SGW-C informs the same to SGW-U by updating FAR with Apply Action as BUFFER in Sx Modification Request for all the PDNs.

3. SGW-U send Sx Modification response after applying Buffering in SGW-U for corresponding PDN.

4. SGW-C sends Release Access Bearer response to MME.

5. First Downlink data arriving in SGW-U triggers Sx Report Request (with Report Type as Downlink Data Report) towards SGW-C.

6. On arrival of Sx Report Request message, the SGW-C initiates Downlink Data Notification request message towards MME.

7. SGW-C sends Sx Report Response message towards SGW-U.

8. If MME is not able to page UE then it can reject Downlink Data Notification Request with relevant cause.

OR

If MME accepts Downlink Data Notification Request. But later sends Downlink Data Notification Failure indication in order to indicate SGW-C that the UE did not respond to paging.

9. SGW-C received DDN failure and hence to stop sending next DDN immediately , SGW-C starts DDN Failure Timer. SGW-C sends Sx Modification Request with DROBU flag to discard buffered packets and Apply Action as DROP to drop subsequent packets.

10. SGW-U sends Sx Modification Response to SGW-C.

11. On DDN Failure Timer Expiry SGW-C initiates Sx Modification with Apply Action as BUFFER in order to start buffering again.

Further steps are continued from Step 3 in the call flow.

No User Connect Timer Support

1. MME sends Release Access Bearer request to SGW-C to release downlink remote TEIDs of all the bearers for that UE.

2. On arrival of Release Access Bearer request, SGW-C informs the same to SGW-U by updating FAR with Apply Action as BUFFER in Sx Modification Request for all the PDNs.

3. SGW-U send Sx Modification response after applying Buffering in SGW-U for corresponding PDN.

4. SGW-C sends Release Access Bearer response to MME.

5. First Downlink data arriving in SGW-U triggers Sx Report Request (with Report Type as Downlink Data Report) towards SGW-C.

6. On arrival of Sx Report Request message, the SGW-C initiates Downlink Data Notification request message towards MME.

7. SGW-C sends Sx Report Response message towards SGW-U.

8. Downlink Data Notification Acknowledgment is received from MME.SGW-C starts no-user-connect timer.

9. If the Modify Bearer request with eNodeB TEID information is not received and no-user-connect timer expires, SGW-C sends Downlink Data Notification again.

10. Downlink Data Notification Acknowledgment is received from MME. SGW-C initiates the no-user-connect timer again.

11. SGW-C initiates Sx Session Modification request towards SGW-U with DROBU flag set in the message. On receiving this flag SGW-U drops the buffered data. New data will be buffered, and the subsequent first packet initiates a Sx Report message for initiating Downlink Data Notification message.

12. SGW-U sends Sx Modification Response.

DDN Delay Timer

1. MME sends Release Access Bearer request to SGW-C to release downlink remote TEIDs of all the bearers for that UE.

2. On arrival of Release Access Bearer request, SGW-C informs the same to SGW-U by updating FAR with Apply Action as BUFFER in Sx Modification Request for all the PDNs.

3. SGW-U send Sx Modification response after applying Buffering in SGW-U for corresponding PDN.

4. SGW-C sends Release Access Bearer response to MME.

5. First Downlink data arriving in SGW-U triggers Sx Report Request (with Report Type as Downlink Data Report) towards SGW-C.

6. On arrival of Sx Report Request message, the SGW-C initiates Downlink Data Notification request message towards MME.

7. SGW-C sends Sx Report Response message towards SGW-U.

8. Downlink Data Notification Acknowledgment is received from MME with DDN Delay Timer value. This timer value will be saved for this peer , and now onwards every Downlink Data notification that we initiate should be after this delay for that peer.

9. On success paging, MME sends a Modify bearer request to the SGW with eNodeB TEIDs that sets up the S1-U connection at the SGW.

10. SGW-C sends Sx Modification Request with updated FAR for new TEID information to SGW-U. SGW-U can now forward all the buffered data to UE via eNodeB.

11. SGW-C sends Modify Bearer Response to MME.

12. SGW-U sends Sx Modification Response to SGW-C.

13. MME sends Release Access Bearer Request to SGW-C to release downlink remote TEIDs of all the bearers for that UE.

14. On arrival of Release Access Bearer Request, SGW-C inform the same to SGW-U via updating FAR with Apply Action as BUFFER in Sx Modification Request for all the PDNs.

15. SGW-U send Sx Modification Response after applying Buffering in SGW-U for corresponding PDN.

16. SGW-C sends Release Access Bearer Response to MME.

17. First Downlink data arriving in SGW-U triggers Sx Report Request (with Report Type as Downlink Data Report) towards SGW-C.

18. On arrival of Sx Report Request message, SGW-C starts DDN Delay Timer. On DDN Delay timer expiry SGW-C Initiates Downlink Data Notification message towards MME.

Sx Interface

Sx Session Level Reporting Procedure

Detection of first Downlink Data for Idle-Mode UE (by SAEGW-U):

When SAEGW-U receives the downlink packet but no S1-bearer for transmission and the buffering is performed by SAEGW-U, it reports the detection of first downlink data to SAEGW-C, for the purpose of paging the UE.

PFCP Session Report Request

The PFCP Session Report Request is sent over the Sxab interface by the User Plane function to report information related to a PFCP session to the Control Plane function.

Information elements	P	Condition / Comment	Appl.				IE Type
			Sxa	Sxb	Sxc	N4	
Report Type	M	This IE shall indicate the type of the report.	X	X	X	X	Report Type
Downlink Data Report	C	This IE shall be present if the Report Type indicates a Downlink Data Report.	X	-	-	X	Downlink Data Report

Downlink Data Report IE within PFCP Session Report Request

The Downlink Data Report grouped IE is encoded as shown in the following table.

Octet 1 and 2		Downlink Data Report IE Type = 83 (decimal)					
Octets 3 and 4		Length = n					
Information elements	P	Condition / Comment	Appl.				IE Type
			Sxa	Sxb	Sxc	N4	

| PDR ID | M | This IE shall identify the PDR for which downlink data packets have been received at the UP function.

More than one IE with this type may be included to represent multiple PDRs having received downlink data packets. | X | - | - | X | PDR ID |
|--------|---|---|---|---|---|---|--------|

Notification to User Plane Function for DDN Failure

The Control Plane function notifies User Plane function for any failure so that buffered packets can be dropped and DDN related flags can be reset through DROBU flag in PFCP Sx Modification message.

PFCPSMReq-Flags	C	DROBU (Drop Buffered Packets): The CP function shall set this flag if the UP function is requested to drop the packets currently buffered for this PFCP session (see NOTE 1).

Limitations

Following are the known limitations of this feature:

- SAEGW Buffering is done for five data packets per PDN session.

- DDN profile configuration is not supported.

- Support for buffered data (data packet stream) that get deleted due to Flow Idle Timeout or other cases, is not present.

DDN Throttling for non-Release 10 Compliant MME

Use the following configuration to configure DDN throttling for a non-release 10 MME:

```
configure
  context context_name
    sgw-service service_name
      ddn throttle arp-watermark arp_value [ rate-limit limit time-factor
seconds throttle-factor percent increment-factor percent [ poll-interval seconds
 ] throttle-time-sec seconds [ throttle-time-min minutes ] [
throttle-time-hour hour ] stab-time-sec seconds [ stab-time-min minutes ] [
stab-time-hour hour ]
      no ddn throttle
      end
```

NOTES:

- **rate-limit**: DDN permitted per second.

- **time-factor**: Time period in seconds over which SGW makes throttling decision (valid range 1-300 seconds.

- **arp-value**: Valid ARP value between 1 and 15. All the packets which have arp greater than the configured values will be throttled as per the throttling factor.

- **throttling-factor**: Percentage of DDN to be dropped upon detecting DDN surge (valid range between 1-100).

- **throttling-time-sec**: Time period in seconds over which DDN are throttled at SGW (valid range between 0-59 seconds).

- **throttling-time-min**: Time period in minutes over which DDN are throttled at SGW (valid range between 0-59 minutes).

- **throttling-time-hour**: Time period in hours over which DDN are throttled at SGW (valid range between 0-310 hours).

- **increment-factor**: Percentage value by which throttling factor is incremented dynamically, if existing throttling factor is insufficient to curb the DDN surge.

- **poll-interval**: Time in seconds (optional argument, default value = 1 second, poll interval < time-factor)

- **stab-time-sec/min/hours**: Stabilization time factor, time period over which if DDN rate returns to normal, then throttling need not be applied over entire throttling time period.

DDN throttling for non-Release-10 compliant MME makes use of existing Release-10 throttling implementation at SGW. By providing a configuration mechanism for SGW service, operator can still apply ddn throttling without needing any feedback from MME. Some salient points of this feature are described below:

1. The CLI configuration is applied per MME/S4-SGSN. Throttling parameters are tracked independently per MME/S4-SGSN.

2. On configuring this feature through CLI, demuxmgr polls each sessmgr for number of DDNs sent. By default, polling is done every second. This time interval can be changed by configuring the poll-interval time. Greater the poll interval time, lesser the number of internal messages within the chassis. However, it would take longer to detect a DDN surge.

3. By configuring time-factor, operator can specify the time interval for S-GW to apply throttling, if needed. It allows for some surge of DDNs if the net DDN rate is within specified limit over time-factor time interval. For example, time-factor= 10 seconds, ddn rate = 1000, poll interval = 2 seconds. Demux would

poll each sessmgr every 2 seconds. Acceptable DDN rate limit is 1000*10 = 10000 DDNs every 10 seconds. Say after 2 seconds, 4000 DDNs were sent, in that case S-GW wouldn't apply throttling till rate limit of 10000 DDNs is crossed within time period of 10 seconds. This allows for intermittent bursts of DDNs.

4. DDN rate limit is configured through CLI. For example, if DDN rate limit is 1000 and poll interval = 1 second, time-factor = 5 seconds, then acceptable rate limit is 5000 DDNs over 5 seconds. If the number of DDNs sent by S-GW is greater than 5000 after 5 seconds, demuxmgr would ask all sessmgrs to initiate throttling.

5. Percentage of DDNs to be throttled is configured through throttling-factor.

6. Operator can specify increment-factor to increment throttling factor if existing throttling factor is insufficient to curb the DDN surge. For example, if throttling-factor = 10%, ddn-rate = 1000, increment-factor=10%. Once throttling is applied, S-GW drops ~10% DDNs. However, if DDN rate is still greater than 1000, S-GW would increase throttling-factor to 20%. If this is still not sufficient, it would be incremented to 30%. After incrementing throttling factor, if number of DDNs dropped are greater than expected, throttling-factor would then be decrement by increment-factor. E.g. in this case, after increasing throttling factor to 30%, if DDNs sent is less than 1000 per second (taking time-factor and poll-interval into consideration), throttling factor would be decremented to 20. The cap for decrementing throttling-factor would be the configured value (10% in this case).

7. Operator can configure the time duration for which throttling is applicable at S-GW. This could be a large value in order of days (for example: 10 days or 240 hours). The operator has an option to stop throttling if DDN rate is well under control by configuring stabilization time factor. In such a case, DDNs won't be needlessly dropped. For example, throttling-time =10 days, stab-time = 8 hours. After S-GW starts DDN throttling, in a time span of 8 hours, DDNs sent + DDNs dropped < ddn-rate * 8 hours, throttling would be stopped.

DDN Throttling for Release 10 Compliant MME

DDN throttling is enabled through Call Control Profile by providing the ARP value. For example, if the ARP value provided is 10, then all bearers with ARP value between 10-15 are treated as low priority bearers and are given throttling treatment. Throttling would not be enabled if ARP value is not provided through S-GW service configuration. Also, ARP IE in DDN message towards MME would not be included unless DDN throttling is configured using S-GW service. If MME is Release 10 compliant, the user need not configure the duration value as the DDN Acknowledgment would have the throttling IE. Otherwise, throttling can be enabled at S-GW by setting the duration value. If it's set to 0, S-GW would apply throttling recurringly. To enable throttling only for a given duration of time (in non Rel-10 compliant MME), user needs to set the value in hours and minutes. From the time of configuration, throttling would be applied at S-GW until the timer duration expires. For example, if user sets hours = 10, minutes = 30, S-GW would apply throttling for next 10 hours 30 minutes.

On re-configuration, all the parameters will be set with new values, but they will be applicable only from the next recalibration except from polling time and time factor.

Use the following configuration to configure DDN throttling for release 10 MME:

```
configure
   context context_name
      sgw-service service_name
```

```
        [ no ] ddn throttle arp-watermark arp_value
        end
```

NOTES:

- **arp-value**: Valid ARP value between 1 and 15. All the packets which have ARP greater than the configured values will be throttled as per the throttling factor.

Idle Timer for SAE-GW Sessions

An Idle Timer is supported to identify and remove idle sessions that occur in the SAE-GW.

A session becomes idle in some cases where the session is removed from other network nodes, but due to a technical mishap the session could still remain on the SAE-GW leading to resources being held by these idle sessions.

The Idle Timer, once configured, removes those sessions that remain idle for longer than the configured time limit effectively utilizing the system capacity.

☞

Important This feature is currently restricted to Pure-P and Collapsed Call.

Limitations

The Idle Timer feature does not support recovery of Idle Timer in case of redundancy events.

Configuring Idle Timer for SAE-GW Sessions

The Idle Timer is configurable at APN level.

Use the following commands to configure the idle timer for SAE-GW sessions:

```
configure
  context context_name
    apn apn_name
      timeout idle timeout_value
      no timeout idle
      default timeout idle
      end
```

- **no**: Disables the idle timer configuration.

- **default**: Configures the default value for subscriber's time out settings. The default idle timeout value is 0.

- **idle** *timeout_value*: Designates the maximum duration a session can remain idle, in seconds, before system automatically terminates the session. Must be followed by number of seconds between 0 and 4294967295. Zero indicates function is disabled.

S-GW Session Idle Timeout

This chapter describes the Idle Timeout Handling feature for S-GW sessions. On the ASR5500 platform, subscriber session is represented by call-line. The S-GW product call-line interfaces to its peers through MME/S4-SGSN on S11/S4 and P-GW on S5/S8. In some scenarios, peer sessions are deleted by respective peers, S-GW does not receive or miss deletion messages, and as a result S-GW session remains idle. Such idle or stale sessions are counted towards valid call-lines in system for effectively consuming resources and causing capacity reduction. In such cases, S-GW triggers to get the new subscriber session, which results in the removal of old session for same subscriber. The Idle Timeout Handling support enables the identification of such sessions and initiates deletion to release the resources.

The following points describes the idle timeout handling for S-GW sessions:

- The subscriber session is idle when there is no data traffic activity for the subscriber. The session manager keeps track of the call-line state, when no data traffic is recorded for call-line, such sessions are moved to idle state.

- Session which is idle for defined timeframe referred as idle timeout is considered for idle timeout handling. In idle timeout session, S-GW initiates the deletion of session towards its peers.

- Idle timeout is configured in seconds depending on the network requirements. The timeout range is 1-4294967295 seconds.

- The idle timeout configuration is applicable on S-GW service level for enabling the idle timeout handling for set of subscribers handled by that service.

Configuring Session Idle Timeout

The session idle timer for S-GW sessions is configurable from S-GW service.

To configure Session Idle Timeout for S-GW, use the following configuration:

```
configure
   context context_name
      sgw-service service_name
         [ no | default] timeout idle timeout_duration
         end
```

NOTES:

- **timeout idle** *timeout_duration*: Specifies the maximum duration a session can remain idle for, in seconds, before the system automatically terminates the session. *timeout_duration* must be an integer in the range of 1-4294967295. 0 disables the feature. By default, it is disabled for the S-GW service.

Show Commands Input and/or Outputs

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

show subscribers user-plane-only full all

The output of this command displays the following fields in support of this feature:

- buffered pkts

- buffered bytes

- buffer overflow drop pkts

- buffer overflow drop bytes

show subscribers user-plane-only full callid <call_id>

Use the following configuration to check the buffering per subscriber:

- DDN buffered packets

- DDN buffered bytes

- DDN buffer overflow drop packets

- DDN buffer overflow drop bytes

When the buffered packets are debuffered the state changes back to zero.

Each time the packet limit reaches 5, the additional packets get dropped as overflow drops.

DSCP Markings For Collapse Calls

Feature Summary and Revision History

Table 28: Summary Data

Applicable Product (s) or Functional Area	5G-UPF
Applicable Platforms	VPC-SI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Table 29: Revision History

Revision Details	Release
First introduced	2021.02.0

Feature Description

Currently, QCI-based DSCP markings are applicable for Pure-S and Pure-P calls. The DSCP markings are based on QCI-QOS-Mapping associated with respective S-GW service or P-GW service. For collapse calls QCI-QOS-Mapping associated with PGW-service is applicable. This feature helps to apply the DSCP markings for collapse calls based on associated S-GW and P-GW services for uplink and downlink traffic. For uplink traffic, DSCP markings associated with logical P-GW service are applicable. For downlink traffic, DSCP markings associated with logical S-GW service are applicable. The DSCP markings are present in IP header

of data traffic as a part of GTP-U header and Inner IP. There's option to enabled or disable this functionality by CLI configuration. When you enable the feature, then only the new functionality is applicable otherwise existing functionality also works. By default, this feature is disabled, so that there's no impact on customers who upgrades to this feature.

DSCP Markings for 5G Calls

The QCI/QOS mapping table in SMF drives the DSCP values being put on the packets. It's similar configuration as in PGW-CP. The DSCP values are sent as part of FAR during call establishment (Sx establishment request/Sx modification request). The DSCP values are being applied to inner packet and outer GTPU packet.

DSCP Markings for 4G Collapsed Datapath Calls

For collapsed Datapath calls:

- On Sxa leg, DSCP marking is sent in FAR from SGWC/Cn-SGWC.

- On N4 leg, DSCP marking is sent in FAR from SMF.

- For uplink packets, UPF applies the DSCP marking as per SMF configuration.

- For downlink packets, UPF applies the DSCP marking as per SGWC/Cn-SGWC configuration in GTPU header.

- For inner packet, DSCP marking is as per SMF configuration.

The SMF configuration/logic of sending the DSCP marking is similar as PGW-C. In case of the DSCP marking for the ECS charging action, priority is given to the charging action configuration for inner packet DSCP marking.

How It Works

Following are the steps that describe the DSCP markings for the collapse calls.

- In case of Collapse call:

  - For ACCESS side, QCI-QOS mapping table associated with SGW-service is used.

  - For CORE side, QCI-QOS mapping table associated with PGW-service is used.

- The preceding conditions apply once you enable the feature, otherwise QCI-QOS mapping table associated with PGW-service is applicable for both sides.

- APN associated QCI-QOS mapping table is preferred over the P-GW service QCI-QOS mapping table.

- APN-Profile associated QCI-QOS mapping table is preferred over SGW-Service QCI-QOS mapping table for ACCESS side DSCP markings.

- In case, only P-GW service has QCI-QOS mapping table configuration, then these DSCP markings is applicable on both ACCESS & CORE side for collapse call.

- In case only S-GW service has QCI-QOS mapping table configuration then these DSCP markings is applicable on ACCESS side for collapse call.

- There's a new configurable parameter inside the SAEGW service which indicates whether the feature is enabled or disable.

- For Pure-P to Collapse HO and conversely, transport layer markings are updated in FAR as a part of Sx Modify request.

- Layer2 markings are also modified based on QCI-QOS mapping table picked for ACCESS and CORE side.

- DSCP markings continues to apply on existing bearers post session recovery.

- DSCP markings continues for the bearers on standby chassis once it switches to active mode.

SessMgr SMF Changes

DSCP markings for Uplink/CORE and Downlink/ACCESS are present at bearer level inside sessmgr_sub_session_t → sessmgr_qci_tab_t.

User datagram DSCP markings are updated in IP header of inner packet. That is, packet sent from UE to Internet and the opposite way.

Encaps header DSCP markings are updated in IP header of outer IP layer having GTP-U header (Outer header).

DSCP markings are sent from SMF to UPF inside FAR IE as follows:

- Transport Level Marking - The DSCP markings is configured in encaps header for ACCESS side and User-datagram on CORE side for collapse call.

- Transport Level Marking Options—Includes two options and are applicable only for outer header:

  - Copy-inner: Copy the inner packets markings to outer header

  - Copy-outer: Relay the DSCP markings for outer header

Inner Packet Marking—DSCP markings is configured in user datagram for ACCESS side. For CORE side, it's N/A for collapse call.

Logic to fetch the DSCP marking has changed for collapse call:

- Fetch the DSCP markings based on QCI and "qrp_pl" for session from the associated SGW Service for ACCESS/downlink side.

- Fetch the DSCP markings based on QCI and "qrp_pl" for session from the associated PGW Service for CORE/uplink side.

- For ACCESS/downlink side, QCI-QOS-mapping table associated with APN-profile takes preference over SGW Service QCI-QOS-mapping table.

- For CORE/uplink side QCI-QOS-mapping table associated with APN config takes preference over PGW Service QCI-QOS-mapping table.

- In case SGW Service QCI-QOS-mapping table isn't configured, then PGW Service QCI-QOS-mapping table is applicable on both ACCESS/CORE side.

- In case PGW Service QCI-QOS-mapping table isn't configured, then SGW Service QCI-QOS-mapping table is applicable on ACCESS side and no DSCP markings are applicable on CORE side.

- DSCP markings are updated on UPF in create/update FAR sent as a part of Sx/N4 Establishment/Modification request from SMF to UPF.

- Update the TLM, IPM, and TLMO in case of HO from Pure-P to Collapse and vice versa in Sx/N4 Modification request as a part of Update FAR IE.

- Update the Layer2 markings in case of HO from Pure-P to Collapse and vice versa in Sx/N4 Modification request as a part of Update FAR IE.

Following table depicts the various possible config combinations and outcome for DSCP markings to be applied on ACCESS and CORE side for COLLAPSE call:

S. No.	Feature Enable / Disable	PGW Service QOS-QCI Table Configured (Q1)	SGW Service QOS-QCI Table Configured (Q2)	APN QOS-QCI Table Configured (Q3)	APN-Profile QOS-QCI Table Configured (Q4)	ACCESS/Downlink DSCP Markings for Collapse Call	CORE/Uplink DSCP Markings for Collapse Call
1	ENABLE	YES	YES	YES	YES	Q4 (APN-Profile)	Q3(APN)
2	ENABLE	YES	YES	YES	NO	Q2 (SGW-Service)	Q3(APN)
3	ENABLE	YES	YES	NO	YES	Q4 (APN-Profile)	Q1 (PGW-service)
4	ENABLE	YES	YES	NO	NO	Q2 (SGW-Service)	Q1 (PGW-service)
5	ENABLE	YES	NO	YES	YES	Q4 (APN-Profile)	Q3(APN)
6	ENABLE	YES	NO	YES	NO	Q3(APN)	Q3(APN)
7	ENABLE	YES	NO	NO	YES	Q4 (APN-Profile)	Q1 (PGW-service)
8	ENABLE	YES	NO	NO	NO	Q1 (PGW-service)	Q1 (PGW-service)
9	ENABLE	NO	YES	YES	YES	Q4 (APN-Profile)	Q3(APN)
10	ENABLE	NO	YES	YES	NO	Q2 (SGW-Service)	Q3(APN)
11	ENABLE	NO	YES	NO	YES	Q4 (APN-Profile)	N/A (NO DSCP)
12	ENABLE	NO	YES	NO	NO	Q2 (SGW-Service)	N/A (NO DSCP)
13	ENABLE	NO	NO	YES	YES	Q4 (APN-Profile)	Q3(APN)
14	ENABLE	NO	NO	YES	NO	Q3(APN)	Q3(APN)
15	ENABLE	NO	NO	NO	YES	Q4 (APN-Profile)	N/A (NO DSCP)
16	ENABLE	NO	NO	NO	NO	N/A (NO DSCP)	N/A (NO DSCP)
17	DISABLE	YES	YES	YES	YES	Q3(APN)	Q3(APN)
18	DISABLE	YES	YES	YES	NO	Q3(APN)	Q3(APN)
19	DISABLE	YES	YES	NO	YES	Q1 (PGW-service)	Q1 (PGW-service)
20	DISABLE	YES	YES	NO	NO	Q1 (PGW-service)	Q1 (PGW-service)
21	DISABLE	YES	NO	YES	YES	Q3(APN)	Q3(APN)
22	DISABLE	YES	NO	YES	NO	Q3(APN)	Q3(APN)

S. No.	Feature Enable / Disable	PGW Service QOS-QCI Table Configured (Q1)	SGW Service QOS-QCI Table Configured (Q2)	APN QOS-QCI Table Configured (Q3)	APN-Profile QOS-QCI Table Configured (Q4)	ACCESS/Downlink DSCP Markings for Collapse Call	CORE/Uplink DSCP Markings for Collapse Call
23	DISABLE	YES	NO	NO	YES	Q1 (PGW-service)	Q1(PGW-service)
24	DISABLE	YES	NO	NO	NO	Q1 (PGW-service)	Q1 (PGW-service)
25	DISABLE	NO	YES	YES	YES	Q3(APN)	Q3(APN)
26	DISABLE	NO	YES	YES	NO	Q3(APN)	Q3(APN)
27	DISABLE	NO	YES	NO	YES	N/A (NO DSCP)	N/A (NO DSCP)
28	DISABLE	NO	YES	NO	NO	N/A (NO DSCP)	N/A (NO DSCP)
29	DISABLE	NO	NO	YES	YES	Q3(APN)	Q3(APN)
30	DISABLE	NO	NO	YES	NO	Q3(APN)	Q3(APN)
31	DISABLE	NO	NO	NO	YES	N/A (NO DSCP)	N/A (NO DSCP)
32	DISABLE	NO	NO	NO	NO	N/A (NO DSCP)	N/A (NO DSCP)

Configuration

Configure the following command inside the SAEGW service to enable/disable this feature.

```
configure
  context egress
    saegw-service saegw_service_name
      downlink-dscp-per-call-type [ enabled | disabled ]
      end
```

Note   For collapsed call, when you enable the feature, use the S-GW service QCI-QOS mapping DSCP markings for downlink. By default, the downlink-DSCP-per-call-type is Disabled.

Monitoring and Troubleshooting

This section provides information on CLI commands that are available for monitoring and troubleshooting for DSCP markings for collapse calls.

Show Commands Outputs

This section provides information about show CLI commands that are available in support of DSCP markings for collapse calls.

show saegw-service all

This show command is to check if the feature is enabled or Disabled.

```
Service name : SAEGW11
Service-Id : 47
Context : EPC1
Status : STARTED
sgw-service : SGW11
pgw-service : PGW11
sx-service : SX11C
User Plane Tunnel GTPU Service : SAEGW11SXU
Newcall policy : n/a
downlink-dscp-per-call-type : enabled
CUPS Enabled : Yes
Service name : SAEGW21
Service-Id : 25
Context : EPC2
Status : STARTED
sgw-service : SGW21
pgw-service : PGW21
sx-service : SX21C
User Plane Tunnel GTPU Service : SAEGW21SXU
Newcall policy : n/a
downlink-dscp-per-call-type : disabled
CUPS Enabled : Yes
```

show sub user-plane-only callid <call_id> far full all

Use this User Plane CLIs to validate the Transport level marking options and inner packet markings for UPLINK/DOWNLINK FAR.

show sub user-plane-only full all

Use this User Plane CLI to see the number of TOS marked packets for U/L and D/L.

Dynamic and Static PCC Rules

This chapter covers the following topics:

Feature Summary and Revision History

Summary Data

Table 30: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

Table 31: Revision History

Revision Details	Release
Support has been added for flow-level policing.	2021.01.0
The maximum number of groups that can be configured per bandwidth policy has been increased.	2021.01.0
First introduced.	2020.02.0

Feature Description

Dynamic PCC rules are provisioned by the PCF to the PCEF via the HTTP interface and may be either predefined/static or dynamically generated in the PCF. Dynamic PCC rules can be installed, modified and removed at any time.

Predefined PCC rules are configured in the PCEF and can be activated or deactivated by the PCF or by the PCEF at any time. Static PCC rules within the PCEF may be grouped allowing the PCF to dynamically activate a set of static PCC rules over the HTTP reference point. Those static PCC rules to be locally activated by the PCEF are not explicitly known in the PCF, but the PCF simply knows identifiers of static PCC rules to be activated from the PCF.

How it Works

Predefined PCC Rules Support

Config URR IDs are applicable for static rules and also predefined rules. When a subscriber call comes up, it traverses the static rules in rule base. The subscriber master URR list with bucket IDs as key updates the corresponding URR buckets for the various interfaces with the charging action configuration. For dynamic rules and predefined rules, URR ID list in PDR creates the URR buckets on the User Plane.

Following are the ecosystem changes to support Cisco SMF and UPF to work independently for Charging Action (vendor agnostic way) to work:

- Configurable "Config URR IDs" at UPF

- UPF to enable the local configuration for thresholds

Provisioning of Predefined PCC Rules

Predefined PCC rule is preconfigured in the SMF (for 5GC). Predefined PCC rules can be activated or deactivated by the PCF at any time. The Predefined PCC rules may be grouped allowing the PCF to dynamically activate a set of PCC rules. The SMF may enforce an activated predefined PCC rule by the PCF in the UPF by:

- Determining the service data filters or application IDs referred by the activated predefined PCC rule(s) and the corresponding QoS and charging control information respectively.

- Creating the necessary PDR(s) to identify the service data flow(s), application(s) that the predefined PCC or ADC rule refer to, if not already existing.

- Creating the necessary QER for the QoS enforcement at service data flow or application-level accordingly.

- Creating the necessary FAR if a new FAR needs to be created as result of QoS flow binding and QoS control for forwarding the detected service data flow or application traffic, or to redirect or to apply traffic steering control if included in the predefined PCC rule.

- Creating the necessary URR(s) for each monitoring key, charging key, combination of charging key and service ID, or combination of charging key, sponsor ID and Application Service Provider ID if included in the predefined PCC rule.

And, later by:

- Associating the created URR(s) to the newly created PDR(s).

- Associating the existing FAR or the new FAR to the newly created PDR(s).

Optionally, the traffic handling policies common to many PFCP sessions (that is, predefined QER(s)/FAR(s)/URR(s)) can be configured in the UPF. The SMF activates these traffic handling policies by including the Activate Predefined Rules IE within one of the following:

- The Create PDR IE in an PFCP Session Establishment Request

- The Create PDR IE in an PFCP Session Modification Request

For traffic matching PDR(s) associated with the activated predefined rules, the UPF enforces the rules. For example, the UPF generates Usage Report(s) and sends it to the SMF, for URR, and the SMF handles the usage reports.

The URR IDs used in reports triggered by a predefined rule in UPF are also preconfigured at the SMF.

Dynamic PCC Rules Support

For dynamic PCC rules multiple flows are supported on per Packet Forwarding Control Protocol (PFCP) session:

- The 5G QoS model allows classification and differentiation of specific services based on subscription-related and invocation-related priority mechanisms. These mechanisms provide abilities such as invoking, modifying, maintaining, and releasing QoS Flows with priority, and delivering QoS Flow packets according to the QoS characteristics under network congestion conditions.

- The 5G QoS model is based on QoS Flows. The 5G QoS model supports both QoS Flows that require guaranteed flow bit rate (GBR QoS Flows) and QoS Flows that do not require guaranteed flow bit rate (Non-GBR QoS Flows).

- The QoS Flow is the finest granularity of QoS differentiation in the PDU session. A QoS Flow ID (QFI) is used to identify a QoS Flow in the 5G System. User Plane traffic with the same QFI within a PDU session receives the same traffic forwarding treatment (Example - scheduling, admission threshold).

- Within the 5GS, a QoS Flow associated with the default QoS rule is required to be established for a PDU session and remains established throughout the lifetime of the PDU session. This QoS Flow must be a Non-GBR QoS Flow.

- A QoS flow is associated with QoS requirements as specified by QoS parameters and QoS characteristics. A QoS flow can either be "GBR" or "Non-GBR" depending on its QoS profile.

  - For each QoS Flow, the QoS profile includes the QoS parameters:

    - 5G QoS Identifier (5QI)

    - Allocation and Retention Priority (ARP)

  - For each GBR QoS flow only, the QoS profile must also include the QoS parameters:

    - Guaranteed Flow Bit Rate (GFBR) - UL and DL

    - Maximum Flow Bit Rate (MFBR) - UL and DL

  - In the case of a GBR QoS Flow only, the QoS profile may also include one or more of the QoS parameters:

    - Notification control

    - Maximum Packet Loss Rate - UL and DL

During PDR creation or modification UPF receives the QER for QoS enforcement on flows.

The QoS enforcement rule correlation ID is assigned by the CP function to correlate QERs from multiple PFCP session contexts. For instance, the enforcement of APN-AMBR in the PGW-U is achieved by setting the same QoS enforcement rule correlation ID to the QERs from different PFCP sessions associated with all the PDRs corresponding to the non-GBR bearers of all the UE's PDN connections to the same APN. The QERs that are associated to the same QoS Enforcement Rule Correlation ID in multiple PFCP sessions will be provisioned with the same QER contents in each of these PFCP sessions. The QoS enforcement rule correlation ID is only used to enforce the APN-AMBR when the UE is in EPC, it may be provided by the CP function over N4 to the UP function for a PDU session may move to EPC in a later stage.

If the UPF receives QoS Enforcement Rule Correlation ID for 5G PFCP sessions, then it will enforce it.

Policing

The policer configuration uses inputs from the session manager, these inputs are received either from PCF as AMBR or from flow-level QoS information. The values received from the PCF are always accepted for session-level AMBR policing. However, the flow-level policing is prioritized, if available, and AMBR policing is applied sequentially. That is to say, the policer engine applies the hierarchical policing—first the flow-level/rule bandwidth limiting and then the session-level bandwidth limiting.

Note AMBR modifications during session run-time through RAR or CCA-U is applicable.

The input values received from the session manager are pushed into a policer configuration and a policer token bucket. For each direction - uplink or downlink, a new record is created for Policer configuration and Policer token bucket.

The Policer configuration is the reference for the policer engine, and the policer token bucket is used for calculation and restoration of values.

Currently, Policing is supported for AMBR received from PCF and rule-level QoS information for dynamic rules. For static and predefined rules, bandwidth limiting is achieved by the bandwidth policy configuration. Extended bit rates configured in bandwidth-policy configuration in Active Charging Service Configuration mode on SMF is provided to the UPF by RCM, and same is applied for policing by the UPF. An example configuration of bandwidth policy, with extended bit rate, is given below:

```
configure
  active-charging service ACS
    bandwidth-policy BWP

      flow limit-for-bandwidth id 1 group-id 2

      flow limit-for-bandwidth id 2 group-id 3
      flow limit-for-bandwidth id 100 group-id 100

      group-id 2 direction uplink peak-data-rate 256000 peak-burst-size 1000 violate-action
 discard
      group-id 3 direction downlink peak-data-rate 256000 peak-burst-size 1000 violate-action
 discard
      group-id 4 direction uplink peak-data-rate 300000 peak-burst-size 1200 violate-action
 lower-ip-precedence
      group-id 5 direction downlink peak-data-rate 300000 peak-burst-size 1200 violate-action
 lower-ip-precedence committed-data-rate 256000 committed-burst-size 1000 exceed-action
lower-ip-precedence
      group-id 100 direction downlink peak-data-rate-kbps 4294967295 peak-burst-size
4294967295 violate-action discard
      group-id 100 direction uplink peak-data-rate-kbps 4294967295 peak-burst-size 4294967295
 violate-action discard
      exit
    charging-action catchall
      flow limit-for-bandwidth id 1
      exit
    rulebase cisco
      bandwidth default-policy BWP
      exit
    end
```

Limitations

In this release, Policing has the following limitations:

- Modification of **bandwidth-policy** isn't supported.

- Interaction with other features, such as token replenishment (both APN-level and ACL-level) isn't supported.

- Currently, policer-based statistics aren't supported. You can verify bandwidth limiting using network performance monitoring tools.

Bandwidth Policy Configuration Limits

The UPF expects the user to configure the bandwidth limits in both SMF and UPF, for both downlink and uplink packets, in all charging actions of predefined PCC rules, even if the bandwidth limitation configuration is the same for all the charging actions.

To optimize these configurations, the user has to define a bandwidth ID to include all bandwidth-related configurations and associate the bandwidth ID under the charging actions.

If the bandwidth value is changed, the new subscribers use the configured bandwidth values while the existing subscribers continue to use the old values.

The following are the bandwidth-policy configuration limits:

- Maximum number of bandwidth policies that can be configured: 64.

- Maximum number of Groups per bandwidth policy that can be configured: 1000.

- Maximum number of bandwidth IDs per bandwidth policy that can be configured: 1000.

- Maximum number of Groups across bandwidth policies that can be configured: 10000.

- Maximum number of bandwidth IDs across bandwidth policies that can be configured: 10000.

Rate Limiting for Static and Predefined Rules

For static and predefined rules, bandwidth limiting is achieved by the bandwidth policy configuration. Bandwidth Policy must be configured on SMF and UPF under Active Charging Service Configuration Mode.

The following is an example configuration of bandwidth policy with extended bit rate:

```
config
   active-charging service ACS
      bandwidth-policy BWP
         flow limit-for-bandwidth id 1 group-id 2
         flow limit-for-bandwidth id 2 group-id 3
         flow limit-for-bandwidth id 100 group-id 100
        group-id 2 direction uplink peak-data-rate 256000 peak-burst-size 1000 violate-action
 discard
         group-id 3 direction downlink peak-data-rate 256000 peak-burst-size 1000
violate-action discard
        group-id 4 direction uplink peak-data-rate 300000 peak-burst-size 1200 violate-action
 lower-ip-precedence
         group-id 5 direction downlink peak-data-rate 300000 peak-burst-size 1200
violate-action
         lower-ip-precedence committed-data-rate 256000 committed-burst-size 1000
exceed-action
         lower-ip-precedence
         group-id 100 direction downlink peak-data-rate-kbps 4294967295 peak-burst-size
4294967295 violate-action discard
         group-id 100 direction uplink peak-data-rate-kbps 4294967295 peak-burst-size
4294967295 violate-action discard
         exit
      charging-action catchall
         flow limit-for-bandwidth id 1
         exit
      rulebase cisco
```

```
bandwidth default-policy BWP
exit
end
```

Note   The modification of bandwidth-policy configuration is not supported.

Rate Limiting for Dynamic Rules

As per 3GPP TS 29.244, the following IE is received from SMF for QoS enforcement in Create QER or Update QER in Session Establishment or Modification Request:

- Maximum Bitrate: This IE is present if an MBR enforcement action is applied to packets matching this PDR. When present, this IE indicates the uplink and/or downlink maximum bit rate to be enforced for packets matching the PDR. For 5GC, this IE may be set to the value of:

  - the Session-AMBR - for a QER that is referenced by all the PDRs of the non-GBR QoS flows of a PDU session.

  - the QoS Flow MBR - for a QER that is referenced by all the PDRs of a QoS Flow.

  - the SDF MBR - for a QER that is referenced by all the PDRs of an SDF.

- Guaranteed Bitrate: This IE is present if a GBR has been authorized to packets matching this PDR. When present, this IE indicates the authorized uplink and/or downlink guaranteed bit rate. This IE may be set to the value of:

  - the aggregate GBR - for a QER that is referenced by all the PDRs of a GBR bearer

  - the QoS Flow GBR - for a QER that is referenced by all the PDRs of a QoS Flow

  - the SDF GBR - for a QER that is referenced by all the PDRs of an SDF

- QoS flow identifier (QFI): This IE is present if the QoS flow identifier is inserted by the UPF.

- Gate Status: This IE indicates whether the packets are allowed to be forwarded (the gate is open) or it is discarded (the gate is closed) in the uplink and/or downlink directions.

- QER Correlation ID: This IE is present if the UP function is required to correlate the QERs of several PFCP sessions, for APN-AMBR enforcement of multiple UE's PDN connections to the same APN.

Note   Although it is not applicable, but if UPF receives QoS Enforcement Rule Correlation ID for 5G PFCP sessions then it will enforce it.

The SMF provisions QoS enforcement in UPF by creating necessary PDRs to represent SDF, QoS Flow and session and associating respective QERs as follows:

- creating QERs for the QoS enforcement at session level, SDF level.

- creating QERs for the QoS enforcement of the aggregate of SDFs with the same GBR QFI.

- associating the session level QER to all the PDRs defined for the session.

- associating the SDF or application QER to the PDRs associated to the SDF or application.

- associating the QER of the aggregate of SDFs to the PDRs associated to SDFs or applications that share the QER.

Standards Compliance

The N4 interface between SMF and UPF is specified in 3GPP TS 23.501 and 3GPP TS 23.502.

Configuring the URR IDs

Following are the steps to achieve the configurable URR IDs:

- Configuration template outside of Charging action to allow URR-Id mapping with "Rating Group" and "Service-ID".

- If a separate RG is configured for Gy, then that RG is applied for Gy bucket. If no separate RG is configured for Gy, then the same Content-id applicable for all interfaces.

- "Service-ID" would be optional for URR-ID mapping.

- URR-ID should be unique (this need to be ensured through **show configuration error** or separate script to validate. Another option would be to check during config time itself, provided this should not lead to bigger configuration loading time). The actual URR ID value on N4 interface will have additional bits along with the config URR ID value.

- For UPF, current logic for URR-ID generation need to be updated to take value from configuration. There are no changes for URR usage/generation logic/call-flow except UPF receiving config URR-ID from configuration rather than PFD message.

- Same configuration values are required at SMF as well. Configuration mistake of SMF and UPF having different URR-ID for same mapping will be avoided once common configuration point to SMF and UPF is available/enabled.

To configure URR-IDs, perform the following steps:

```
configure
    active-charging service service_name
        urr-list list_name
            rating-group group_number { service-identifier service_number | urr-id
    id_range }
            end
```

Note	• **urr-list** *list_name:* Configures the active charging service URR list. *list_name* must be an alphanumeric string of 1 to 63 characters.
	• **rating-group** *group_number* **:** Specifies the rating ID used in prepaid charging. *group_number* must be an integer in the range of 0 to 2147483647.
	• **service-identifier** *service_number* **:** Specifies the number given to the service.
	• **urr-id** *id_range* **:** Specifies the URR identifier for rating/service group. *id_range* must be an integer in the range of 1-134217727.

Threshold Configuration

The GTPP group configuration is required for threshold calculation at UPF.

UPF uses GTPP group name available from APN configuration. Only one GTPP group should be associated under APN configuration.

The following is a sample configuration:

```
configure
  context context_name
    apn apn_name
      gtpp group group_name
      ip context-name name
      exit
    gtpp group group_name
      gtpp egcdr service-data-flow threshold interval interval
      gtpp egcdr service-data-flow threshold volume downlink bytes
      gtpp egcdr service-data-flow threshold volume uplink bytes
      gtpp egcdr service-data-flow threshold total bytes
      end
```

If any one of the above service-data-flow thresholds is hit for offline URR, the UPF sends SX_SESSION_REPORT_REQUEST towards SMF reporting the data volume.

C H A P T E R 16

GTP-U Support

This chapter covers the following topics:

Feature Summary and Revision History

Summary Data

Table 32: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Table 33: Revision History

Revision Details	Release
Optimization of UDP checksum is added in this release.	2021.02.0
First introduced	2020.02.0

Feature Description

3GPP specifies provisions for UEs capable of supporting both 5G and 4G NAS to connect to E-UTRAN and 5G core network.

To forward data (G-PDUs and End Marker packets) during an EPS to 5GS handover, the SMF:

- Provisions one PDR per E-RAB (that supports data forwarding for at least one QoS flow).

- Creates and associate one QER with each PDR, including the QFI IE set to the QFI value of one of the QoS flows mapped to the E-RAB, to request the UPF to insert a GTP-U PDU Session Container extension header including the QFI.

Data forwarding during handovers between 5GS and EPS is supported as follows (see, 3GPP TS 38.300):

- For 5G to 4G handover, the source NG-RAN node sends one or several end-markers including one QFI of those QoS flows mapped to the same E-RAB and sends the end-marker packets to the UPF over the PDU session tunnel. UPF removes the QFI and maps to an appropriate E-RAB tunnel towards SGW.

- For 4G to 5G handover, the source eNB forwards the received end markers in the EPS bearer tunnel to the SGW, which forwards them to the UPF. The UPF adds one QFI among the QoS flows mapped to that E-RAB to the end-markers and sends those end-markers to the target NG-RAN node in the per PDU session tunnel.

Error Indication and GTP-U Path Failure

The UPF notifies an Error Indication message for a GTP-U peer to the sender when a GTP-PDU is received with a TEID that does not exist. This ensures that there are no stale sessions or bearers, and maintains consistency in the network.

Error Indication and GTP-U Path Failure between SMF and UPF nodes are supported over N4 interface. For the neighbor nodes, it is supported over the S1u/S5u interfaces.

Behavior variations of local-purge or signal-peer for Error Indication and GTP-U Path Failure are considered in this implementation.

- When Error Indication is received, the UPF communicates the TEID and GTPU-peer information with the SMF to ensure deletion or modification of the GTPU-peer.

- On receiving GTP-U packet with non-existing TEID, the UPF generates and sends Error Indication with TEID and GTP-U peer entries.

- The deletion of a session or a bearer is decided based on the Path Failure detection at SMF or UPF.

- GTP-U Path Failure is detected using GTP-U echo messages between UPF nodes, and between the UPF and SMF nodes.

How it Works

Call Flows

Initial Attach on E-UTRAN via MME and S-GW

Initial attach on E-UTRAN/EPS follows the procedure defined in 3GPP TS 23.401, Section 5.3.2.1.

The following diagram shows the call flow derived from 3GPP reference for initial attach on E-UTRAN/EPS.

Table 34: Initial Attach on E-UTRAN via 5G Core Call Flow

Step	Description
1	At Step 9, SMF+PGW-C perform a UPF selection and perform N4 Session Establishment procedure. Since this session is a 4G session connecting to SMF+PGW-C, separate CN tunnel is created for each bearer and QFI is not sent in the QER and PDR, correlation ID might be present.
2	At Step 18, SMF+PGW-C performs N4 Session Modification to update the eNodeB TEID on the data path to the UPF.

The 3GPP specifications provide mechanisms to achieve mobility of a UE from LTE to 5G NR and vice versa. This mobility is achieved in two different architectures – with and without N26 interface between AMF and MME.

5G to EPS Handover with N26 Interface

5G to EPS handover with N26 interface is defined in 3GPP TS 23.502, Section 4.11.1.2.1. The following diagram shows the detailed call flow for N26 interface.
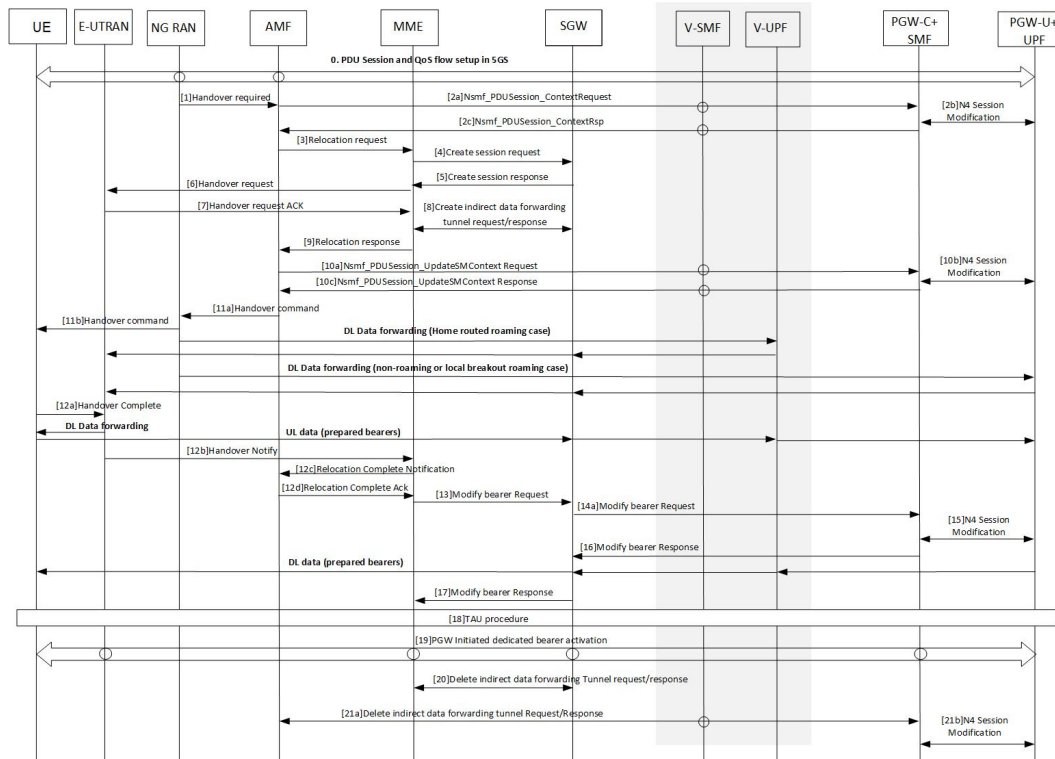
Table 35: 5G to EPS Handover with N26 Interface Call Flow

Step	Description
1	In Step 2b, the SMF+PGW-C sends the N4 Session modification to the UPF to establish the CN tunnel for each EPS bearer. The bearer mapping to the 5G QoS and PCC rules received from PCC is already present with SMF. The SMF also contains the bearer IDs obtained from the Bearer ID Allocation procedure. SMF+PGW-C creates new PDRs for the N4 session and gets TEID allocated for each bearer as required by the 4G system.
2	In Step 10b, SMF+PGW-C sends N4 Modification Request to UPF to create additional PDRs and FARs to receive the redirected DL data over the indirect tunnel from NG RAN and forward them to eNodeB. The uplink PDRs in this case has the QFI to match forwarded DL data from NG RAN and the associated QER does not have the QFI as data needs to be forwarded to eNodeB. Also, the FAR redirects the received data to eNodeB over appropriate tunnel based on the QFI.
3	At Step 11, for the QoS flows indicated in QoS Flows for Data Forwarding, NG-RAN initiates data forwarding through the UPF based on the CN Tunnel Info for Data Forwarding per PDU Session. Then the UPF maps data received from the data forwarding tunnel(s) in the 5GS to the data forwarding tunnel(s) in EPS and sends the data to the target eNodeB through the Serving GW.

Step	Description
4	In Step 15, the SMF sends N4 Modification Request to UPF to activate the DL data path to E-UTRAN. At this time, both the indirect tunnel and the direct DL path are activated towards eNodeB.
5	At Step 21, the SMF sends N4 Modification Request to the UPF to delete the indirect forwarding tunnel.

Other call flows related to EPS to 5G and 5G to EPS handover with N26 interface, or without N26 interface are defined in 3GPP 23.502, Section 4.11.1.2.1 and Section 4.11.2.

Error Indication Handling on UPF

UPF, on receiving Error Indication, initiates a PFCP Session Report Request with Error Indication Report that includes remote F-TEID containing TEID and GTP-U Peer address.

- For PGW-U, Error Indication message is sent or received over S5u.

- For SAEGW-U, Error Indication message is sent or received over S1u.

- For SGW-U, Error Indication message is sent and received over S1u and S5u.

UPF generates Error Indication with TEID and GTP-U Peer Address towards a peer when a data packet is received with TEID for which a session or bearer doesn't exist.

GTP-U Path Failure Support at UPF

GTP-U Echo Requests is initiated and sent periodically as per the configured interval on UPF. GTP-U Echo Response is sent for the GTP-U Echo Request received from SMF over GTP-U tunnel.

If Response is not received for the GTP-U Echo Request, the UPF retries Echo Requests based on configured retransmission timeout and maximum retries. When retries are exhausted, the UPF initiates PFCP node

Report Request including (Node ID, Node Report Type, User Plane Path Failure Report including Remote GTP-U Peer).

If UPF receives PFCP Node Report Response and PFCP Session Deletion Request to delete the session, it responds to the deletion request with usage reports.

Disabling UDP Checksum

This functionality disables the UDP checksum in UDP header of the GTP-U packet. The value of the UDP checksum is set to zero.

Disabling UDP Checksum

Use the following configuration to disable the UDP checksum in UDP header of the GTP-U packet.

```
configure
   context context_name
      gtpu-service service_name
```

```
no udp-checksum
end
```

Heartbeat Support for N4/Sx Interface

Feature Summary and Revision History

Summary Data

Table 36: Summary Data

Applicable Product (s) or Functional Area	5G-UPF
Applicable Platforms	VPC-SI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

Revision Details	Release
First introduced.	2021.01.0

Feature Description

In accordance with 3GPP TS 29.244, support has been added for node-level Heartbeat procedures between the Session Management Function (SMF) and User Plane Function (UPF) over N4/Sx interface.

The Heartbeat procedure contains the following two messages:

1. Heartbeat Request

2. Heartbeat Response

Heartbeat Request

The SMF or the UPF sends a Heartbeat Request on a path to the peer node to find out if it is alive. The Heartbeat Request messages are sent for each peer with which a Packet Forwarding Control Protocol (PFCP) control association is established.

For each peer with which a PFCP control association is established, a SMF or UPF is prepared to receive a Heartbeat Request at any time, and replies with a Heartbeat Response.

Heartbeat Response

This message is sent as a response to a Heartbeat Request.

How It Works

The SMF and UPF sends Heartbeat messages after configurable time duration. If the peer does not respond, the message is retried for configured number of times with the retry-interval and then the configured action is taken for the calls associated with the corresponding peer.

Recovery Time Stamp Information Element (IE), which contains the start time of the node, is supported by both Heartbeat Request and Heartbeat Response. Heartbeat Request contains its own Recovery Time Stamp value and sends it to the peer while Heartbeat Response contains the peers Recovery Time Stamp value.

Path Failure Detection

Path failure is detected in following conditions:

1. Heartbeat failure: This condition occurs when the peer does not respond to the Heartbeat that is sent and also retires.

2. Recovery Time stamp change in Heartbeat: This condition occurs when the Heartbeat Request or Heartbeat Response has a new larger value than the previously received value.

3. Recovery Time stamp change in N4/Sx Association message: This condition occurs when the N4/Sx association message is received again from the peer with a new Recovery Time Stamp.

Path Failure Handling

When the Recovery Time Stamp value received is more than the previously received value, then the peer restart is detected. If the Recovery Time Stamp value is lower than the previously received value then the value is ignored and peer restart is not detected.

When a peer restart is detected, an SNMP Trap is generated to indicate the path failure for the peer. Also, based on the path failure configuration (refer), all the calls connected to that peer can be cleared.

Configuring Heartbeat for N4/Sx Interface

This section provides information about the CLI commands available in support of this feature.

Enabling Heartbeat for Sx Interface

Use the following commands under Sx Service Configuration mode to enable Heartbeat parameters for N4/Sx interface.

```
configure
  context context_name
    sx-service service_name
      [ default ] sx-protocol heartbeat { interval seconds |
max-retransmissions number | path-failure detection-policy {
control-recovery-timestamp-change | heartbeat-retry-failure |
heartbeat-recovery-timestamp-change } | retransmission-timeout seconds }
      no sx-protocol heartbeat { interval | path-failure detection-policy
 { control-recovery-timestamp-change | heartbeat-retry-failure |
heartbeat-recovery-timestamp-change }
        end
```

Notes:

- **default**: Sets/restores default value assigned for specified parameter.

- **no**: Disables the followed option.

- **heartbeat**: Configures N4/Sx Heartbeat parameters.

- **interval** *seconds*: Configures Heartbeat interval (in seconds) for N4/Sx Service. *seconds* must be an integer in the range of 1 to 3600.

- **max-retransmissions** *number*: Configures maximum retries for N4/Sx Heartbeat request. Must be followed by integer, ranging from 0 to 15. Default is 4.

- **retransmission-timeout** *seconds*: Configures the Heartbeat retransmission timeout for N4/Sx service, in seconds, ranging from 1 to 20. Default is 5.

- **path-failure**: Specifies the policy to be used when path failure happens through Heartbeat request timeout.

Configuring Detection Policy for Path Failure

Use the following commands under Sx Service Configuration mode to specify detection policy to be used for path failure.

```
configure
  context context_name
    sx-service service_name
      [ default | no ] sx-protocol heartbeat path-failure detection-policy
{ control-recovery-time-stamp-change | heartbeat-retry-failure |
heartbeat-recovery-
timestamp-change }
        end
```

NOTES:

- **default**: Sets/restores default value assigned for specified parameter.

- **no**: Disables the followed option.

- **detection-policy**: Specifies the policy to be used. Default action is to do cleanup upon Heartbeat request timeout.

- **control-recovery-time-stamp-change**: Path failure is detected when the recovery timestamp in control request/response message changes.

- **heartbeat-retry-failure**: Path failure is detected when the retries of Heartbeat messages times out.

- **heartbeat-recovery-timestamp-change**: Path failure is detected when the recovery timestamp in Heartbeat request/response message changes.

Monitoring and Troubleshooting

This section provides information about CLI commands available for 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 sx-service all

The output of this show command has been enhanced to include the following fields introduced in support of this feature.

- SX Heartbeat

    - Interval

    - Retransmission Timeout

    - Max Retransmission

- SX path failure detection policy

- Heartbeat Timeout

- Heartbeat Req/Rsp Recovery timestamp change

- Control Msg Recovery timestamp counter change

show sx-service statistics all

The output of this show command has been enhanced to include the following fields introduced in support of this feature.

- Heartbeat Request

    - Total TX

    - Total RX

    - Initial TX

    - Initial RX

    - Retrans TX

- Heartbeat Response

    - Total TX

    - Total RX

Disconnect Reasons

The following disconnect reason has been added in support of this feature:

- sx-path-failure - When the Recovery timestamp changes or heartbeat failure is detected, based on the configuration, calls are cleared with this disconnect reason.

SNMP Traps

The following SNMP traps have been added in support of this feature:

- SxPathFailure - This trap is generated when the peer path failure is detected.

- SxPathFailureClear - This trap is generated when the path is restored for the peer.

CHAPTER 18

Idle Mode Buffering and Paging

This chapter covers the following topics:

Feature Summary and Revision History

Summary Data

Table 37: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Table 38: Revision History

Revision Details	Release
First introduced.	2020.02.0

Feature Description

A Buffering Action Rule (BAR) provides instructions to control the buffering behavior of the User Plane Function (UPF). The BAR controls the buffering behavior for all Forwarding Action Rules (FARs) of the Packet Forwarding Control Protocol (PFCP) session. This control is applicable when the PFCP session is set with an Apply Action parameter, which requests packets to be buffered and associated with the respective BAR.

How it Works

If the User Plane Function indicates the support of the feature UL or DL Buffering Control (UDBC), the SMF provides the buffering packet count IE in a BAR. The buffering count IE is created during a PFCP Session Establishment procedure or a PFCP Session Modification procedure. The SMF modifies it in a subsequent PFCP session modification request, "and" or "or" a PFCP Session Report Response message. The same BAR associates with all the FARs in a PFCP session to indicate that all service data flows in the PFCP session shares the same buffer in the UPF for the PFCP session. One BAR is created per PFCP session.

Provisioning of Buffering Action Rule in the UPF

The SMF provisions multiple buffering parameters in a BAR. It is in Create BAR or Update BAR in various PFCP messages.

Currently, UPF supports the following IE:

- The suggested buffering packet count IE - If the UPF indicates the support of the feature UDBC to indicate the number of packets. It includes both uplink or downlink that the SMF suggests buffering in the UPF, until it receives new instructions from the SMF. Example: when the new quota is granted.

- DL buffering suggested packet count IE – This IE is received with update BAR from SMF in Session Report Response message, if SMF wants more DL packets to be buffered on UPF.

The UPF does not apply the DL buffering duration and DL buffering suggested packet count parameters and deletes these parameters from the BAR (without explicit request from the SMF) when extended buffering of downlink data packets ends in the UPF. The UPF does not apply buffering when it receives the new instruction from the SMF. The buffered packets are either dropped or forwarded following the packet forwarding model and considering that the buffered packets are already processed earlier.

Buffering Action Rule Call Flow

This section describes the provisioning of buffering action rule in the UPF call flow.

Figure 9: Buffering Action Rule

Downlink Data Report for First DL Packet

When instructed to buffer and notify the SMF about the arrival of a DL packet, the UPF notifies the SMF, when it receives a first downlink packet for a given FAR. The UPF notifies the DL packet arrival by sending a PFCP Session Report Request including a Downlink Data Report IE identifying the PDR(s) for which downlink packets was received.

Paging Policy Differentiation

The UPF supports the Paging Policy Differentiation, for each PDR and for each packet that triggers a Downlink Data Notification, the UPF function copies the value of the DSCP in TOS (IPv4) or TC (IPv6) information received in the IP payload in Downlink Data Service Information IE.

For each PDR and for each packet that triggers a Downlink Data Notification, if the QFI of the downlink data packet is available, the QFI is also sent in Downlink Data Service Information IE.

Paging Policy Indicator (PPI)

The SMF sends the PPI value in Create QER or Update QER, if UPF needs to set Paging Policy Indicator in outgoing PDU packets.

Frame Format for the PDU Session User Plane Protocol

Downlink PDU Session Information (PDU Type 0) - This frame format is defined to allow the NG-RAN to receive some control information elements which are associated with the transfer of a packet over the interface. The following figure shows the respective DL PDU SESSION INFORMATION frame.

Figure 10: DL PDU SESSION INFORMATION (PDU Type 0) Format

Bits								Number of Octets
7	6	5	4	3	2	1	0	
PDU Type (=0)				Spare				1
PPP	RQI		QoS Flow Identifier					1
PPI			Spare					0 or 1
Padding								0-3

QoS Flow Identifier (QFI)

Description: When present, the QoS Flow Identifier (QFI) parameter indicates the QoS Flow Identifier of the QoS flow to which the transferred packet belongs.

Value Range: The value range is between 0 to $2^6 -1$.

Field Length: 6 bits.

Paging Policy Presence

Description: The Paging Policy Presence (PPP) parameter indicates the presence of the Paging Policy Indicator (PPI).

Value Range: A value of 0 indicates that Paging Policy Indicator is not present and 1 indicates that Paging Policy Indicator is present.

Field Length: 1 bit.

Paging Policy Indicator

Description: When present, the Paging Policy Indicator (PPI) is used for paging policy differentiation (see details in 3GPP TS 23.501). This field applies to PDU sessions of IP type.

Value Range: the value range is between 0 to $2^3 -1$.

Field Length: 3 bits.

CHAPTER 19

Indirect Forwarding Tunnel

Revision History

Revision Details	Release
First introduced	2021.02.0

Feature Description

The UPF supports Indirect Forwarding Tunnel (IDFT) procedures for creation and deletion, which are applicable for Pure-S and Collapsed calls with dedicated bearers. This feature is applicable for IDFT support with S-GW Relocation.

> **Note**   The IDFT in UPF is a CLI-controlled feature. By default, the IDFT feature in UPF is disabled.

How It Works

Call Flow

The following call flow illustrates the IDFT support with S-GW Relocation.

Figure 11: IDFT Support with S-GW Relocation

The above call flow describes the IDFT tunnels establishment and deletion with S-GW relocation and without MME change.

If IDFT tunnels are not deleted by MME, then S-GW initiates the local delete of IDFT tunnels.

This feature supports the following scenarios for the Pure-S and Collapse calls:

- S-GW relocation with same MME

- S-GW relocation with same MME and different eNodeB

- S-GW relocation with different MME

- S1-based eNodeB Handoff

- EUTRAN to UTRAN Handoff

- EUTRAN to UTRAN Handoff with S-GW relocation

- UTRAN to EUTRAN Handoff

- UTRAN to EUTRAN Handoff with S-GW relocation

- Tracking Area Update (TAU) with S-GW change and indirect data forwarding

- Radio Access Bearer (RAB) during Active IDFT

- Sx transaction timeout during IDFT setup or removal

- Pending Sx transaction (event from PCRF or OCS) and IDFT request comes in

- Create Bearer Request (CBR) during Active IDFT

- Update Bearer Request (UBR) during Active IDFT

- Delete Bearer Request (DBR) during Active IDFT

- Modify Bearer Request (MBR) behavior on other PDN during Active IDFT

- Source MME path failure

- Target MME path failure

- MME path failure with NTSR enabled

- eGTP-C S5 path failure

- eGTP-C S5 path failure with P-GW restart notification enabled

- Sx path failure (clean IDFT and calls)

- Abort session (clear sub all, local abort, and so on.)

- CBR, UBR on other PDN during Active IDFT

- DBR on other PDN/bearer during Active IDFT

- S1-u path failure for target eNodeB

- S1-u path failure for source eNodeB

- S-GW path failure for source S-GW

- S-GW path failure for target S-GW

- S1-u error indication on the default bearer while Active IDFT

- S1-u error indication on the dedicated bearer while Active IDFT

- S1-u error indication from the target S-GW to source S-GW bearer

- S1-u error indication from the target eNodeB to target S-GW bearer

- Sending End Marker when tearing down IDFT tunnel after failure

- If SMF ICSR/SR leads to the cleanup of IDFTs, the UPF also cleans the IDFT PDRs/FARs

5G to 4G Handover with IDFT

In compliance with 3GPP TS 23.502 v15.5.1, the 5G to 4G handover with IDFT is supported in UPF. Refer the "5GS to EPS handover using N26 interface" section in the 3GPP specification for details about call/datapath flow for IDFT.

This functionality isn't CLI-controlled in UPF.

4G to 5G Handover with IDFT

In compliance with 3GPP TS 23.502 v15.5.1, the 4G to 5G handover with IDFT is supported in UPF. Refer the "EPS to 5GS handover using N26 interface" (preparation and execution phase) section in the 3GPP specification for details about the call flow.

This functionality isn't CLI-controlled in UPF.

Supported Functionality

The IDFT feature supports the following functionality:

- Create IDFT request for Collapsed, Pure-S, combination of Collapsed and Pure-S multi-PDN calls with multiple bearers.

- Data transfer on downlink and uplink IDFT bearers.

- Deletion of IDFT request from MME. Also, timer-based deletion of IDFT bearer after expiration of a default value of 100 seconds, if the MME does not send an IDFT request for deletion.

- Deletion of IDFT PDN, including Clear/Delete subscribers from MME/P-GW, when normal PDN goes down.

- IDFT creation of Sx Failure Handling for Pure-S and Collapsed PDN.

☞

Important   Transport GTP-U address capability is assumed to be same across eNodeB and S-GW.

Limitations

The IDFT feature has the following limitations:

- Message interaction and collision during IDFT PDN establishment or deletion with any other procedure is not supported.

- S11/S5 and Sx Path Failure Handling on non-IDFT PDN is not supported when IDFT PDN is Active.

- Deletion of partial dedicated bearers in IDFT connected-state is not supported.

Configuring Indirect Forwarding Tunnel

This section describes the CLI commands available in support of IDFT feature.

Enabling Indirect Forwarding Tunnel Feature

On SMF, use the following CLI commands to enable or disable the IDFT feature.

```
configure
  context context_name
    sgw-service service_name
      [ default | no ] egtp idft-support
      end
```

NOTES:

- **idft-support**: Enables/Disables the IDFT feature in UPF.

- By default, the IDFT feature is disabled and this CLI command is applicable on run-time change.

Verifying the Indirect Forwarding Tunnel Feature

show sgw-service name <service_name>

On SMF, the output of this CLI command has been enhanced to display if the IDFT feature is enabled or disabled.

- IDFT-Feature Support for CUPS: Enabled/Disabled

Monitoring and Troubleshooting

This section provides information regarding the CLI commands available in support of monitoring and troubleshooting the feature.

Show Commands Input and/or Outputs

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

show subscribers saegw-only full all

On UPF, use this command to see the IDFT Local and Remote TEID data. The following is a sample output:

```
Indirect Fwding  : Active
DL fwd local  addr: 1.1.1.4                DL fwd remote   addr: 1.1.1.2
DL fwd local  teid: [0x80028004]           DL fwd remote   teid: [0x2002d2e5]
UL fwd local  addr: 1.1.1.4                UL fwd remote   addr: 1.1.1.2
UL fwd local  teid: [0x8002a004]           UL fwd remote   teid: [0x20042bca]
```

show subscribers user-plane-only callid <call_id> pdr all

On UPF, use this command to see the PDRs created for IDFT. The following is a sample output:

☞

Important   IDFT PDRs will have ACCESS as the source and destination interface type.

```
+-----Source Interface:        (C) - Core        (A) - Access
|-----Type                     (P) - CP-function  (.) - Unknown
|
|+----Destination Interface:   (C) - Core        (A) - Access
||----Type                     (P) - CP-function  (.) - Unknown
||
||
||+---Rule-Type:               (S) - Static       (P) - Predefined
|||---Type                     (D) - Dynamic      (.) - Unknown
|||
|||
vvv   PDR-ID     Associated FAR-ID   Associated URR-ID(s)  Associated QER-ID(s)
---   ------     -----------------   --------------------  --------------------
ACS   0x0001       0x8001                n/a              0x80000001
CAS   0x0002       0x8002                n/a              0x80000001
ACD   0x0003       0x0003             0x00000007          0x00000002
                                        n/a              0x80000003
CAD   0x0004       0x0004             0x00000007          0x00000002
                                        n/a              0x80000003
CAD   0x0005       0x0005             0x00000000          n/a
ACD   0x0006       0x0006             0x00000000          n/a
CAD   0x0007       0x0007             0x00000000          n/a
ACD   0x0008       0x0008             0x00000000          n/a
AAD   0x0009       0x0009             0x00000000          n/a
AAD   0x000A       0x000A             0x00000000          n/a
AAD   0x000B       0x000B             0x00000000          n/a
AAD   0x000C       0x000C             0x00000000          n/a

Total subscribers matching specified criteria: 1
```

Similarly, you can use the **show subscribers user-plane-only callid** *call_id* **far all** CLI command to see the FARs created for IDFT

show subscribers user-plane-only full all

☞

Important Data statistics on IDFT PDRs are captured in the same way as existing PDR statistics. However, it is captured with a limitation – Statistics for DL and UL IDFT will be incremented in Pkts-Down and Bytes-Down category.

The following is sample output:

```
Static & Predef Rule Match stats:
Rule Name            Pkts-Down  Bytes-Down Pkts-Up    Bytes-Up   Hits       Match-Bypassed
  FP-Down(Pkts/Bytes)  FP-Up(Pkts/Bytes)
-------------------- ---------- ---------- ---------- ---------- ---------- --------------
   ------------------- -----------------
catchall                      0          0          3       1368          3              0
        0/0                0/0

Dynamic Rule Match stats:
PDR Id    Pkts-Down  Bytes-Down Pkts-Up    Bytes-Up   Hits       Match-Bypassed
FP-Down(Pkts/Bytes)     FP-Up(Pkts/Bytes)
-------- ---------- ---------- ---------- ---------- ---------- --------------
   ------------------   -----------------
0x0004         2        856          0          0          2              0          0/0
                  0/0
0x000b         2        856          0          0          2              0          0/0
                  0/0
0x000c         2        168          0          0          2              0          0/0
                  0/0
```

CHAPTER 20

LTE - Wi-Fi Seamless Handover

Feature Summary and Revision History

Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI
Feature Default Setting	Enabled - Always-on
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Revision Details	Release
First introduced.	2021.02.0

Feature Description

Seamless handovers between LTE and Wi-Fi (S2b), for UEs that need continuity with their ongoing data session, is supported in the 5G UPF architecture.

When handover is initiated from LTE to Wi-Fi, the Delete Bearer Request (DBR) is sent over the LTE tunnel immediately when the Create Session Response (CSR) is sent on the Wi-Fi tunnel. This causes some packet loss because of the IPSec tunnel establishment delay at the ePDG. To address the issue of packet loss, a Delete Bearer Request is sent on LTE tunnel only on expiry of the configured handover timer. If the LTE tunnel is active, uplink and downlink data are exchanged on the LTE tunnel. When handover is complete, uplink and

downlink data is exchanged on the Wi-Fi tunnel. This prevents packet loss. During Wi-Fi to LTE handover, if the Modify Bearer Request is received with HI=1, it initiates a tunnel switch from Wi-Fi to LTE as per the specification.

With this feature, the following benefits are seen:

- Minimum packet loss during LTE to Wi-Fi (S2bGTP) handover and making the handover seamless (that is, MAKE before BREAK).

- LTE procedures are handled gracefully over the LTE tunnel when both tunnels are established with the P-GW.

- Wi-Fi procedures are handled gracefully over the Wi-Fi tunnel when both tunnels are established with the P-GW.

☞

Important

- In an LTE to Wi-Fi or Wi-Fi to LTE handover, a tunnel identifier is allocated for new access traffic type for experiencing seamless handover.

How It Works

EPC to Non-3GPP Untrusted Wi-Fi Handover Call Flow

This section describes the EPC to non-3GPP untrusted Wi-Fi handover call flow.

Figure 12: EPC to Non-3GPP Untrusted Wi-Fi Handover Call Flow

Table 39: EPC to Non-3GPP Untrusted Wi-Fi Handover Call Flow Description

Step	Description
1	The UE is attached to the 3GPP access network.
	The SMF+PGW-C communicates with UPF, PCF, and CHF for IPv4, IPv6, or dual-stack to establish 4G LTE PDU session. The PCF sends the Policy Control Request trigger, which is the SM policy decision, in response to SM policy control create. The CHF provides session-level or rating-group-level triggers to the SMF in Charging Data Create response.
2	The UE connects to an untrusted non-3GPP access and an ePDG is selected through the ePDG selection process. Then, the UE initiates the handover attach procedure as defined in *3GPP TS 23.402, section 8.6.2.1*. After the IKE tunnel is established between the UE and ePDG and after the UE is authenticated over SWm interface with AAA server, the UE initiates IKE authentication (IKE_AUTH). The IKE_AUTH includes configuration parameters of the earlier assigned IPv4 or IPv6 addresses in the EPC and P-CSCF and the DNS options.

Step	Description
3	The ePDG sends a Create Session Request to the P-GW. This request includes the following details: • IMSI • APN • Handover indication • RAT type • ePDG TEID of the Control Plane • ePDG address for the User Plane • ePDG TEID of the User Plane • EPS bearer identity • User location The RAT type indicates the non-3GPP access technology type. If the UE supports the IP address preservation and is included in the port analyzer adapter (PAA), then the ePDG configures the handover indication in the Create Session Request to allow the PDN gateway to reallocate the same IP address or the prefix assigned to the UE. This IP address or prefix is assigned while UE is connected to the 3GPP IP access and initiates the policy modification procedure with PCF.
4a	The SMF performs UDM registration by updating the PGW-C FQDN with UDM. The UDM registration does not occur during the session establishment with EPC.
4b	The SMF detects the charging triggers with the information available in Step 3 against the charging triggers that are received during EPC session establishment.
4c	The SMF detects the PCF triggers with the information available in Step 3 against the Request Policy Control triggers that are received in the communication with PCF during EPC session establishment.
5	Based on the detected armed Policy Control Triggers that are received in Step 4b, the SMF sends the SM Policy Control Update request with the detected access parameters in Step 3 to the PCF.
6	The PCF includes new or updated PCC rules and sends the SM Policy Control Update response. The Update response includes information on the SM policy decision.
7	Based on the information received in Step 6 and existing policy data of EPC session, SMF prepares the information for the new or updated PCC rules.
8	If new PCC rules are received in Step 6 with new Rating Group that requires quota information, SMF sends the Charging Update request to CHF. SMF also includes new access parameters for the PDU session information.
9	CHF sends the Charging Update Response with multi-unit information that contains quota information for the requested rating-group in Step 8 to SMF. CHF may also send the new quota information for the existing rating-group of EPC session.
10	SMF processes the information that is received as Charging Update response from CHF.

Step	Description
11	SMF sends the N4 session modification request to UPF for Wi-Fi tunnel. This request includes details on creation of uplink PDR, creation of QER, creation of URR for received new rating-group quota information, and update on URR for modified quota information.
12	UPF sends the UL tunnel information that is in created PDR as the N4 session modification response to SMF.
13	SMF sends the GTPv2 Create Session response to S-GW. This response details on request accepted or request accepted partially, P-GW S2b F-TEID, PAA, APN-AMBR, bearer context creation, charging gateway address, and APCO.
14	SMF sends the GTPv2 Create Bearer request to S-GW. This request includes information on bearer context list, which contains DL tunnel information to end-user, to be created.
15	S-GW sends the GTPv2 Create Bearer response to SMF. The response includes details on request accepted or request accepted partially and bearer contexts.
16	SMF processes the Create Bearer response and derives the DL tunnel Information for the established bearer and the the failed EBI list, if any. SMF sends the N4 session modification request to UPF for Wi-Fi tunnel. This request is to create the DL PDR and DL FAR with DL tunnel information for each bearer, RAT modification information, and to delete resources for the 4G tunnel. SMF also deletes the N4 resources of Wi-Fi tunnel for the received failed EBI list or the failed QFI list.
17	UPF sends the usage report as N4 Session Modification response to SMF.
18	SMF+PGW-C sends the GTPv2 DB request to S-GW. This request includes EBI or list of EBIs.
19	S-GW sends the GTPv2 DB response to SMF+PGW-C.
20	SMF sends the Charging Update request to CHF. This request includes the PDU session information with the new access params and multi-usage report containing details on the access params and usage report that is received in Step 8
21	CHF sends the multi-unit information as Charging Update response to SMF. The multi-unit information may include new quota information for the existing rating-groups.
22	SMF sends the SM Policy Control Update request to UPF. This request includes the new access params and rule report for failed QFI list that is received from AMF as part of Create Bearer response. PCF sends the SM policy decision as SM Policy Control Update response. SMF processes the SM policy decision and handles it as PCF Initiation Modify procedure as defined in *3GPP 23.502, section 4.3.3.2*.

Non-3GPP Untrusted Wi-Fi to EPC Handover Call Flow

This section describes the non-3GPP untrusted Wi-Fi to EPC handover call flow.

Figure 13: Non-3GPP Untrusted Wi-Fi to EPC Handover Call Flow

Table 40: Non-3GPP Untrusted Wi-Fi to EPC Handover Call Flow Description

Step	Description
1	One or more PDU sessions are established between UE and ePDG through untrusted non-3GPP access. With the 5G NAS capability of UE, ePDG selects a combined PGW+SMF. UE sends the PDU session ID to the PGW+SMF.
2	UE discovers the E-UTRAN access and hands over the sessions from the currently used non-3GPP access system to E-UTRAN. For details on UE discovery of the 3GPP access system, see *3GPP TS 23.401*, section 4.8.
	UE sends an Attach request to MME for the Handover Attach request type. E-UTRAN routes the messages received from UE to MME as defined in *3GPP TS 23.401*. UE includes the one of the APNs which are corresponding to the PDN connections in the source non-3GPP access. The APN is provided as defined in *3GPP TS 23.401*.

Step	Description
3	MME and HSS perform authentication, which is followed by location update procedure and subscriber data retrieval to receive the APN information.
	The MME selects an APN, an SGW and PDN gateway as defined in *3GPP TS 23.401*. MME sends a Create Session Request message to SGW. This request includes information on IMSI, MME context ID, PDN-GW address, handover indication for the "handover" request type, and APN.
4	SGW sends a Create Session Request, which is handover indication, message to PDN-GW in the HPLMN as described in *3GPP TS 23.401*. As the MME includes the handover indication information in the Create Session Request message, the SGW sends the GTPv2 Create Session Request message to PDN GW. This message includes details on IMSI, APN, handover indication, RAT type, S5-C TEID, S5-U TEID of the user plane, EBI, and user location information. The RAT type indicates the 3GPP IP access E-UTRAN technology type. If the UE supports IP address preservation and is included in PAA, the SGW configures the handover indication in the Creation Session Request. With this configuration, the PDN GW re-allocates the same IP address or prefix that was assigned to the UE while it was connected to the 3GPP IP access. With this configuration, SGW initiates the Policy Modification Procedure to the PCF.
	As the handover indication is includes, the PDN GW does not switch the tunnel from non-3GPP IP access to 3GPP access system at this point.
	SMF does not perform the UDM Registration as the registration happens during the Wi-Fi session establishment.
4a	SMF detects the charging triggers with the information available in Step 3 against the charging triggers that are received during EPC session establishment.
4b	SMF detects the PCF triggers with the information available in Step 3 against the Request Policy Control triggers that are received in the communication with PCF during EPC session establishment.
5	Based on the detected armed Policy Control Triggers that are received in Step 4b, SMF sends the SM Policy Control Update request with the detected access parameters in Step 3 to PCF.
6	PCF sends the SM Policy Control Update response, which is the SM policy decision, by including new or updated PCC rules.
7	Based on the information received in Step 6 and existing policy data of EPC session, SMF prepares the information for the new or updated PCC rules.
8	If SMF receives new PCC rules in Step 6, the SMF sends the Charging Update request, with the new rating-group having quota information, to CHF. This request includes the PDU session information with the new access params.
9	CHF sends the multi-unit information as Charging Update response to SMF. The multi-unit information includes new quota information for the rating-group and the existing rating-group of EPC session, if any.
10	SMF prepares the charging data of the received Charging Update Response that CHF sent.
11	SMF sends the N4 Session Modification Request to UPF. This request includes the details on creation of UL and DL PDR, creation of QER, creation of URR for received new rating-group quota information, updated URR for modified quota information, and creation of FAR.
12	UPF sends the UL tunnel information in the created PDR as N4 Session Modification response to SMF.

Step	Description
13	SMF sends the GTPv2 Create Session response to S-GW. This response details on request accepted or request accepted partially, P-GW S2b F-TEID, PAA, APN-AMBR, bearer context creation, charging gateway address, and APCO.
14	SGW sends the Modification Bearer request with handover indication to PGW for data path switching from Wi-Fi tunnel to 4G tunnel.
15	PGW sends the N4 Session Modification request to delete the Wi-Fi tunnel and to configure DL tunnel information that is received in GTPv2 Create Session request for 4G tunnel in Step 4.
16	UPF sends the N4 Session Modification response to SMF.
17	SMF sends the GTPv2 Create Session request, which includes the bearer context list, to SGW. This list includes the DL Tunnel information for the end-user.
18	SGW sends the GTPv2 Create Session response to SMF. This response includes details on request accepted or request accepted partially and bearer contexts.
19	ePDG sends the GTPv2 Create Bearer resp (accepted EBIs with DL tunnel info to SMF
20	SMF processes the Create Bearer response and derives the DL tunnel Information for the established bearer and the failed EBI list, if any. SMF sends the N4 session modification request to UPF for Wi-Fi tunnel. This request is to update the DL FAR with the DL tunnel information, RAT modification information, and to delete resources for the 4G tunnel. SMF also deletes the N4 resources of Wi-Fi tunnel for the received failed EBI list or the failed QFI list.
21	UPF sends the N4 Session Modification Response with usage report to SMF.
22	SMF sends the Charging Update request to CHF. This request includes the PDU session information with new access params and multi-usage report consisting of access-params and usage report that is received in Step 8.
23	CHF sends the Charging Update Response with multi-unit information that contains quota information for the existing rating-groups to SMF.
24	SMF+PGW-C initiates the GTPv2 DB Request toward SGW by including EBI or EBI list.
25	SGW sends the GTPv2 DB Response toward SMF+PGW-C.
26	SMF sends the SM Policy Control Update request to UPF. This request includes the new access params and rule report for failed QFI list that is received from AMF as part of Create Bearer response. PCF sends the SM policy decision as SM Policy Control Update response. SMF processes the SM policy decision and handles it as PCF Initiation Modify procedure as defined in 3GPP 23.502 section 4.3.3.2.

ICSR and Session Recovery

- At Session Management Function, during transition, the most recent is considered as the stable state and a full checkpoint is triggered once handover is complete from LTE to Wi-Fi (S2BGTP) or vice-versa. This is applicable to Session Recovery and ICSR. User Plane has individual session recovery and ICSR check pointing on every message received.

- During handover failure, that is, when SMF and UPF are out of sync, the SMF session is recovered on the most recently accessed state and UPF is recovered in the new transition state. This behavior is applicable during UPF failure.

Limitations

The LTE - Wi-Fi Seamless Handover feature does not support LTE to eHRPD and Wi-Fi to eHRPD handover and hand back.

Standards Compliance

The LTE – Wi-Fi Seamless Handover feature is compliant with the following standards:

- 3GPP TS 23.214
- 3GPP TS 29.244
- 3GPP TS 23.401
- 3GPP TS 23.402

CHAPTER 21

Monitor Subscriber

Feature Summary and Revision History

Summary Data

Table 41: Summary Data

Applicable Product (s) or Functional Area	5G-UPF
Applicable Platforms	VPC-SI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

> **Note**   Revision history details are not provided for features introduced before release 21.24.

Revision Details	Release
First introduced	Pre 21.24

Feature Description

The Monitor Subscriber (MonSub) feature enables tracing of subscriber related information which includes user and control traffic, and events such as charging and internal events that are useful for debugging. By default, this information is displayed on the Control Plane console, where the user executes MonSub tracing CLI command and captured in a Packet Capture (PCAP) file on the User Plane.

User traffic is carried on slowpath where packets traverse to the application or fastpath where packets do not have to traverse up to the application but are offloaded to fastpath processing (VPP). Slowpath mode was the default mode until fastpath offload (VPP) into SAEGW, was introduced.

Monitor Subscriber provides the following functionality:

- Continuous capture of user traffic from fastpath in PCAP files on the User Plane.

- The non-user traffic information, that is, control event traffic and other related information are displayed in Control Plane console and are captured in separate PCAP files on the User Plane.

- New option UP PCAP trace [W - UP PCAP Trace (ON)] is introduced for CUPS on Control Plane and User Plane in MonSub CLI. The new option is like the D option in the ICUPS. The slow-path and fast-path PCAP generates only when this option is ON.

- There are a maximum of four subscriber tracing sessions per NPUMGR instance. The NPUMGR (per User Plane instance) enforces the maximum tracing session limit. Slow-path capture naming convention contains the MonSub tracing session ID on SMGR instance, whereas fast-path tracing session contains the PSN as session ID. If there are already four tracing sessions running at SESSMGR instance, then slow-path capture is by name "S4". It continues until the time NPUMGR rejects the tracing session due to max tracing limit reached.

Following are some of the important definitions related to this feature:

- **Chassis Traffic Volume**: The total volume of packet throughput on the chassis.

- **Monitored Traffic Volume**: Monitoring of the total throughput of all the subscribers through MonSub across all the MonSub sessions.

- PCAP success: The percentage of the MonSub traffic capture request and the successful capture in the PCAP files.

Packet Processing Throughput

Following are the scenarios that impacts the packet processing throughput:

- When VPP utilization is above 80%, MonSub may have an impact to packet processing throughput. The impact is in proportion to the monitored traffic volume.

- Specifically, when the monitored traffic volume approaches 10% of the chassis traffic volume, there may be impact on the VPP throughput causing subscriber packet loss.

- The impact to packet processing throughput is higher when using monitor priorities above 0 (zero).

⚠

Caution   you must be cautious during the packet processing. When VPP is running at 80% utilization and handling approximately 10Gbps chassis traffic volume, there's impact on the packet processing, if the set of MonSub sessions is collectively monitoring the subscribers, totaling more than 1Gbps of monitored traffic volume.

PCAP Success

The PCAP success depends on the following factors:

- The level of PCAP success depends on several factors, including monitored traffic volume, VPP utilization, MonSub monitor priority, and background disk I/O.

- In general, the PCAP success rates are greater for the following cases:

    - When the VPP utilization is low and/or MonSub monitor priority is above best-effort.

    - When the monitored traffic volume is less than 10% of the chassis traffic volume.

Example: When VPP is running at 80% utilization, handling approximately 10Gbps chassis traffic volume, monitored traffic volume up to 1Gbps is likely to yield high PCAP success percentages.

How It Works

The Monitor Subscriber feature is discussed in detail in the following sections:

UPF SessMgr Functionality

Following are the modification done in the UPF SessMgr to support this feature.

- Provide services to the CLI for enabling or disabling the MonSub tracing.

- Control NPUMgr to connect/start/stop/add/delete streams/TEP bearers and disconnect.

- The SessMgr maintains the PSN from the NPUMgr (as part of CONNECT API) and sub session ID, which is SessMgr (local to SMGR instance) specific. The SessMgr sends all requests with PSN and sub session ID to NPUMgr for a monitor subscriber tracing session.

- Based on the instructions from the CLI, configures panopticon (through NPUMgr) for changes, such as packet size and priority.

- Read the "hex dump module" configurations and store them locally. Pass the relevant parameters (such as filename) to Session Manager Co-Proc.

- Instantiate SessMgr Co-Proc and then instruct it to copy panopticon generated PCAP files to hard disk. Also handle the termination of SessMgr Co-Proc when MonSub session is over.

- Handle file copy message from SessMgr Co-Proc and inform panopticon about the copied bundle.

- If the file copy fails or there are problems with SessMgr Co-Proc instantiation, raises the SNMP alarms.

- Handle the buffer full indications from panopticon and copy the PCAP from the RAM disk to the configured destination directory.

- Capture the control/slowpath packets. Pass them to SessMgr Co-Proc to publish it as a separate PCAP.

- This feature supports a maximum of four monitor subscriber tracing sessions for a UPF instance. The NPUMgr enforces the tracing limit.

- The MonSub tracing session terminates in the absence of no space on hard disk or no hard disk.

- There are C-Proc (file copy and logging) per UPF SessMgr instance, when monitor subscriber tracing is initiated for that SessMgr instance.

- The MonSub session tear down takes time depending on the final poll timer and disconnect responses from Co-Proc/NPUMgr.

Multi PDN Multi Trace

For a multi-PDN call, when you start the MonSub with Multi-trace=OFF, then it traces the only one PDN as a part of that MonSub session. When new PDN is initiated, then existing PDN tracing stops and new PDN tracing starts. For this, first the new PDN tracing is started and then existing PDN tracing is stopped and hence new PSN and SessMgr sub-session ID is allocated.

For a multi-PDN call, when you start the MonSub with Multi-trace=ON, then it traces the new PDN as part of new fastpath tracing session (that is MonSub session). Hence, after tracing the four PDN, MonSub CLI shows max tracing session reached. Tracing of the each PDN takes place as a separate MonSub session.

MonSub Statistics

A new mechanism is added to publish the statistics regarding the quality of FASTPATH PCAP capture on MonSub CLI. The new mechanism publishes the statistics whenever it receives the buffer full MEH indication at SessMgr, throttled at every five seconds. The feature supports a maximum of four buffers for a FASTPATH PCAP corresponding to MonSub session. The feature does not publish the statistics by default, and needs to be enabled through debug CLI on UPF.

- **debug uplane monsub-stats disabled**

- **debug uplane monsub-stats enabled**

The stats contains the following informations:

```
Packet accepted: 14250000        Packet rejected: 62297
Congestion Short Term: 0                 Congestion Longer Term: 0
Throttled: 0                             PCAP File Transfer Rate: 9.91 mbps
```

The PCAP file transfer rate is the rate at which copy Co-Proc writes the PCAP from RAM-FS to HD-RAID.

X-Header

This feature supports the X-Header capture in slowpath PCAP. When UPF inserts the X-Header for Uplink packet, the UPF captures the packet at entry and exit interfaces. So, the exit packet sent to N6 contains the inserted X-header.

When UPF inserts the X-header for Downlink packet, the UPF captures the packet at entry and exit interfaces. So, the exit packet sent to N3 contains the inserted x-header.

Configuration Procedure for Monitor Subscriber

The protocol monitor can be used to display information for a specific subscriber session that is currently being processed. Depending on the number of protocols monitored, and the number of sessions in progress, a significant amount of data is generated. It is highly recommended that logging be enabled on your terminal client in order to capture all of the information that is generated.

Follow the instructions in this section to invoke and configure the protocol monitoring tool for a specific subscriber session.

Step 1   Invoke the monitor subscriber command from the Exec mode by entering the **monitor subscriber** CLI command.

```
[local]host_name# monitor subscriber { callid | imei | imsi | ipaddr | ipv6addr |
msid | msisdn | next-call | pcf | peer-fa | peer-lac | sgsn-address | type |
username }
```

An output listing all the currently available protocols, each with an assigned number, is displayed. Specify the method the monitor should use by entering the appropriate keyword.

Step 2   Specify the method the monitor should use by entering the appropriate keyword.

Select other options and/or enter the appropriate information for the selected keyword.

Step 3   Select other options and/or enter the appropriate information for the selected keyword.

If no session matching the specified criteria was being processed when the monitor was invoked, a screen of available monitoring options appears.

Step 4   Configure the amount of information that is displayed by the monitor. To enable or disable options, enter the letter or 2-digit number associated with that option (C, D, E, 11, 12, etc.). To increase or decrease the verbosity, use the plus (+) or minus (-) keys.

The current state, ON (enabled) or OFF (disabled), is shown to the right of each option.

Option **Y** for performing multi-call traces is only supported for use with the GGSN.

```
WARNING!!! You have selected options that can DISRUPT USER SERVICE
Existing CALLS MAY BE DROPPED and/or new CALLS MAY FAIL!!!
(Under heavy call load, some debugging output may not be displayed)
Proceed? - Select (Y)es or (N)o
```

Step 5   Repeat step 6 as needed to enable or disable multiple protocols.
Step 6   Press the **Enter** key to refresh the screen and begin monitoring.

The monitor remains active until disabled. To quit the protocol monitor and return to the prompt, press **q**.

Monsub CLI Options

Monitor Subscriber CLI – New Options

The following options with their default value are added to existing **monitor subscriber** command.

UPF Monitor Subscriber CLI

Following are the options:

- **W** - UP PCAP Trace (ON): This parameter is used to create PCAP trace for slowpath and fastpath and is only applicable for CUPS.

- **U** - Mon Display (ON): The non-protocol events (such as statistics and charging information from ECS and so on) are also captured in slowpath PCAP files and are displayed on U-PLANE monitor console.

- **V** - PCAP Hexdump (ON): This flag must be set to ON to capture the protocol packets in a text file in hexdump format on U-PLANE.

Note Currently, UP PCAP Trace flag must be set to ON to capture fastpath and slowpath PCAP files on CUPS.

Monitor Subscriber CLI – New Options

The following options with their default value are added to existing **monitor subscriber** command.

- **F - Packet Capture (Full Pkt)**: Captures all packets from fastpath.

  Using this option, operators can choose between full and partial packet captures. By entering **F**, the packet capture type can be changed to either full or partial. With partial packet capture, users can enter packet sizes from 1 to 16384 bytes. For example, if input is given as 20, only the first 20 bytes of fastpath packets will be captured and the remaining packets will be dropped.

  Note When opening the PCAP file, the summary view will display full length of the packet, but the detailed view will show only the truncated packet.

- **/ - Priority (0)**: The value is in the range from "0 – Best Effort" to "7 – Guaranteed"

  - 0 - Best Effort

  - 1 - Low

  - 2 - Med-Low

  - 3 - Medium

  - 4 - Med-High

  - 5 - High

  - 6 - Critical

  - 7 - Guaranteed

> ⚠️ **Caution**   It is strongly recommended to not change the default value. It can adversely affect the system performance.

• **N - MEH Header (OFF)** : The MEH header is stripped from the IP packet if this option is configured

Show Monitor Subscriber Sessions

Following is the new CLI to show the ongoing MonSub session.

You can trigger the CLI **show monitor subscriber fastpath session all** from UPF.

- • **SessId:** This is the local session id for MonSub session on UP Sessmgr.

- • **CallID:** Call id on Userplane.

- • **PSN:** This is panopticon sequence no. There is a maximum of four MonSub fastpath tracing sessions on one UP with PSN ranging from 0-3.

- • **Start time:** Time at which MonSub tracing session starts.

- • **Interface Type:** This is to identify the call type for which MonSub fastpath tracing session was started, whether it is Sxa, Sxb or Sxab.

Disconnect Monitor Subscriber Sessions

Following is the new CLI to disconnect the ongoing MonSub session. You can trigger the CLI from UPF.

monitor subscriber fastpath disconnect sessmgr-instance <UP SMGR Instance ID> session-id <Local Monitor subscriber Session ID at SMGR instance level>

If the MonSub session disconnect is successful, the following message dispalys on console.

```
Session Disconnected Successfully
```

If the MonSub session disconnect fails, the following message dispalys on console.

```
Monitor Subscriber session does not exist
```

> ✎ **Note**   Only security administrator can execute the monitor disconnect CLI.

Context, CDRMOD and Hexdump Interaction for Monitor Subscriber

Hexdump module must be configured to provide operators the provision to configure Files names and Poll timers. The Hexdump module is one of the modules such as—EDR, UDR and so on, that are part of the CDRMOD functionality. It is recommended to configure the hexdump in a non-local context such as the ECS context. Hexdump modules are not supported in local context.

For more information on Hexdump module and its configuration, refer to the *Packet Capture (PCAP) Trace* chapter in the *ASR5500 System Administration Guide*

PCAP File Name Convention

The naming conventions for PCAP files are discussed in the following sections

Note Only **monitor-subscriber-file-name** and **rotation** options are used in naming PCAP files.

Slowpath File Name Convention

The slow path file names appear in the following format:

```
curr_slowpath_{SMGR Mon Sub Session
Id}_{monsub_file_name_option_val}_{Timestamp}_{RotationCount}.pcap
```

or

```
slowpath_{SMGR Mon Sub Session
Id}_{monsub_file_name_option_val}_{Timestamp}_{RotationCount}.pcap
```

File with 'curr_' prefix is the file, that is currently being written to, that is still not closed. When files are to be rotated (depending on the file rotation parameters), file without the 'curr_' prefix are copied to hard disk.

The SMGR MonSub Session Id – This is the session Id for MonSub session created on Uplane SMGR instance ID, which created this PCAP. This Id is local to SMGR instance, so there could be two SLOWPATH pcap captured with same ID.

When files are to be copied to hard disk, The monsub_file_name_option_val is replaced by:

- IMSI value if **monitor-subscriber-file-name** is set to "imsi".

- Call ID value if **monitor-subscriber-file-name** is set to "call-id"

- Username value if **monitor-subscriber-file-name** is set to 'username'

Timestamp is in the following format "MMDDYYYYHHMMSS", where:

- MM - Month, DD - Date and YYYY - Year.

- HH -Hour, MM - Minutes and SS - Seconds.

RotationCount is a 9-digit value that is incremented every time an old file is rotated, and a new file is generated.

00000000 for the first file, 00000001 for the second file and so on.

Rotation of slowpath files is determined by following option in **hexdump-module file** configuration:

rotation { num-records *number* | **time** *seconds* | **volume** *bytes* **}**

- **num-records:** num-records specifies the number of packets after which a new file is generated and 'RotationCount' in the filename is incremented. The range of number is between 100 to 10240, and the default value is 1024.

- **time:** time specifies the time to wait in seconds before a new file is generated and 'RotationCount' in the filename is incremented. seconds must be an integer from 30 through 86400. The default value is 3600.

- **volume:** volume specifies the number of bytes after which a new file is generated and 'RotationCount' in the filename is incremented. bytes must be an integer from 51200 through 62914560. The default value is 102400.

> **Note**
>
> The **tarriff-time** parameter under rotation is ignored as it is not suitable for PCAP file capture.

The following are examples of the file naming conventions for slowpath PCAP files.

- For the 'imsi' option where IMSI is '112233445566778', slowpath files are named as:

  `slowpath_S0_112233445566778_07152019050907_000000000.pcap`

- For 'call_id' option where Call Id is '01317b22', slowpath files are named as:

  `slowpath_S0_01317b22_07152019050907_000000000.pcap`

> **Note**
>
> The parameter **tarrif-time** is not applicable for PCAP file capture.

Fastpath File Name Convention

The fastpath file names appear in the following format:

`vpp_{S}_{B}_{encap}_{monsub_file_name_option}_{Timestamp}_{FileCount}.pcap`

- S is replaced by either 'S1', 'S2','S3', or 'S4'.

- B is replaced by either 'B0','B1','B2', or 'B3' depending on the bundle generated by Panopticon.

- monsub_file_name_option is replaced by:

  - IMSI value if **monitor-subscriber-file-name** is set to "imsi".

  - Call ID value if **monitor-subscriber-file-name** is set to "call-id"

  - Username value if **monitor-subscriber-file-name** is set to 'username'

Timestamp is in the following format "MMDDYYYYHHMMSS", where:

- MM - Month, DD - Date and YYYY - Year.

- HH -Hour, MM - Minutes and SS - Seconds.

RotationCount is a 9-digit value that is incremented every time an old file is rotated, and a new file is generated.

00000000 for the first file, 00000001 for the second file and so on.

Fast path "FileCount" is not the same as the slowpath "RotationCount" parameters and hence 'hexdump-module file rotation' parameters are ignored while naming fastpath files.

In Phases 1 of the feature, fastpath generated file names are like 'vpp_S1_B0_ip.pcap' or 'vpp_S1_B1_ip.pcap', they are renamed to following when being copied over to non-volatile storage:

- `vpp_S1_B0_ip_01317b22_07152019050907_000000000.pcap`

- `vpp_S1_B1_ip_01317b22_07152019050908_000000001.pcap`

- `vpp_S1_B0_ip_01317b22_07152019050908_000000002.pcap`

In MonSub phase 3, a PCAP "bundle" is replaced with a single PCAP file that uses Ethernet encapsulation.

In Phase 3, each fastpath session file is captured in the Ethernet PCAP file that is 'vpp_S0_B0_eth.pcap' and they are renamed to following when being copied to a non-volatile storage:

```
vpp_S0_B0_eth_01317b22_07152019050907_000000000.pcap
```

For 'callid' option where Call Id is '12345678ef':

- `slowpath_S0_12345678ef _07152019050907_000000000.pcap`

- `vpp_S1_B0_eth_12345678ef _07152019050907_000000000.pcap`

For 'username' option where username is '9890098900':

- `slowpath_S0_07152019050907_000000000_9890098900.pcap`

- `vpp_S1_B0_eth_07152019050907_000000000_9890098900.pcap`

PCAP File Location

Fastpath PCAP files are written to the `/records/pcap` directory in same card and CPU complex where the SMGR owns the subscriber session resides.

`/records` directory is mapped to the "tmpfs" filesystem that is mapped to RAM. In this state, the files are suffixed with a ".pending" extension. For example:

```
-rw-rw-r-- 1 root root 268599296 Sep 23 14:04 vpp_S1_B0_eth.pending
```

Note The files size at this stage is not the actual file size when it is written to a persistent storage.

Once the fastpath tracing mechanism has written the files, they are converted to '.pcap' files and renamed as given below. Additionally, there is a file that ends with a ".done" extension:

```
-rw-rw-r--   1 root    root      8689188 Oct 16 22:06 vpp_S0_B0_eth.pcap
```

After the PCAP files are written by fastpath tracing mechanism, the Co-Proc functionality instantiates and copies the files to a hard disk or a persistent storage.

The above file location process for Fastpath is also applicable to Slowpath.

The target file location in all cases is: `/hd_raid/records/hexdump`, except for the case in the hexdump module configuration where **use-harddisk** is enabled and the **directory** option under the **hexdump file** is to a custom value. For example, if the **directory** option is set to a value "abc" then the target location for the PCAP file will be: `/hd_raid/records/hexdump/abc/`.

In this feature implementation, a predefined location is set for PCAP files

- To make sure that `/records/pcap` directory is not populated when issues are encountered with the use of **use-harddisk** and **hexdump module** configurations.

- For regular cleanup from `/hd_raid/records/hexdump` directory.

File Transfer to External Location

Once the files have been copied to the hard disk, they can be copied over to external server using the command: **transfer-mode** option under the **hexdump** command in the **hexdump-module** configuration.

Apart from **transfer-mode**, other relevant options under **hexdump** can be used for external file transfer. Operators can use these commands to avoid excessive storage during fastpath processing.

Limitations

Following are the Limitations:

- SR/ICSR is not supported for this feature.

- Restarting trace immediately after quitting may result in fastpath files in `/records/pcap` directory to be overwritten. It is recommended to restart the session after a brief moment (a few seconds).

- When MonSub trace is stopped, the tear down process can take a few seconds, so it is recommended to wait for few seconds. A maximum of (5sec, hexdump poll timer value in sec) before toggling the MonSub trace to start, else operators may observe MAX TRACING SESSIONS REACHED momentarily.

- Show monitor subscriber fastpath sessions CLI does not display the MonSub sessions that are being stopped. Hence there is a transient period where new MonSub sessions can be rejected due to max sessions reached, whereas show CLI shows less sessions. It is recommended that operators wait for some time before starting a new MonSub trace session.

- Changing fastpath configuration options is only possible when **UP Pcap Trace** is set to OFF.

- When MT=ON in the Multi-PDN, then once MT=OFF, new PDN tracing is not started due to MAX TRACING REACHED, and then all other tracing is STOPPED. This is because the first new PDN tracing is started and then all previous PDN's were STOPPED for MT=OFF case.

- It is recommended to not to launch the same UE MonSub sessions from different CLIs.

- In slowpath PCAP, the egress DL packets does not show the GTPU-U header because the functionality to add GTP-U is with fastpath. So, ingress and egress DL packets shows up the duplicates, unless there is some packet modification like HTTP X-headers applied over the ingress packets.

- Toggling C and D options does not impact the PCAP captures in UPF.

- For Multi-PDN, the fastpath filenames does not use the Call Id, because, by definition the multi-PDN case has more than one call id and hence a higher-level configuration such as IMSI is more suitable for naming the files.

- Only the named options explicitly mentioned in this document are supported from *hexdump-module file* configuration.

- Number of streams that can be traced in fastpath is limited to 5000. Stream is defined as a TCP or UDP flow which is made up of {source IP address, destination IP address, source port, and destination port, transport protocol such as TCP or UDP).

- Fastpath packets cannot be streamed to an external server. They are stored on the hard-disk and transferred (either manually) or by using **transfer-mode** options.

- The UP PCAP trace must be set to ON to capture fastpath and slowpath PCAP files.

- MonSub CLI option '<SPACE> Pause' is only to pause console events. There is no impact on other tracing events (slowpath PCAP, fastpath PCAP and protocol packets tracing in a text file in hexdump format) with this option.

- The UP trace PCAP file does not contain the initial PFCP Sx Request/Response, due to race condition.

- The ICMP Packets and a first packets of TCP and UDP streams flow through both slowpath and fastpath. Default values of GTPU (option 26) and User L3 (Option 19) are set to OFF. As a result, these packets are not captured in slowpath captures. If Option 26 is set to ON then these packets are captured in slowpath PCAP captures. As mentioned in previous point, option 19 has no effect on slowpath PCAP capture.

- Data Events flag must be set to ON to capture fastpath and slowpath PCAP files.

- The Mon sub tracing is not supported for option **Next-SAEGW Call** on UP.

- The Mon sub tracing is not supported for option **Next call by APN** for Pure-S call type.

- On ASR-5500 setup with the default value of poll-timer, all the packets may not be captured due to a known issue. To avoid large number of packets to be rejected, it is recommended to change the poll-timer value to the lowest possible (10ms).

- If context replacement occurs (if the same subscriber reattaches without a detach) then the slowpath captures for the new call continues to be in the old slowpath files.

Configuring the Hexdump Module for MonSub in UPF

Configuring MonSub Poll Timer

Use this configuration to set the frequency of PCAP file capture check.

```
configure
   context context_name
      hexdump-module
         hexdump monitor-subscriber-poll-timeout poll_timer_value
         end
```

NOTES:

- **hexdump monitor-subscriber-poll-timeout** : This option specifies how frequently the check for newly captured PCAP files in the volatile storage must be done before they are copied to persistent storage.

- *poll_timer_value*: Specifies the poll timer value in milliseconds. It must be an integer in the range of 10 ms to 60 seconds. Default: 30 seconds.

> **Note**   It is strongly recommended to not configure the timer with a value less than 5 seconds.

- This option is only applicable when MonSub is enabled for the products that have fastpath functionality - PGW, SAEGW on ASR-5500 and VPC-SI.

Configuring MonSub File Name

Use the following configuration to specify the file name of the PCAP file which contains IMSI, Call ID or Username.

```
configure
  context context_name
    hexdump-module
      file rotation { num-records number | tariff-time minute minutes hour
 hours | time seconds| volume bytes | monitor-subscriber-file-name { imsi |
username | call-id }
      end
```

NOTES:

- **monitor-subscriber-file-name { imsi | username | call-id }**: This option specifies if the name of the captured PCAP files will contain IMSI, Call Id or Username. This option is only applicable on products that have fastpath functionality (PGW, SAEGW on ASR 5500 and VPC-SI) AND only when Monitor Subscriber functionality is enabled. Default: IMSI.

- **rotation { num-records** *number* **| tariff-time minute** *minutes* **hour** *hours* **| time** *seconds* **| volume** *bytes* **}**: Specifies when to close a hexdump file and create a new one.

    - **num-records** *number* : Specifies the maximum number of records that should be added to a hexdump file. When the number of records in the file reaches this value, the file is complete.

      number must be an integer from 100 through 10240. Default: 1024

    - **tariff-time minute** *minutes* **hour** *hours* : Specifies to close the current hexdump file and create a new one based on the tariff time (in minutes and hours).

      minutes must be an integer from 0 through 59.

      hours must be an integer from 0 through 23.

    - **time** *seconds* : Specifies the period of time to wait (in seconds) before closing the current hexdump file and creating a new one.

      seconds must be an integer from 30 through 86400. Default: 3600

      ☞

      Important   It is recommended to set the rotation time to 30 seconds.

    - **volume** *bytes* : Specifies the maximum size of the hexdump file (in bytes) before closing it and creating a new one.

      bytes must be an integer from 51200 through 62914560. Note that a higher setting may improve the compression ratio when the compression keyword is set to gzip. Default: 102400

Monitoring and Troubleshooting

This section provides information regarding monitoring and troubleshooting the Monitor Subscriber feature.

SNMP Traps

The following SNMP trap(s) are added in support of the Monitor Subscriber feature:

- MonSubProcessInitFailure: This trap is triggered when MonSub handler process has failed for a particular process and service.

CHAPTER **22**

Multiple cnSGW Support

Feature Summary and Revision History

Summary Data

Table 42: Summary Data

Applicable Product (s) or Functional Area	5G-UPF
Applicable Platforms	VPC-SI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

Revision Details	Release
First introduced.	2021.02.0

Feature Description

The Multiple SMF/cnSGW Support feature enables a single UPF to establish multiple N4/Sx interfaces with cnSGWs and their paired SMFs. Integration of multiple SMF and cnSGW combinations with a single UPF results in optimal usage of resources.

Architecture

The following illustration depicts the architecture of multiple cnSGWs/SMFs.

Relationship to Other Features

The Multiple cnSGW Support feature is related to *Multiple N4/Sx Interface* feature.

How it Works

The functionality of Multiple cnSGW feature involves:

- Single UPF has multiple N4/Sx interface associations with SMF/cnSGWs.

- There's no slicing of configuration in UPF per individual SMF.

- Cumulatively, a maximum of four peers—combination of cnSGW/SMF or individual cnSGW/SMF as per the need—are connected to a single UPF.

- Individual N4/Sx association release purges sessions of the impacted peer.

- UPF redundancy works seamlessly.

- All cnSGWs paired with a UPF is associated with a single user plane service.

Configuring Multiple SMF/cnSGWs

This section provides information about CLI commands that are available in support of this feature.

Configuring Multiple SMF/cnSGWs on UPF

Use the following CLI commands to configure multiple SMF/cnSGWs on UPF by adding multiple peer node under Control Plane Group Configuration mode.

```
configure
  user-plane-service service_name
    associate control-plane-group group_name
  control-plane-group group_name
    peer-node-id ipv4-address ipv4_address interface n4
    peer-node-id ipv4-address ipv4_address interface n4
    peer-node-id ipv4-address ipv4_address
    peer-node-id ipv4-address ipv4_address
    end
```

Monitoring and Troubleshooting

This section provides information about monitoring and troubleshooting the Multiple cnSGW feature.

Show Commands and/or Outputs

This section describes the show commands that are available in support of this feature.

show subscribers user-plane-only full all

The output of this CLI command is enhanced to display the corresponding Control Plane address.

show sx peers

The output of this CLI command is enhanced to display the peer ID with corresponding number of sessions.

show sx peers

Multiple N4/Sx Interface

Feature Summary and Revision History

Summary Data

Table 43: Summary Data

Applicable Product (s) or Functional Area	5G-UPF
Applicable Platforms	VPC-SI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

Revision Details	Release
First introduced.	2021.01.0

Feature Description

The Multiple N4 Interface feature enables a single UPF to establish multiple N4 interfaces with as many SMFs. Integration of multiple SMFs with a single UPF results in optimal usage of resources.

Architecture

The following illustration depicts the architecture of Multiple N4 Interface.

How it Works

The functionality of Multiple N4 Interface feature involves:

- Single UPF has multiple N4/Sx interface associations with each SMF.

- There is no slicing of configuration in UPF per individual SMF.

- The ECS/ACS configuration at the UPF is a union of all the individual SMF-specific configurations. For example:

  SMF1 has rulebase *RB1* and no *RB2*.

  SMF2 has rulebase *RB2* and no *RB1*.

  The UPF has both rulebase, *RB1* and *RB2* to cater the sessions from *RB1* and *RB2*.

- A maximum of four SMF peers are connected to a single UPF.

- Overlapping IP pools from multiple SMFs are segregated based on the VRF ID.

- Individual N4 association release purges sessions of the impacted SMF peer.

- UPF redundancy works seamlessly.

- In rare instance of any conflict amongst different SMF configurations, it will not be resolved at the UPF and will be installed in the sequence in which such CLIs were configured.

Configuring Multiple N4 Interface

This section provides information about CLI commands that are available in support of this feature.

Configuring Multiple SMF on UPF

Use the following CLI commands to configure multiple SMF on UPF by adding multiple peer node under Control Plane Group Configuration mode.

```
configure
  user-plane-service service_name
    associate control-plane-group group_name
  control-plane-group group_name
    peer-node-id ipv4-address ipv4_address interface n4
    peer-node-id ipv4-address ipv4_address interface n4
    . . .
    . . .
    . . .
    end
```

Monitoring and Troubleshooting

This section provides information about monitoring and troubleshooting the Multiple N4 Interface feature.

Show Commands and/or Outputs

This section describes the show commands that are available in support of this feature.

show ip chunks

The output of this CLI command is enhanced to display the IP pools pushed to the UPF from multiple SMFs in Gi context.

show ipv6 chunks

The output of this CLI command is enhanced to display the IPv6 pools pushed to the UPF from multiple SMFs in Gi context.

show subscribers user-plane-only full all

The output of this CLI command is enhanced to display the corresponding Control Plane address.

show sx peers

The output of this CLI command is enhanced to display the peer ID with corresponding number of sessions.

show user-plane-service statistics peer-address <address>

The output of this CLI command is enhanced to display per peer statistics in SMF.

N:M Redundancy and Redundancy Configuration Manager

Feature Summary and Revision History

Summary Data

Table 44: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G RCM Configuration and Administration Guide*

Revision History

Table 45: Revision History

Revision Details	Release
Support has been added to defer SSH IP installation.	2021.02.2
Support has been added for configuring the RCM through Network Services Orchestrator (NSO).	2021.02.0
First introduced.	2021.01.0

Feature Description

The Redundancy Configuration Manager (RCM) is a Cisco proprietary node/network function (NF) that provides redundancy of StarOS-based UP/UPFs. The RCM provides N:M redundancy of UP/UPFs wherein "N" is the number of Active UPs/UPFs and is less than 10, and "M" is the number of Standby UP/UPF in the redundancy group.

For details, refer the Redundancy Configuration Manager - Configuration and Administration Guide.

N3 Transfer of PDU Session Information

- Feature Summary and Revision History, on page 223
- Feature Description, on page 223

Feature Summary and Revision History

Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Revision Details	Release
First introduced.	2020.02.0

Feature Description

The N3 transfer of PDU session information involves the inclusion of QoS Field Identifier (QFI) IE in the GTP-U extension header while performing GTP-U encapsulation toward gNodeB on the N3 interface, and removal of the GTP-U extension header while performing GTP-U decapsulation when packets are received from the gNodeB.

The QFI IE detects traffic pertaining to specific QoS sessions. It is used to send control information between the gNodeB and the UPF.

How it Works

This section describes the transfer of PDU session Information procedures between the gNodeB and UPF for Uplink and Downlink packets.

Transfer of PDU Session Information for Downlink Data Packets

The Transfer of PDU Session Information for downlink data packets involves transfer of control information elements related to the PDU Session from UPF/NG-RAN to NG-RAN.

A PDU session user plane instance that makes use of this transfer procedure is associated to a single PDU Session. The procedure is invoked whenever packets for that particular PDU Session need to be transferred across the related interface instance.

The DL PDU SESSION INFORMATION frame includes a QoS Flow Identifier (QFI) field associated with the transferred packet. The NG-RAN uses the received QFI to determine the QoS flow and QoS profile which are associated with the received packet.

The following frame shows the respective DL PDU SESSION INFORMATION.

NOTE: In current implementation, the Reflective QoS Indicator (RQI) and Paging Policy Presence (PPP) in DL PDU SESSION INFORMATION frame is not supported.

Transfer of PDU Session Information for Uplink Data Packets

The Transfer of PDU Session Information for uplink data packets involves transfer of control information elements related to the PDU Session from NG-RAN to UPF.

An UL PDU Session user plane instance that makes use of the transfer procedure is associated to a single PDU Session. This procedure is invoked whenever packets for that particular PDU Session need to be transferred across the related interface instance.

The UL PDU SESSION INFORMATION frame includes a QoS Flow Identifier (QFI) field associated with the transferred packet.

The following frame shows the respective UL PDU SESSION INFORMATION.

PDU Session Information Frame IEs

The following table describes the Information Elements present in the PDU Session Information frame.

Information Element	Description
PDU Type	The PDU Type indicates the structure of the PDU session UP frame. The field takes the value of the PDU Type it identifies: "0" for PDU Type 0. The PDU type is in bit 4 to bit 7 in the first octet of the frame. **Value range**: {0= DL PDU SESSION INFORMATION, 1=UL PDU SESSION INFORMATION, 2-15=reserved for future PDU type extensions} **Field length**: 4 bits
Spare	The spare field is set to "0" by the sender and should not be interpreted by the receiver. This field is reserved for later versions. **Value Range**: (0–2n-1) **Field Length**: n bits
QoS Flow Identifier	When this IE is present, this parameter indicates the QoS Flow Identifier of the QoS flow to which the transferred packet belongs. **Value range**: {0 to 2^6-1} **Field length**: 6 bits
Padding	The padding is included at the end of the frame to ensure that the PDU Session user plane protocol PDU length (including padding and the future extension) is (n*4– 2) octets, where n is a positive integer. **Field Length**: 0–3 octets.

Standards Compliance

The feature complies with the following standard: 3GPP TS 38.415 V15.2.0 (NG-RAN; PDU Session User Plane Protocol).

Limitations

The following are the known limitations to this feature in this release:

- Reflective QoS Indicator (RQI) is not supported in this release.

C H A P T E R 26

N4 Interface Compliance with 3GPP Specification

This chapter covers the following topics:

Feature Summary and Revision History

Summary Data

Table 46: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI
Feature Default Setting	Enabled - Always-on
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Table 47: Revision History

Revision Details	Release
In this release, PFCP library is upgraded to support the latest version of Outer Header IE.	2020.02.5
First introduced.	2020.02.0

Feature Description

In compliance with 3GPP TS 29.244, the User Plane Function (UPF) supports the following IEs:

- Averaging Window

- Paging Policy Indicator (PPI)

- Outer Header Creation

- Outer Header Removal

Averaging Window

Averaging window IE contains the duration over which the GBR and MBR is calculated. It is sent from SMF to UPF with Create QER or Update QER parent IE, if the default pre-configured value under UPF needs to be overridden.

The following format is used for encoding and decoding of the IE:

	Bits							
Octets	8	7	6	5	4	3	2	1
1 to 2	Type = 157 (decimal)							
3 to 4	Length = n							
5 to 8	Averaging Window							
9 to (n+4)	These octet(s) is/are present only if explicitly specified							

NOTE: The value should be in milliseconds.

Paging Policy Indicator

The SMF sends PPI value in Create QER or Update QER, if UPF requires to set Paging Policy Indicator in outgoing packets.

In the case of Network Triggered Service Request and UPF buffering downlink data packet, the UPF includes the DSCP in TOS (IPv4) / TC (IPv6) value from the IP header of the downlink data packet. It also sends an indication of the corresponding QoS Flow in the data notification message to the SMF. When PPD applies, the SMF determines the Paging Policy Indicator (PPI) based on the DSCP received from the UPF.

In the case of Network Triggered Service Request and SMF buffering downlink data packet, when PPD applies, the SMF determines the PPI based on the DSCP in TOS (IPv4) / TC (IPv6) value from the IP header of the received downlink data packet and identifies the corresponding QoS Flow from the QFI of the received downlink data packet.

The following format is used for encoding and decoding of the IE:

	Bits							
Octets	8	7	6	5	4	3	2	1

1 to 2	Type = 158 (decimal)	
3 to 4	Length = n	
5	Spare	PPI value
6 to (n+4)	These octet(s) is/are present only if explicitly specified	

NOTE: The PPI should be encoded as a 3-bit value between 0 and 7.

Outer Header Creation

Per 3GPP TS 29.244 v16.4.0, the Outer Header Creation Description field, when present, is encoded as specified in following table. It takes the form of a bitmask where each bit indicates the outer header to be created in the outgoing packet. Spare bits are ignored by the receiver.

Octet / Bit	Outer Header Created in the Outgoing Packet
5/1	GTP-U/UDP/IPv4
5/2	GTP-U/UDP/IPv6
5/3	UDP/IPv4
5/4	UDP/IPv6
5/5	IPv4
5/6	IPv6
5/7	C-TAG
5/8	S-TAG
6/1	N19 Indication
6/2	N6 Indication
6/3	TCP/IPv4
6/4	TCP/IPv6

NOTE:

- Currently, the UP/UPF doesn't support the following values of Outer Header Creation Description:

  - IPv4

  - IPv6

  - C-TAG

  - S-TAG

  - N19 Indication

  - N6 Indication

- Third and fourth bits of sixth Octet (that is, 6/3 and 6/4) are spare bits (that is, not part of 3GPP TS) used for LI over TCP.

👉

Important   If SMF/CP uses older version for Outer Header Creation, then undefined behavior (including crashes) can be seen.

C H A P T E R 27

N4 Interface Configuration

This chapter covers the following topics:

Feature Summary and Revision History

Summary Data

Table 48: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI
Feature Default Setting	Disabled - Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

Table 49: Revision History

Revision Details	Release
New IEs are supported in UPF in compliance with 3GPP TS 29.244.	2021.01.0
First introduced.	2020.02.0

Feature Description

This chapter provides the configuration information to identify a peer node to be an N4 interface, and the configuration to modify N4 parameters in an Sx-Service.

Configuring N4 Interface

This section describes the following configurations:

- Identifying N4 Interface

- Adding N4-type and Modification of N4 Parameters in Sx Service

Identifying an N4 Interface

Use the following configuration to identify if a peer node is an N4 interface type.

```
configure
   control-plane-group group_name
      peer-node-id [ ipv4-address ipv4_address | ipv6-address ipv6_address ]
interface n4
         end
```

NOTES:

- To enable the **n4 interface** CLI command, you need the **require upf** CLI command on the UPF, which depends on the UPF license.

- **[ipv4-address** *ipv4_address* | **ipv6-address** *ipv6_address* **]** :

  ipv4-address *ipv4_address*: Specifies the IPv4 address of the peer node.

  ipv6-address *ipv6_address*: Specifies the IPv6 address of the peer node.

- **interface n4**: Identifies the N4 interface.

Modification of N4-type Parameters in an Sx Service

Use the following configuration to modify N4-type parameters in an Sx Service.

```
configure
   context context_name
      sx-service service_name
         n4 [ max-retransmissions max_retransmission_value |
retransmission-timeout-ms timeout_value ]
         end
```

NOTES:

- **n4**: Allows modifications to N4 parameters.

- [**max-retransmissions***max_retransmission_value* | **retransmission-timeout-ms** *timeout_value*]:

    max-retransmissions *max_retransmission_value* Configures maximum retries for Sx control packets. *max_retransmission_value* must be an integer in the range of 0 to 15. The default value is 4.

    retransmission-timeout-ms: Configures the control packet retransmission timeout in Sx in milliseconds. *timeout_value* must be an integer in the range of 1000 to 20000 milliseconds. The timeout value must be configured in steps of 100; for example: 1000, 1100, 1200, and so on. The default value is 5000 milliseconds.

Statistics

This section provides information on show commands and their output available in support of this feature.

show control-plane-group

The output of this command displays the following fields for this feature:

- Interface Type – This field indicates if the peer interface is N4. It is not displayed for non-N4 interfaces.

show sx-service all

The output of this command displays the following fields for this feature:

- N4

    - N4 Retransmission Timeout

    - N4 Maximum Request Retransmission

show subscribers user-plane-only all

The output of this command displays the following fields for this feature:

- Interface

    - N4

show user-plane-service statistics all

The output of this command displays the following fields for this feature:

- N4 interface-type PDNs

    - Active

    - Setup

    - Released

show subscribers user-plane-only seid number pdr all

The output of this command displays the following fields for this feature:

• Associated-QFIs

show subscribers user-plane-only callid number pdr full all

The output of this command displays the following fields for this feature:

• QoS Flow Identifier

N4 Session Management, Node Level, and Reporting Procedures

This chapter covers the following topics:

Feature Summary and Revision History

Summary Data

Table 50: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Revision Details	Release
show user-plane-service statistics all command has been enhanced to display PFCP IEs received in Sx Establishment or Modification messages over N4 interface.	2021.02.2
PDN Update procedure is introduced in this release.	2021.02.0
First introduced.	2020.02.0

Feature Description

N4 Session Management, Node Level, and Reporting Procedures

N4 Node-level Procedures

The N4 Node-level procedures in User Plane Function (UPF) involves the following processes:

- N4 Association Set up Procedure – The procedure used for setting up an N4 association between the Session Management Function (SMF) and UPF.

- N4 Association Update Procedure – The procedure used for modifying an existing N4 association between the SMF and UPF.

- N4 Association Release Procedure – The procedure used for terminating the N4 association between the SMF and UPF.

- N4 Heartbeat Procedure – The procedure used for sending and receiving the Heartbeat request and response.

- N4 Reporting Procedure – The procedure used for reporting echo request and response for the GTP-u path failure.

N4 Session Management

N4 session management procedures are used to control the functionality of the UPF. SMF can create, update, and remove the N4 session context in the UPF, which is described in 3GPP TS 23.501, clause 5.8.2.

The following procedures are performed in N4 Session Management:

- N4 Session Establishment

- N4 Session Modification

- N4 Session Deletion

NOTE: The SMF initiates all the above procedures.

N4 Session/Node-level Reporting Procedures

Whenever the data path between UPF and gNB is down, it is detected and reported to the SMF for corrective actions. The mechanism to detect and report it to SMF is clearly defined in 3GPP specifications. The reporting happens per GTP-u Tunnel level or per GTP-u endpoint level.

Relationships

The following features support the N4 session management, node level, and reporting procedures.

End Marker Support

The UPF sends the End Marker packets to support the reordering function in the target Radio Access Network (RAN). The UPF constructs the End Marker packets that are required for the reordering function.

Constructing the End Marker Packets through UPF

At the time of the handover procedure, the PDU session for the UE – which comprises of an UPF node – acts as a PDU session anchor and an intermediate UPF terminating N3 reference point. The SMF sends an N4 Session Modification Request message with the new AN Tunnel Info of NG RAN to specify the UPF to switch to the N3 paths. In addition, the SMF also specifies the UPF to send the End Marker packets on the old N3 user plane path.

After the UPF receives the indication, the End Markers are constructed and sent to each N3 GTP-U tunnel toward the source NG RAN, after sending the last PDU on the old path.

UEs IPv4, IPv6, and IPv4v6 Support

The UPF supports UE's IPv4, IPv6, and IPv4v6 sessions.

The N4 Session Establishment and Modification procedure for IPv6 sessions is the same as for IPv4 sessions. After the session is established, the SMF sends Router Advertisement (RA) message to UE announcing the IPv6 prefix to be used for traffic. Optionally, to get the IPv6 parameter from SMF faster, the UE can also initiate IPv6 Router Solicitation (RS).

The N4 Session Establishment and Modification procedure for IPv4v6 Session are similar to the IPv4 or IPv6 sessions except for the allocation of two UE IP addresses - one for IPv4 and the other for IPv6. The SMF sends the Router Advertisement message to the UE announcing the IPv6 prefix used for traffic after the session is established. Optionally, the UE can also initiate the IPv6 Router Solicitation to receive the IPv6 parameter from the SMF quickly.

How it Works

This section describes the N4 node-level, session management, and reporting procedures and associated call flows.

N4 Node-level Procedure Call Flows

N4 Association Setup Procedure Call Flow

The N4 Association Setup procedure creates the N4 association between the SMF and the UPF, which enables the SMF to use the UPF resources to establish N4 sessions. The N4 association setup procedure involves the following steps:

1. The UPF initiates the procedure by sending N4 Association Setup Request to the SMF.

2. The SMF sends an N4 Association Setup Response after it receives the request from the UPF.

The following call flow describes the UPF-initiated N4 Association Setup procedure:

The UPF sends the following PFCP Association Setup Request:

- Node ID (UPF).

- Supported optional features in UPF. The UPF supports F-TEID allocation and release, sending of End Marker, and so on.

N4 Association Update Procedure Call Flow

The N4 Association Update procedure modifies an existing N4 association between the SMF and the UPF. It can be initiated either by the UPF or by the SMF to update the supported features or available UPF resources.

The following call flow depicts the SMF-initiated N4 Association Update procedure:

The following call flow depicts the UPF-initiated N4 Association Update procedure:
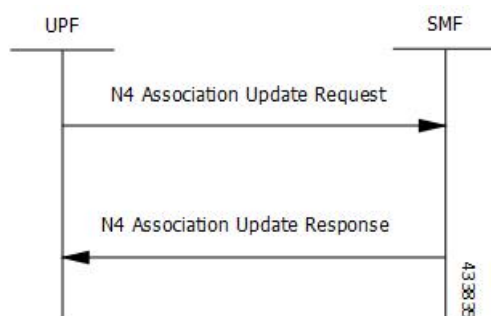
N4 Association Release Procedure Call Flow

The N4 Association Release procedure terminates the N4 association between the SMF and the UPF. It can be initiated either by the SMF or by the UPF. The UPF requests the SMF to perform the release of PFCP

association by sending a PFCP Association Update Request. The SMF then initiates a PFCP Association Release Request to release the PFCP association.

The following call flow depicts the SMF-initiated N4 Association Release procedure:

The following call flow depicts the UPF-initiated N4 Association Release procedure:

N4 Heartbeat Procedure

The PFCP Heartbeat procedure includes the following messages:

- Heartbeat Request
- Heartbeat Response

N4 Session Management Procedures Call Flows

The following section describes the N4 Session Management procedures.

N4 Session Establishment Call Flow

N4 Session Establishment is used to create the initial N4 session context for a PDU session at the UPF. SMF assigns a new N4 session ID and provides it to the UPF. The N4 session ID is stored by both entities and used to identify the N4 session context during their interaction. SMF also stores the relation between the N4 session ID and PDU session for a UE.
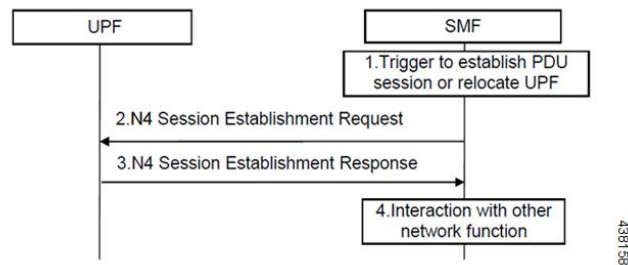
Step	Description
1	SMF receives the trigger to establish a new PDU session or change the UPF for an established PDU session.
2	SMF sends an N4 session establishment request message to the UPF that contains the structured control information which defines how the UPF needs to behave.
3	UPF responds with an N4 session establishment response message containing any information that the UPF has to provide to the SMF in response to the control information received.
4	SMF interacts with the network function which triggered this procedure. For example, AMF or PCF.

N4 Session Modification Call Flow

N4 Session Modification is used to update the N4 session context of an existing PDU session at the UPF, which is executed between SMF and UPF whenever PDU session-related parameters have to be modified.

Step	Description
1	SMF receives the trigger to modify the existing PDU session.
2	SMF sends an N4 session modification request message to the UPF which contains the update for the structured control information that defines how the UPF needs to behave.
3	UPF identifies the N4 session context to be modified by the N4 session ID and updates the parameters of this N4 session context according to the list of parameters that are sent by the SMF. UPF responds with an N4 session modification response message containing any information that the UPF has to provide to the SMF in response to the control information received.

Step	Description
4	SMF interacts with the network entity which triggered this procedure. For example, AMF or PCF.

N4 Session Delete Call Flow

N4 Session Delete is used to remove the N4 session context of an existing PDU session at the UPF.

Step	Description
1	SMF receives the trigger to remove the N4 session context for the PDU session.
2	SMF sends an N4 session delete request message to the UPF.
3	UPF identifies the N4 session context to be removed by the N4 Session ID and removes the whole session context. UPF responds with an N4 session delete response message containing any information that the UPF has to provide to the SMF.
4	SMF interacts with the network entity which triggered this procedure. For example, AMF or PCF.

N4 Session/Node Level Reporting Procedure Call Flows

Session Level Reporting Due to the GTP-u Error Indication Call Flow

When the UPF receives the GTP-u Error Indication from gNB, it detects the PFCP session and sends the PFCP Session Report request to the SMF handling that session along with the Error Indication Report IE. The Error Indication IE also includes the remote F-TEID IE, which contains the GTP-u peer address and the TEID received from the GTP-u Error Indication IE.

Step	Description
1	A PFCP session is established and the data traffic starts running.
2	When gNB clears up the TEID details locally for some reason, it sends the GTP-u Error Indication to UPF for unknown TEID.
3	Once the Error Indication is received, the UPF identifies the PFCP session and sends the PFCP session report request to the SMF.
4	The session report request contains the TEID and the remote IP address from where the Error Indication is received.

Node-level Reporting Procedure due to GTP-u Path Failure Call Flow

When the UPF enables GTP-u Echo procedure for GTP-u endpoints and identifies a data path failure because of no response, it sends a PFCP Node Report Request to the SMF. The Node Report Type in the PFCP Node Report Request is set to User Plane Path Failure Report when it is sent to the SMF. The Node Report procedure

includes only the peer IP address in Remote GTP-u Peer IE – the child IE of the User Plane Path Failure IE – since it is not specific to any PFCP session.

Step	Description
1	Once the PFCP session is established and GTP-u Echo procedure is configured, UPF initiates the GTP-u Echo Request for each peer with at least one GTP-u Tunnel.
2	If there is no GTP-u Echo Response received after a specified number of retries, then the UPF sends the PFCP Node Report Request to the SMF.
3	Only the peer IP address is sent in the Node Report request since it is not a GTP-u tunnel-specific failure.
4	Once the message is received, the SMF sends a Delete and Modify request for all the PFCP Sessions for that gNB to UPF.

PDN Update Procedure - eNodeB F-TEIDu

Feature Description

For S-GW or SAEGW, a procedure to initiate an N4/Sx Modification Request is implemented for:

- eNodeB F-TEIDu update

- Release Access Bearer (RAB) Request for an eNodeB release

How it Works

The PDN update procedure includes the following events for an eNodeB F-TEIDu Update/Release:

- For eNodeB F-TEIDu Update:

  1. The SGW-C initiates N4/Sx Session Modification Request towards SGW-U on receiving a Modify Bearer Request for eNodeB F-TEIDu Update from the MME.

  2. The N4/Sx Modification Request for eNodeB F-TEIDu update contains Update FAR with Apply Action as "Forward" and the updated eNodeB IPv4/IPv6 address in Outer Header Creation, which is a part of the Update Forwarding Parameters IE.

- For eNodeB F-TEIDu Release:

  1. The SGW-C initiates N4/Sx Modification Request towards SGW-U on receiving a RAB Request from the MME.

  2. RAB is a UE-level message. If the UE has multiple PDN connections, then the N4/Sx Modification Request is sent to each PDN connection separately.

  3. SGW-C initiates N4/Sx Session Modification Request towards SGW-U for the N4/Sx session with Update FAR with destination interface as ACCESS. Update FAR contains: FAR ID and Apply Action as Drop. FAR with the destination interface as CORE is not updated.

Standards Compliance

The PDN Update procedure complies with the following standards:

- 3GPP TS 23.401: "General Packet Radio Service (GPRS) enhancements for Evolved Universal Terrestrial Radio Access Network (E-UTRAN) access".

- 3GPP TS 29.274: "3GPP Evolved Packet System (EPS); Evolved General Packet Radio Service (GPRS) Tunnelling Protocol for Control plane (GTPv2-C); Stage 3".

- 3GPP TS 29.244: "Interface between the Control Plane and the User Plane of EPC ".

- 3GPP TS 23.214: "Architecture enhancements for control and user plane separation of EPC nodes; Stage 2"

- 3GPP TS 23.714: "Study on control and user plane separation of EPC nodes"

UEs IPv4, IPv6, and IPv4v6 Support Call Flows

N4 Session Establishment and Modification Procedure for IPv6 Call Flow

The following call flow provides a high-level description of the N4 session establishment and modification procedure for IPv6.

Step	Description
1	The SMF receives the trigger to establish a new PDU session or change the UPF for an established PDU session.
2	The SMF sends N4 Session Establishment Request message to the UPF, in which Create PDR IE has IPv6 UE address, and SDF filter has IPv6 filters to select the PDRs.
3	The UPF responds with an N4 Session Establishment Response message containing any information that the UPF must provide to the SMF in response to the control information received.
4	The SMF interacts with the network function and triggers to modify the PDU session.
5	The SMF sends an N4 Session Modification Request message to the UPF that contains the update for the structured control information which defines how the UPF needs to behave.
6	The UPF identifies the N4 session context to be modified by the N4 Session ID. Then, the UPF updates the parameters of this N4 session context according to the list of parameters sent by the SMF. The UPF responds with an N4 session modification response message containing any information that the UPF must provide to the SMF in response to the control information received.
7	An extra procedure is required for IPv6 sessions, which is Router Solicitation (RS) and Router Advertisement (RA). After the session is established, UE sends the RS message to network to get the link layer address. The UPF forwards this message to SMF. The SMF sends RA message with required parameters to configure the IPv6 address of UE.

Step	Description
8	The RS and RA between UPF and SMF is GTP-u encapsulated and SMF sends an extra pair of PDRs during session establishment and modification procedure, in which GTP-u Tunnel IDs are exchanged for GTP-u tunneling.
9	The additional pair of PDRs that are sent from the SMF are as follows: • One PDR has Source interface as Access, and Destination Interface as CP-Function, to forward IPv6 RS from UE to SMF. The SDF filter is present so that UPF can select this PDR for IPv6 RS from UE • Example: Permit in 58 from any tp ff01::2 any • Another PDR has source interface as CP-Function, and Destination Interface as Access, to forward the IPv6 RA from SMF to UE. The SDF filter is present so that UPF can select this PDR for all the IPv6 RAs from SMF. • Example: Permit out 58 from any to ff01::2 any
10	After RS and RA procedure is completed, the UE sends IPv6 traffic to PDN.

N4 Session Establishment and Modification Procedure for IPv4v6 Call Flow

The following call flow provides a high-level description of the N4 session establishment and modification procedure for IPv4v6.

The IPv4v6 session establishment and modification procedure are similar to the IPv6 session establishment and modification procedure, except for the following procedures:

Step	Description
1	In the session establishment request, a PDR IE is created to include both IPv4 and IPv6 UE addresses. The SDF filter also includes the IPv4 and IPv6 filters for selecting the PDRs.
2	When the IPv6 address is assigned to the UE, an extra procedure – Router Solicitation and Router Advertisement – is required. The UE sends the Router Solicitation message to the network to receive the link layer address, once the session is established. The UPF forwards this message to the SMF and the SMF sends the Router Advertisement Message with the required parameters for configuring the IPv6 address of the UE.
3	The RS/RA between UPF and SMF is GTP-u encapsulated. In addition, the SMF sends an extra pair of PDRs during session establishment and modification procedure, in which the GTP-u tunnel IDs are exchanged for GTP-u tunneling.
4	The additional pair of PDRs that are sent from the SMF are as follows: • For forwarding IPv6 Router Solicitation from UE to SMF, one PDR's source interface is set to access, and its destination interface is set to CP-Function. The SDF filter is present in such a way that the UPF selects this specific PDR for the IPv6 Router Solicitation from the UE. • For instance, permit in 58 from any to ff01::2 any.
5	The UE sends the IPv6 traffic to PDN once the RS/RA procedure is complete.
6	No additional procedures are required for the IPv4 traffic for this PFCP session.

Configuring the N4 Session/Node Level Reporting Procedures

This section describes how to configure the N4 Session/Node Level Reporting procedures.

Enabling the GTP-u Echo Request Procedure

The existing CLI (Command Line Interface) in **gtpu-service** is used to enable the GTP-u Echo request procedure.

```
configure
   gtpu-service service_name
      echo-interval seconds
      echo-retransmission-timeout seconds
      max-retransmissions num
      path-failure detection-policy gtp echo
      end
```

NOTES:

- **gtpu-service** *service_name*: Creates a GTP-u service enters the GTP-u Service Configuration Mode for the current context. *service_name* specifies the name of the GTP-u service.

- **echo-interval** *seconds*: Configures the rate at which GTP v1-u echo packets are sent. *seconds* specifies the number of seconds between the sending of a GTP-uv1 echo packet. It must be an integer in the range of 60–3600.

- **echo-retransmission-timeout** *seconds*: Configures the timeout for GTP-u echo message retransmissions for this service. *seconds* specifies the echo retransmission timeout, in seconds, for the GTP-u service. It must be an integer in the range of 1–20. The default value is 5.

- **max-retransmissions** *num*: Configures the maximum retry limit for GTP-u echo retransmissions. *num* specifies the number of GTP-u echo message retransmissions allowed before triggering a path failure error condition. It must be an integer in the range of 0–15.

- **path-failure detection-policy gtp echo**: Configures a path failure detection policy on GTP-u echo messages that have been retransmitted the maximum number of retry times**. gtp echo** sets the detection policy to detect a failure upon reaching the maximum number of GTP-u echo message retransmissions.

The following is a sample configuration for enabling GTP-u Echo request procedure.

```
configure
   gtpu-service n3-gtpu-service
   echo-interval 60
   echo-retransmission-timeout 5
   max-retransmissions 5
   path-failure detection-policy gtp echo
   end
```

Verifying the N4 Session/Node Level Reporting Procedure Configuration

This section describes how to verify the N4 Session/Node Level Reporting Procedure configuration.

N4 Session Node Level Reporting Procedure OA and M Support

Use the **show gtpu statistics command** to display the GTP-u statistics for Error Indication and GTP-u Echo Request and Response. The following is a sample output from the **show gtpu statistics** command.

```
show gtpu statistics
Path Management Messages:
   Echo Request Rx:                    0 Echo Response Rx:            7
   Echo Request Tx:                   19 Echo Response Tx:            0
   SuppExtnHdr Tx:                     0  SuppExtnHdr   Rx:           0

 Peer Stats:
   Total GTPU Peers:                             1
   Total GTPU Peers with Stats:                  1

 Tunnel Management Messages:
   Error Indication Tx:                          0
   Error Indication Rx:                          0
   Error Indication Rx Discarded:
```

Use the **show sx-service statistics all** command to display the Node report request and response statistics. The following is a sample output of the **show sx-service statistics all** command.

```
show sx-service statistics all
Node Report Request:
 Total TX:                       1    Total RX:                         0
 Initial TX:                     1    Initial RX:                       0
 Retrans TX:                     0    Retrans RX:                       0
 No Rsp received TX:             0    Discarded:                        0
```

```
Node Report Response:
 Total TX:                              0    Total RX:                             1
 Initial TX:                            0    Initial RX:                           1
 Accepted:                              0    Accepted:                             1
 Denied:                                0    Denied:                               0
 Retrans TX:                            0    Discarded:                            0
```

Use the **show user-plane-service statistics all** command to display the statistics of N4 PFCP message parameters. The following is a sample output of the **show user-plane-service statistics all** command.

```
show user-plane-service statistics all
N4 Statistics:
    URR : Created        :           0
          Deleted        :           0
          Queried by ID  :           0
          Queried by all :           0
          Total Queried  :           0
    FAR : Created        :           0
          Updated        :           0
          Removed        :           0
    PDR : Predef rule    :           0
```

![Note pencil icon]

Note   Statistics are cumulative, and are displayed as a total of all session managers.

The descriptions of the fields are as follows:

- URR Created: Displays the total number of URRs created either locally or as requested by the SMF.

- URR Deleted: Displays the total number of URRs removed either locally or as requested by the SMF.

- URR Queried by ID: Displays the total messages received to query a subset of URRs with their specific IDs.

- URR Queried by all: Displays the total messages received to query all the URRs of a session.

- URR Total Queried: Displays the total number of URRs reported to the SMF in response to "Queried by ID" and "Queried by all".

- FAR Created: Displays the total number of FARs created.

- FAR Updated: Displays the total number of FARs updated.

- FAR Removed: Displays the total number of FARs removed.

- PDR Predef rule: Displays the total number of CREATE PDRs that have predefined rules set. For one predefined rule two PDRs are received, and so the counter is updated as 2.

![Note pencil icon]

Note   Total Predef counter isn't updated when a predefined rule is removed. SMF only sends Remove PDR message, which doesn't contain predefined rule name. Also, during HO, for exiting predef rule, PDRs are removed and created again, so it's counted twice for the same session.

SNMP Traps

The following traps are available to track status and conditions GTP-u path failure.

- EGTPUPathFailure: This trap is generated when no response is received for GTP-U ECHO requests and data path failure is detected towards peer EPC Node.

- EGTPUPathFailureClear: This trap is generated when the data path towards the peer node is available.

New Standard QCI Support

Feature Summary and Revision History

Summary Data

Table 51: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

Revision Details	Release
First introduced.	2021.02.0

Feature Description

The 5G-UPF supports new standard QoS Class Index (QCIs) based on 3GPP TS 23.203 Release 12, for Mission Critical and Push-to-Talk (MC/PTT) applications.

As part of this feature, the following functionalities are supported:

- Establishing a LTE/WiFi/5G-NR call with default bearer QCI/QFI with new standard non-GBR QCIs (69, 70 and 80). These are MC-PTT and 5G-NSA QCIs.

- Establishment of dynamic rule with new standard GBR/non-GBR QCI (65, 66, 69, 70, 80, 82, 83).

- Support for Extended QoS Bit Rates for DCNR-enabled UEs with new standard QCIs (80, 82, and 83).

- LTE to 5G and 5G to LTE HO, WiFi to 5G and 5G to WiFi HO, and LTE to WiFi and WiFi to LTE HO are supported.

- DSCP Marking in UL and DL direction based on new standard QCIs.

Limitations

In this release, predefined rules with the new standard QCIs aren't supported.

Configurations

There is no configuration (or License) required at SMF or UPF to enable new standard QCIs.

At SMF, the following configuration is required to enable Extended QoS Bit Rates.

```
configure
  profile dnn name
    dcnr { true | false }
    end
```

- **dcnr**: Specifies to enable dual connectivity with new radio (DCNR).

C H A P T E R **30**

UPF Ingress Interface

This chapter covers the following topics:

Feature Summary and Revision History

Summary Data

Table 52: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Table 53: Revision History

Revision Details	Release
First Introduced.	2020.02.0

Feature Description

This release supports the upf-ingress interface, which the UPF requires for user plane service to start. The user plane service must be associated with GTP-U service. This can be achieved using the **associate gtpu-service** CLI command in User Plane Service configuration mode.

Note To enable **upf-ingress** CLI, you need the **require upf** CLI on the UPF. However, to enable the **require upf** CLI, you need the UPF license.

Configuring UPF Ingress Interface Type Support

To associate the GTPU service with the User Plane Service, use the following configuration:

```
configure
   context context_name
      user-plane-service service_name
         [ no ] associate gtpu-service gtpu_service_name upf-ingress
         end
```

NOTES:

- **associate gtpu-service** *gtpu_service_name*: Associates the GTP-U service with the user plane service.

- **upf-ingress**: Configures the interface type as UPF ingress.

Verifying the UPF Ingress Interface Type Feature Configuration

Run the **show user-plane-service all** command to view the output.

```
[local]qvpc-si# show user-plane-service all

Service name                        : user-plane-service
  Service-Id                        : 4
  Context                           : ingress
  Status                            : STARTED
  UPF Ingress GTPU Service          : n3-gtpu-service
  SGW Ingress GTPU Service          : Not defined
  SGW Egress GTPU Service           : Not defined
  Control Plane Tunnel GTPU Service : control_gtpu
  Sx Service                        : sxu
  Control Plane Group               : g1
  Fast-Path service                 : Disabled
```

NOTES:

- Only one of the interface types **pgw-ingress** or **upf-ingress** can be configured in a single user plane service.

CHAPTER 31

UPF Local Configuration

This chapter covers the following topics:

Feature Summary and Revision History

Summary Data

Table 54: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Table 55: Revision History

Revision Details	Release
First Introduced.	2020.02.0

Feature Description

The support for processing static and predefined rules in Control and User Plane Separation of EPC nodes (CUPS) architecture is dependent on the ruledef, rulebase, and charging action. For processing L3/L4 static and predefined rules, this information is made available at the control-plane in CUPS architecture. The control plane sends all these information to the associated user-plane using the PFD management message. The UPF cannot use the PFD management message to work with CN-SNF. With this feature, the local configuration support for the User Plane Function (UPF) is enabled, which allows the UPF to work with CN-SNF.

How it Works

The Access Control System (ACS) command line interface (CLI) is configured on the user-plane and the CLI module sends it to the ACS Controller (ACSCtrl). The ACSCtrl verifies the CLI and sends it to the Session Controller (SessCtrl). The SessCtrl stores the configuration in the SCT.

The SessCtrl maintains and stores different configuration types in a skiplist. When the length of the skiplist reaches the maximum (BULK configuration length) for a particular configuration type, the entire list is pushed in BULK from the Sessctrl to the Session Manager (SessMgr). As a result, the number of messenger event/message transactions between proclets is greatly reduced since the configurations are sent in BULK in a single message. On the expiry of the bulk configuration timeout (2 seconds), the Bulk Configuration timer – which runs constantly at the Session Controller – pushes the different types of configurations to the SessMgrs.

- The following configuration types are supported for the Bulk Configuration push:

- Ruledef

- Charging Action

- Action Priority Lines

- Group of Ruledef Configuration

- Rule in Group of Ruledef Configuration

- Rulebase L3/L4/L7 Info Configuration

- APN Configuration

- ACS service Configuration

The configurations are pushed only through the bulk push mechanism for configurations that are either added or modified. On the other hand, when configurations are deleted, it is removed immediately without waiting for any response from the Bulk configuration push timer. The deleted configuration is removed from the SCT and other SessMgrs immediately.

Note   The Bulk configuration timeout function is invoked forcefully to push all the pending configurations to the SessMgrs before pushing the configuration delete to avoid any race conditions.

- The configuration changes applied to all the new and existing calls are listed in Table as follows

Table 56: Configuration Changes on New and Existing Call Flows

Change in Configuration	Impact on Existing Calls Current Flows	Impact on Existing Calls New Flows	Impact on New Calls
Existing ruledef contents/New rule addition	Rule match is not enforced on existing flows after configuration change. TRM is not disengaged on existing flows. This may lead to billing issues if ruledef contents were changed for ongoing flows.	The configuration changes apply on new flows. For new flows, anyways fresh rule match would happen and the ruledef changes are applied on new flows for existing calls.	The configuration changes apply on new calls. For new flows, anyways fresh rule match would happen and the ruledef changes are applied on flows for new calls.
No Ruledef	Rule in use cannot be deleted.	Rule in use cannot be deleted	Rule in use cannot be deleted
New Group of Ruledefs/Changes to existing Group of Ruledefs contents (Add or Delete Rule in Group of Ruledefs)	Rule match is not enforced on existing flows after configuration change. TRM is not disengaged on existing flows. This may lead to billing issues if Group of Ruledefs contents were changed for ongoing flows.	The configuration changes apply on new flows. For new flows, anyways fresh rule match would happen and the Group of Ruledefs changes are applied on new flows for existing calls.	The configuration changes apply on new calls. For new flows, anyways fresh rule match would happen and the Group of Ruledefs changes are applied on flows for new calls.
No Group of Ruledefs	Group of Ruledefs in use cannot be deleted	Group of Ruledefs in use cannot be deleted	Group of Ruledefs in use cannot be deleted
No Rule in GoR	Flows continue to match the ruledef defined in Group of Ruledefs unless the ruledef itself is deleted	New flows go through a fresh rule match and configuration change takes effect.	New flows go through a fresh rule match and configuration change takes effect.
Action Priority Changes/Action Priority addition	TRM is not disengaged for ongoing flows. configuration changes do not apply on existing flows	Configuration changes apply on new flows.	Configuration changes apply on new calls.
No Action Priority	No Impact on existing flows	Configuration changes apply on new flows.	Configuration changes apply on new calls.
Rulebase parameters change	Some parameter changes apply on existing calls	Some parameter changes apply on existing calls	Configuration changes apply on new calls
No Rulebase	No Rulebase is not supported	No Rulebase is not supported	No Rulebase is not supported
No APN	No APN is not supported	No APN is not supported	No APN is not supported
IP source violation	No impact on existing calls	No impact on existing calls	Configuration changes apply on new calls.

Configuring the Local Configuration Support for UPF

Use the following CLI commands to configure the User Plane Function (UPF) locally.

```
configure
  require upf
  end
```

UPF Reporting of Load Control Over N4 Interface

Feature Summary and Revision History

Summary Data

Applicable Product(s) or Functional Area	UPF
Applicable Platform(s)	VPC-SI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Revision Details	Release
First introduced.	2020.02.0

Feature Description

Load control enables the UPF to send its load information to the SMF in order to balance PFCP session load across the UPF according to their effective load. The load information reflects the operating status of the resources of the UPF. Load control allows for better balancing of the PFCP session load to prevent overload.

NOTE: Overload mitigation actions are not triggered even if the UPF reports high load.

Supported IE and Messages

To report Load Control Information (LCI) to the SMF, 3GPP specification has defined the following IEs:

- Load Control Information IE – The load control Information IE is as follows: It contains the sequence number IE and load metric IE. This IE is sent in Session Establishment Response, Session Modification Response, Session Deletion Response, and Session Report Request messages sent from UPF.

- Sequence Number IE – The Sequence Number IE contains an Unsigned32 binary integer value. The Load Control Sequence Number increases whenever the load control information changes.

- Load Metric IE – The Load Metric parameter indicates the current load level of the originating node. The computation of the Load Metric happens at the implementation basis. The node considers the various aspects, such as:

  - The used capacity of the UPF

  - The load in the node. For example, memory or CPU usage in relationship to the total memory or CPU available, and so on.

  The Metric IE encoding is as follows: It indicates a percentage and takes binary coded integer values from and including 0 up to and including 100. Considers the other values as 0.

Reporting Load Information to SMF

The UPF sends its load control information to reflect the operating status of its resources at the node level. It allows the SMF to use this information to augment the UPF selection procedures. The load control information is piggybacked in PFCP request or response messages such that the exchange of load control information does not trigger extra signaling.

Considering the processing requirement of the receiver of the load control information, a larger variation in the Load Metric, example 5 or more units are reasonable value to send the new load control information.

The following criteria is used to send the Load Control Information IE:

- Whenever there is an increase or decrease in the load by 5% or more

- At 95% or above, LCI is reported for any increase

- At 5% or below, no LCI is reported

- At 100%, LCI is reported in all messages.

- System timestamp is used as Sequence Number.

  NOTES:

- Currently, only session-load is considered to calculate Load Metric in the UPF.

- Multiple SessMgrs report same value of Load Metric with same sequence number.

Configuring the Max Sessions

Based on various deployment scenarios, if you do not want to load UPF to its maximum capacity in terms of the count of sessions, especially, given that in 5G a single user-session can go up to 5 Gbps. To alter the max session supported in UPF, a CLI command is available under the User Plane Service configuration. It allows the operator to configure the required number of max-sessions that are supported on UPF so that the SMF can load balance the sessions across the UPF. The following is a sample configuration:

```
configure
   context context_name
      user-plane-service user_plane_service
         load-control capacity session_value
         end
```

NOTES:

- *session_value* must be an integer in the range of 1 through the maximum value that is allowed in the platform.

- The use of this configuration is only for the LCI reporting to SMF, and not for any other purpose, such as congestion control, and so on.

The following is an example configuration:

```
configure
   context ingress
      user-plane-service ups1
         load-control capacity 2500
         end
```

Show Command Support

The output of the show command to display the User Plane Service includes the value of configured max sessions.

show user-plane-service all

```
Service name                 : user-plane-service
Service-Id                   : 4
Context                      : ingress
Status                       : STARTED
UPF Ingress GTPU Service     : n3-gtpu-service
SGW Ingress GTPU Service     : Not defined
SGW Egress GTPU Service      : Not defined
Control Plane Tunnel GTPU Service : n4-gtpu-service
Sx Service                   : n4-sx
Control Plane Group          : g1
Load Control Parameters
           Capacity          : 1000
```

CHAPTER **33**

Session Recovery

- Feature Summary and Revision History, on page 263
- Feature Description, on page 263
- How it Works, on page 264
- Configuring the System to Support Session Recovery, on page 264

Feature Summary and Revision History

Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI
Feature Default Setting	Disabled – License Required
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Revision Details	Release
First introduced.	2020.02.0

Feature Description

With robust hardware failover and redundancy protection, any hardware or software failures on the system can quickly be corrected. However, software failures can occur for numerous reasons, often without prior indication.

This chapter describes the Session Recovery feature that provides seamless failover and reconstruction of subscriber session information in the event of a hardware or software fault.

☞

Important Session Recovery is a licensed Cisco feature. A separate feature license may be required. Contact your Cisco Account representative for detailed information on specific licensing requirements.

How it Works

This section provides an overview of how this feature is implemented and the recovery process.

The Session Recovery feature provides seamless failover and reconstruction of subscriber session information in the event of a hardware or software fault within the system preventing a fully connected user session from being disconnected.

Session recovery is performed by mirroring key software processes (for example, session manager and AAA manager) within the system. These mirrored processes remain in an idle state (standby-mode) wherein they perform no processing, until they may be needed in the event of a software failure (for example, a session manager task aborts).

There are some situations wherein session recovery may not operate properly. Additional software or hardware failures occur during the session recovery operation. For example, an AAA manager fails while the state information it contained was being used to populate the newly activated session manager task.

☞

Important After a session recovery operation, some statistics, such as those collected and maintained on a per manager basis (AAA Manager, Session Manager, etc.) are in general not recovered, only accounting and billing related information is checkpointed and recovered.

Configuring the System to Support Session Recovery

The following procedures allow you to configure the session recovery feature for either an operational system that is currently in-service (able to accept incoming calls) or a system that is out-of-service (not part of your production network and, therefore, not processing any live subscriber/customer data).

☞

Important The session recovery feature, even when the feature use key is present, is disabled by default on the system.

Enabling Session Recovery

As noted earlier, session recovery can be enabled on a system that is out-of-service (OOS) and does not yet have any contexts configured, or on an in-service system that is currently capable of processing calls. However, if the system is in-service, it must be restarted before the session recovery feature takes effect.

Enabling Session Recovery on an Out-of-Service System

The following procedure is for a system that does not have any contexts configured.

To enable the session recovery feature on an out-of-service system, follow the procedure below. This procedure assumes that you begin at the Exec mode prompt.

Step 1 At the Exec mode prompt, verify that the session recovery feature is enabled via the session and feature use licenses on the system by running the **show license info** command.

If the current status of the Session Recovery feature is Disabled, you cannot enable this feature until a license key is installed in the system.

Step 2 Use the following configuration example to enable session recovery.

```
configure
   require session recovery
   end
```

Note After you configure this command, you must save the configuration and then reload the chassis for the command to take effect. For information on saving the configuration file and reloading the chassis, refer to the System Administration Guide for your deployment.

Step 3 Save your configuration as described in *Verifying and Saving Your Configuration*.

The system, when started, enables session recovery, creates all mirrored "standby-mode" tasks, and performs packet processing card reservations and other operations automatically.

Step 4 After the system has been configured and placed in-service, you should verify the preparedness of the system to support this feature as described in *Viewing Session Recovery Status* section.

Enabling Session Recovery on an In-Service System

When enabling session recovery on a system that already has a saved configuration, the session recovery commands are automatically placed before any service configuration commands in the configuration file.

To enable the session recovery feature on an in-service system, follow the procedure below. This procedure assumes that you begin at the Exec mode prompt.

Step 1 At the Exec mode prompt, verify that the session recovery feature is enabled via the session and feature use licenses on the system by running the **show license info** command:

If the current status of the Session Recovery feature is Disabled, You cannot enable this feature until a license key is installed in the system.

Step 2 Use the following configuration example to enable session recovery.

```
configure
   require session recovery
   end
```

This feature does <u>not</u> take effect until after the system has been restarted.

Step 3 Save your configuration as described in *Verifying and Saving Your Configuration*.

Step 4 Perform a system restart by entering the **reload** command:

The following prompt appears:

```
Are you sure&quest; [Yes|No]:
```

Confirm your desire to perform a system restart by entering **yes**.

The system, when restarted, enables session recovery and creates all mirrored "standby-mode" tasks, performs packet processing card reservations, and other operations automatically.

Step 5   After the system has been restarted, you should verify the preparedness of the system to support this feature as described in *Viewing Session Recovery Status* section.

More advanced users may opt to simply insert the **require session recovery** command syntax into an existing configuration file using a text editor or other means, and then applying the configuration file manually. Exercise caution when doing this to ensure that this command is placed among the first few lines of any existing configuration file; it must appear before the creation of any non-local context.

Disabling the Session Recovery Feature

To disable the session recovery feature on a system, enter the **no require session recovery** command from the Global Configuration mode prompt.

☞

Important   If this command is issued on an in-service system, then the system must be restarted by issuing the **reload** command.

Viewing Session Recovery Status

To determine if the system is capable of performing session recovery, when enabled, enter the **show session recovery status verbose** command from the Exec mode prompt.

The output of this command should be similar to the examples shown below.

```
[local]host_name# show session recovery status
Session Recovery Status:
    Overall Status           :  SESSMGR Not Ready For Recovery
    Last Status Update       :  1 second ago

[local]host_name# show session recovery status
Session Recovery Status:
    Overall Status           :  Ready For Recovery
    Last Status Update       :  8 seconds ago

[local]host_name# show session recovery status verbose
Session Recovery Status:
    Overall Status           :  Ready For Recovery
    Last Status Update       :  2 seconds ago

            ----sessmgr----        ----aaamgr----        demux
  cpu state active   standby      active   standby      active   status
  ---- ------- ------  -------      ------  -------      ------  ------------
  1/0 Active  7        1            7        1            7        Good
[local]host_name#
```

Viewing Recreated Session Information

To view session state information and any session recreation status, enter the following command:

`show subscriber debug-info callid` *`id`*

The following example shows the output of this command both before and after a session recovery operation has been performed. The "Redundancy Status" fields in this example have been bold-faced for clarity.

```
username: user1              callid: 01ca11b1              msid: 0000100003
    Card/Cpu: 4/2
    Sessmgr Instance: 7
    Primary callline:
    Redundancy Status: Original Session
      Checkpoints       Attempts        Success       Last-Attempt      Last-Success
         Full:               69              68           29800ms           29800ms
         Micro:             206             206           20100ms           20100ms
      Current state: SMGR_STATE_CONNECTED
      FSM Event trace:
            State                            Event
            SMGR_STATE_OPEN                  SMGR_EVT_NEWCALL
            SMGR_STATE_NEWCALL_ARRIVED       SMGR_EVT_ANSWER_CALL
            SMGR_STATE_NEWCALL_ANSWERED      SMGR_EVT_LINE_CONNECTED
            SMGR_STATE_LINE_CONNECTED        SMGR_EVT_LINK_CONTROL_UP
            SMGR_STATE_LINE_CONNECTED        SMGR_EVT_AUTH_REQ
            SMGR_STATE_LINE_CONNECTED        SMGR_EVT_IPADDR_ALLOC_SUCCESS
            SMGR_STATE_LINE_CONNECTED        SMGR_EVT_AUTH_SUCCESS
            SMGR_STATE_LINE_CONNECTED        SMGR_EVT_UPDATE_SESS_CONFIG
            SMGR_STATE_LINE_CONNECTED        SMGR_EVT_LOWER_LAYER_UP
    Data Reorder statistics
    Total timer expiry:            0       Total flush (tmr expiry):   0
          Total no buffers:        0       Total flush (no buffers):   0
          Total flush (queue full): 0      Total flush (out of range): 0
          Total flush (svc change): 0      Total out-of-seq pkt drop:  0
          Total out-of-seq arrived: 0
    IPv4 Reassembly Statistics:
          Success:                 0       In Progress: 0
          Failure (timeout):       0       Failure (no buffers): 0
          Failure (other reasons): 0
    Redirected Session Entries:          Allowed:
2000          Current:                              0
              Added:                                0              Deleted:
                     0
              Revoked for use by different subscriber: 0
    Peer callline:
     Redundancy Status: Recreated Session
    Checkpoints       Attempts        Success       Last-Attempt      Last-Success
         Full:               0               0             0ms               0ms
         Micro:              0               0             0ms               0ms
      Current state: SMGR_STATE_CONNECTED
      FSM Event trace:
            State                                         Event
            SMGR_STATE_LINE_CONNECTED        SMGR_EVT_LOWER_LAYER_UP
            SMGR_STATE_CONNECTED             SMGR_EVT_AUTH_REQ
            SMGR_STATE_CONNECTED             SMGR_EVT_AUTH_SUCCESS
            SMGR_STATE_CONNECTED             SMGR_EVT_REQ_SUB_SESSION
            SMGR_STATE_CONNECTED             SMGR_EVT_RSP_SUB_SESSION
            SMGR_STATE_CONNECTED             SMGR_EVT_ADD_SUB_SESSION
            SMGR_STATE_CONNECTED             SMGR_EVT_AUTH_REQ
            SMGR_STATE_CONNECTED             SMGR_EVT_AUTH_SUCCESS
            SMGR_STATE_CONNECTED             SMGR_EVT_AUTH_REQ
            SMGR_STATE_CONNECTED             SMGR_EVT_AUTH_SUCCESS
            SMGR_STATE_CONNECTED             SMGR_EVT_AUTH_REQ
```

```
            SMGR_STATE_CONNECTED                    SMGR_EVT_AUTH_SUCCESS
            SMGR_STATE_CONNECTED                    SMGR_EVT_AUTH_REQ
            SMGR_STATE_CONNECTED                    SMGR_EVT_AUTH_SUCCESS
            SMGR_STATE_CONNECTED                    SMGR_EVT_AUTH_REQ
            SMGR_STATE_CONNECTED                    SMGR_EVT_AUTH_SUCCESS
Data Reorder statistics
      Total timer expiry:        0         Total flush (tmr expiry):  0
      Total no buffers:          0         Total flush (no buffers):  0
      Total flush (queue full):  0         Total flush (out of range):0
      Total flush (svc change):  0         Total out-of-seq pkt drop: 0
      Total out-of-seq arrived:  0
IPv4 Reassembly Statistics:
      Success:                   0         In Progress:               0
      Failure (timeout):         0         Failure (no buffers):      0
      Failure (other reasons):   0
Redirected Session Entries:
      Allowed:                2000         Current:                   0
      Added:                               Deleted:                   0
      Revoked for use by different subscriber:  0
```

CHAPTER 34

Session Report Rejection Procedure

- Feature Summary and Revision History, on page 269
- Feature Description, on page 269
- OAM Support, on page 271

Feature Summary and Revision History

Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G SMF Configuration and Administration Guide*

Revision History

Revision Details	Release
First introduced.	2021.02.0

Feature Description

The SMF rejects the UPF-originated Session Report Request with a specific cause code during any mismatch in the charging configuration of SMF and UPF.

For any session report rejection by the SMF, the UPF locally purges the sessions. The SMF is unaware of the purging operation and continues to send the N4 message to the UPF. This action triggers the UPF to send "context not found" message to the SMF for the locally purged sessions.

This behavior impacts the UE experience and results in the loss of charging data. So, the current implementation of handling the session report errors is modified to avoid local purging of sessions on the UPF and also to support graceful clearing of sessions.

With this modification, the UPF ignores the Session Report Error Response. The SMF triggers the Session Deletion Request followed by the rejection of Session Report. The UPF responds to the delete request and clears the session gracefully.

Relationships to Other Features

This feature involves implementing some behavioral changes to the SMF and the UPF. The new CLI configuration in SMF aids in controlling this behavior. For details on the SMF behavioral changes, see the *Support for Session Report Rejection Procedure* chapter in the *UCC 5G SMF Configuration and Administration Guide*.

Call Flow

The following figure explains the flow of Session Report rejection prcedure.

Table 57: Session Report Rejection Procedure Call Flow

Step	Description
1	UPF sends Session Report message to SMF. (SX_SESSION_REPORT_REQUEST)
2	SMF rejects the message with some error code. (SX_SESSION_REPORT_RESPONSE with error code)
3	UPF ignores the Session Report error response. SMF initiates Session Delete Request to tear down the session gracefully. (SX_SESSION_DELETION_REQUEST)

Step	Description
4	UPF responds to the delete request and clears the session gracefully. (SX_SESSION_DELETION_RESPONSE)

OAM Support

This section describes operations, administration, and maintenance information for this feature.

Show Command(s) and/or Output(s)

show user-plane-service statistics all

The counter, "Skipped Local Purge", is added to the **show user-plane-service statistics all** CLI command under the section "PDNs Released By Reason".

- **Skipped Local Purge**: Increments whenever Session Report error is triggered from SMF/CP to UPF/UP.

CHAPTER 35

Smart Licensing

Feature Summary and Revision History

Summary Data

Table 58: Summary Data

Applicable Product (s) or Functional Area	5G-UPF
Applicable Platforms	VPC-SI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

Revision Details	Release
First introduced.	2021.01.0

Overview

Ultra Cloud Core 5G User Plane Function (UPF) supports Smart Licensing. Smart Licensing is a cloud-based approach to licensing that simplifies the purchase, deployment, and management of Cisco software assets.

Entitlements are purchased through your Cisco account via Cisco Commerce Workspace (CCW) and immediately deposited into your Virtual Account for usage. This eliminates the need to install license files on every device. Products that are smart-enabled, communicate directly to Cisco to report consumption. A single location is available to customers to manage Cisco software licenses—the Cisco Smart Software Manager (CSSM). License ownership and consumption are readily available to help make better purchase decision based on consumption or business need.

See https://www.cisco.com/c/en/us/buy/smart-accounts/software-licensing.html for more information about Cisco Smart Licensing.

Comparison Between Legacy Licensing and Smart Licensing

Cisco employs two types of license models - Legacy Licensing and Smart Software Licensing. **Legacy Licensing** consists of software activation by installing Product Activation Keys (PAK) on to the Cisco product. A Product Activation Key is a purchasable item, ordered in the same manner as other Cisco equipment and used to obtain license files for feature set on Cisco Products. **Smart Software Licensing** is a cloud-based licensing of the end-to-end platform leveraging few tools that authorize and deliver license reporting. Smart Software Licensing functionality incorporated into Network Functions (NFs) completes the product registration, authorization resulting in reporting services available to the end customer.

Evaluation Period

A 90-day evaluation period is granted for all licenses in use. During this period, feature licenses can be used without limitation, and up to one counting license each can be used. The evaluation period ends when the system registers successfully with the CSSM or Cisco.com. Licensed functionality is blocked when this 90-day period expires.

UPF performs license enforcement for on/off feature licenses. Each on/off feature license is tied to service licenses, which potentially use those on/off features. When an Out of Compliance (OOC) is detected for an on/off license, new calls for the corresponding services will be dropped, subject to the following conditions:

- Each on/off feature license is given a 90-day grace (evaluation) period. During this period, the system generates SNMP traps to inform of the unavailability of valid licenses. To resolve the OOC, corrective action is needed such as purchasing and registering licenses for this feature, or disabling the feature.

- If the feature is still OOC after the 90-day grace period, UPF enforces the OOC state based on a predefined policy for each license. If enforcement is required, new calls for the services corresponding to the on/off licenses are dropped.

The following CLI commands can be used to display details about the enforcement of Smart Licenses in use:

```
show license enforcement policy
show license enforcement status [ allowed | blocked ] [ feature | service
 ]
```

Cisco Smart Software Manager

Cisco Smart Software Manager (CSSM) enables the management of software licenses and Smart Account from a single portal. The interface allows you to activate your product, manage entitlements, and renew and upgrade software. A functioning Smart Account is required to complete the registration process. To access the Cisco Smart Software Manager, see https://software.cisco.com.

Smart Accounts/Virtual Accounts

A Smart Account provides a single location for all Smart-enabled products and entitlements. It helps speed procurement, deployment, and maintenance of Cisco Software. When creating a Smart Account, you must have the authority to represent the requesting organization. After submitting, the request goes through a brief approval process.

A Virtual Account exists as a sub-account withing the Smart Account. Virtual Accounts are a customer-defined structure based on organizational layout, business function, geography or any defined hierarchy. They are created and maintained by the Smart Account administrator.

See https://software.cisco.com to learn about, set up, or manage Smart Accounts.

Smart Licensing Mode

The Smart Licensing Mode is categorized as follows:

- **Reporting Licenses (Parent Licenses)**: The Parent Licenses are reported to backend license server (CSSM) and accounted for usage of licenses. For each Parent Licenses, the entitlement tags are created and the same is used to identify the type service or feature.

- **Non-Reporting Licenses (Child Licenses)**: The Child Licenses are not reported to backend license server (CSSM) and these licenses are enabled by default with the Parent Licenses. For Child Licenses, the entitlement tags are not created.

That is to say, Smart License enables all Parent and Child Licenses based on the Product Type that is configured. However, the reporting is done only for Parent Licenses.

The state of Smart Licensing Agent is persistent across reboot and crashes.

Request a Cisco Smart Account

A Cisco Smart Account is an account where all products enabled for Smart Licensing are deposited. A Cisco Smart Account allows you to manage and activate your licenses to devices, monitor license use, and track Cisco license purchases. Through transparent access, you have a real-time view into your Smart Licensing products. IT administrators can manage licenses and account users within your organization's Smart Account through the Smart Software Manager.

Step 1 In a browser window, enter the following URL:

```
https://software.cisco.com
```

Step 2 Log in using your credentials, and then click **Request a Smart Account** in the **Administration** area.

The **Smart Account Request** window is displayed.

Step 3 Under **Create Account**, select one of the following options:

- **Yes, I have authority to represent my company and want to create the Smart Account** – If you select this option, you agree to authorization to create and manage product and service entitlements, users, and roles on behalf of your organization.

- **No, the person specified below will create the account** – If you select this option, you must enter the email address of the person who will create the Smart Account.

Step 4   Under **Account Information**:

   a) Click **Edit** beside **Account Domain Identifier**.

   b) In the **Edit Account Identifier** dialog box, enter the domain, and click **OK**. By default, the domain is based on the email address of the person creating the account and must belong to the company that will own this account.

   c) Enter the **Account Name** (typically, the company name).

Step 5   Click **Continue**.
The Smart Account request will be in pending status until it has been approved by the Account Domain Identifier. After approval, you will receive an email confirmation with instructions for completing the setup process.

Software Tags and Entitlement Tags

Tags for the following software and entitlements have been created to identify, report, and enforce licenses.

Software Tags

Software tags uniquely identify each licenseable software product or product suite on a device. The following software tags exist for UPF.

Product Type / Description	Software Tag
UPF Ultra Cloud Core - User Plane Function (UPF)	regid.2020-04.com.cisco.UPF, 1.0_bc18a9ff-e0ea-4476-a250-04ebf7839c4c

Reporting (Parent) Entitlement Tags for UPF

The following entitlement tags indentify licenses in use for UPF.

License Display Name/Description	Entitlement Tag	Tag Name
UCC 5G UPF Base Lic Ultra Cloud Core - User Plane Function (UPF), Base Minimum	regid.2020-08.com.cisco.F_UPF_BASE, 1.0_776395f3-8b8d-46e1-ac6e-0bd2306ef3b6	F_UPF_BASE
UCC 5G UPF Instance Ultra Cloud Core - User Plane Function (UPF) Instance	regid.2020-08.com.cisco.F_UPF_INS, 1.0_5cd68c07-152a-48c6-b143-4dc60eb111e5	F_UPF_INS
UCC 5G UPF 1K Sess Ultra Cloud Core - User Plane Function (UPF), 1K Sessions	regid.2020-08.com.cisco.L_UPF_SAE_1K, 1.0_5d16e2f6-808a-45ff-8691-f215d5ba2bea	L_UPF_SAE_1K

Non-reporting (Child) License List

In this release, the following Child Licenses are enabled by default when the Parent Licenses are enabled.

License Description	License Type
PGW 1k Sessions	Counting
SGW 1k Sessions	Counting
GGSN 1k Sessions	Counting
Per Subscriber Stateful Firewall 1k Sessions	Counting
ENAT 1k Sessions	Counting
Enhanced Charging Bundle 1	Counting
Enhanced charging bundle 2	On/Off
Dynamic policy interface	On/Off
Enhanced LI service	On/Off
Lawful intercept	On/Off
Session recover	On/Off
Radius AAA server group	On/Off
IPv6	On/Off
Intelligent Traffic Control	On/Off
DIAMETER Closed-Loop Charging Interface	On/Off
Per-Subscriber Traffic Policing/Shaping	On/Off
Dynamic Radius extensions (CoA and PoD)	On/Off
Proxy MIP	On/Off
FA	On/Off
IPSec	On/Off
Inter-Chassis Session Recovery	On/Off
ICSR/SR Performance Improvements	On/Off
ICSR Enhanced Recovery for Data and Control Plane, 1K Sessions	On/Off
MPLS	On/Off
TACACS+	On/Off
NAT/PAT With DPI	On/Off
Rate Limiting Function (Throttling)	On/Off
Overcharging Protection for EPC-GW	On/Off
Overcharging Protection Upgrade for EPC-GW	On/Off
ADC Trigger Over Gx, 1K Sessions	On/Off

License Description	License Type
Gx Based Virtual APN Selection, 1K Sessions	On/Off
EPC-GW Support for Wi-Fi Integration, 1K Sessions	On/Off
EPC-GW Non-Standard QCI Support, 1K Sessions	On/Off
Local Policy Decision Engine	On/Off
Header Enrichment	On/Off
HTTP Header Encryption	On/Off
HTTP Header Enrichment and Encryption	On/Off
Broadcast & Multicast Services	On/Off
Integrated Content Filtering Provisioned Service	On/Off
Application Detection and Control 1k Sessions	Counting
5G NSA Feature Set 100K Sess VPCSW Active 1k Sessions	Counting
5G NSA Enablement Fee, Network Wide	On/Off
Multimedia Priority Service Feature Set,1K Sessions	On/Off
EPC Gw VoLTE enhancements	On/Off
DNS Snooping	On/Off

Configuring Smart Licensing

Before you begin, ensure you have:

- Created a Smart Licensing account on https://software.cisco.com.

- Registered your products on https://software.cisco.com using the Product Instance Registration tokens created as part of Smart Account/Virtual Account.

- Enabled a communication path between the UPF system to the CSSM server or Cisco.com.

Enable Smart Licensing

By default, Smart Licensing is disabled in UPF. To enable Smart Licensing, enter the following Global Configuration mode commands:

```
configure
  license smart product upf
  license smart enable
  end
```

NOTE: Before enabling Smart Licensing, Product Type must be configured to enable default licenses that are based on product type.

Enter the following command to verify the configuration:

```
show configuration | grep license
```

Register the Device with Cisco

Using the Product Instance Registration token ID provided when you registered the products on https://software.cisco.com, register the system using the following Exec mode command:

license smart register idtoken *token*

The system now automatically reports entitlement usage count to the CSSM server and receives a compliance status. This also removes the system from "Evaluation Mode".

To show the compliance status, enter any of the following Exec mode commands:

show license status
show license summary
show license statistics

The registration for the system is renewed automatically every 180 days. If needed, use the following Exec mode command to renew the registration information manually:

license smart renew id

The license authorization for the system is renewed automatically every 30 days. If needed, use the following Exec mode command to renew the license authorization manually:

license smart renew auth

To unregister a device, enter the following Exec mode command:

license smart deregister

Changing Smart Transport URL

Smart Agent uses Smart Transport to communicate to Cisco CSSM server. Smart Transport uses the configured URL to identify destination URL where CSSM is reachable. This will not initiate any communication with Cisco. If needed, enter the following Configuration mode commands:

configure
  license smart transport smart
  license smart url *https_link*

Handling Out of Compliance

If there are not enough licenses in the virtual account for a given SKU, CSSM sends an Out Of Compliance (OOC) message to the device. The system stops allowing additional sessions until the OOC state is cleared. The OOC state is cleared when the device receives an authorized response.

Monitoring and Troubleshooting Smart Licensing

Enter the following Exec mode command to verify the Smart Licensing configuration:

show configuration | grep license

The following Exec mode commands display information about Smart Licensing:

show license { all | enforcement | smart-tags | statistics | status | summary | tech-support | udi | usage }

NOTES:

- **all** - Shows a superset of information that includes show status, show usage, show UDI, as well as the Smart Licensing agent version.

- **enforcement { policy | status [allowed | blocked] [feature | service] }** - Shows the enforcement policy applied or current enforcement status of Smart Licenses. Status information can be filtered to show only the licenses which are currently allowed or blocked, or by type (feature license or service license).

- **smart-tags [feature | service]** - Shows the features and services that are currently supported and the corresponding Smart Entitlement Tag.

- **statistics [verbose]** - Shows individual feature license status.

- **status** - Shows overall Smart Licensing status information.

- **summary** - Shows summary of Smart Licensing status.

- **tech-support** - Shows information useful for debugging issues with Smart Licensing.

- **udi** - Shows details for all Unique Device Identifiers (UDI).

- **usage** - Shows the usage information for all entitlements that are currently in use.

CHAPTER 36

Virtual Routing and Forwarding

Feature Summary and Revision History

Summary Data

Table 59: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI
Feature Default Setting	Disabled - Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

Table 60: Revision History

Revision Details	Release
Support is added for the following functionality: • Overlapping IP Pools • Removal of mandatory VRF ordering between SMF and UPF.	2021.01.0
First introduced.	2020.02.0

Feature Description

Virtual Routing and Forwarding (VRF) is a technology that allows multiple instances of a routing table to co-exist within the same router at the same time. As the routing instances are independent, the same or overlapping IP addresses can be used without conflicting with each other.

In UPF, the feature enables association of IP address pools with VRF. These IP pools are chunked like any pools. The chunks from this pool are allocated to the UPFs that are configured to use these pools. VRF-associated pools in UPF can only be of type—STATIC or PRIVATE.

The chunks from the PRIVATE VRF pool are allocated when the UPF comes up for registration, similar to the normal private pools. For STATIC VRF pool, the SMF does chunk allocation to UPF during the time of configuration. An SX-Route-Update message is sent for pre-allocated static-chunks during UPF registration.

Overlapping IP Pool

Overlapping pools share and use an IP address range. Overlapping pools can either be of type STATIC or PRIVATE. No public pools can be configured as overlapping pools. Each overlapping pool is part of different VRF (routing domain) and pool-group. Since an APN can use only one pool-group, overlapping pools are part of different APN as well.

Without this functionality, overlapping pools can be configured at SMF. However, chunks from two overlapping pools can't be sent to the same UPF. That is, the UP can't handle chunks from two different overlapping pools. So, same number of UPFs and overlapping pools are required for sharing same IP address range.

With this functionality, UPF can handle chunks from two different overlapping pools. So, a single UP can handle any number of overlapping pools sharing the same IP range.

The functionality of overlapping pools in same UPF includes:

- When a chunk from particular pool is installed on an UP, its corresponding vrf-name is sent along with the chunk.

- The UPs are made VRF-aware of chunks and therefore, UPs install chunks on the corresponding VRFs and the chunk database is populated under the VRFs.

- During call allocation, release, recovery, or any communication towards VPNMgr, the corresponding SessMgr at UP includes vrf-id. This enables VPNMgr to pick the correct chunk for that IP under the provided vrf-id for processing.

A custom IE, UE IP VRF, is introduced to encapsulate VRF name of UE IP in N4 SESSION ESTABISHMENT REQUEST message.

SessMgr in UPF converts the received vrf- name into UP VRF CONTEXT ID and passes it on to UPFs VPNMgr in IP allocation request. This UP VRF CONTEXT ID is also used when release request is sent to UPFs VPNMgr. UPFs VPNMgr is made aware of the VRF to which that chunk belongs to by sending vrf-name in each chunk related communication between SMF and UPF. This enables UPFs VPNMgr to create database of chunks under each VRF enabling support of overlapping pools in same UPF.

UE IP VRF Information Element

The following is the IE format of the private UE IP VRF.

Table 61: UE IP VRF Format

	Bits							
Octets	8	7	6	5	4	3	2	1
1 to 2	Type = 242 (decimal)							
3 to 4	Length = n							
5	Spare					Identical VRF flag	IPv6 VRF Valid	IPv4 VRF Valid
m to m+1	VRF-1 Name Length = p							
m+1 to m+1+p	VRF-1 Name							
n to n+1	VRF-2 Name Length = q							
m+1 to m+1+q	VRF-2 Name							

The following table shows the possible values of the "UE IP VRF" fields.

Cases	UE IP VRF	Value (binary)		
		Bit 3	Bit 2	Bit 1
1	None of the IPv4 and IPv6 UE IP address are associated to VRF.	0	0	0
2	Only IPv4 UE IP address is associated to an VRF	0	0	1
3	Only IPv6 UE IP address is associated to VRF	0	1	0
4	Both IPv4 and IPv6 UE IP address are associated to VRF and are different VRF.	0	1	1
5	Both IPv4 and IPv6 UE IP address are associated to VRF and is common VRF.	1	1	1

VRF Name as Identifier

Prior to this feature, the communication between SMF and UPF, related to VRF, was done through vrf-id. This required the operator to have all the VRFs configured in both SMF and UPF, and also in the same order.

With this feature, vrf-name is used as identifier in all the communication between SMF and UPF with respect to VRFs, eliminating the requirement of configuring all the VRFs in UPF. Operator can configure VRFs in different order at SMF and UPF, and still can identify the VRF since vrf-name is same in both the nodes.

Limitations and Restrictions

The following are the known limitations/restrictions of the feature in UPF:

- VRF limit per UPF is 205.

- Only VRF-based overlapping pools are supported in UPF. Other flavors of overlapping pools, like NH-based, VLAN-based, and so on, are not supported.

- PDN Type IPv4v6-based call is not permitted on static IP pools with multiple UP in same UP-group.

- Dynamic update of the VRF is not supported.

Configuring VRF

Follow these steps to implement VRF support in UPF.

At SMF:

1. Create APN/DNN profile.

2. Create overlapping IP pools and associate the respective APN/DNN and VRF at context-level.

3. Associate APN/DNN to UPF profile.

The following is an example of the SMF configuration:

```
profile dnn intershat1
.
.
.
   upf apn mpls1.com
exit
profile dnn intershat2
.
.
.
   upf apn mpls2.com
exit
profile network-element upf upf1
.
.
.
   dnn-list [ intershat1 intershat2 ]
exit
profile network-element upf upf2
.
.
.
   dnn-list [ intershat1 intershat2 ]
exit
ipam
   source local
   address-pool pool-intershat1
      vrf-name mpls-vrf-1@isp
      tags
         dnn intershat1
      exit
      ipv4
         address-range 192.0.0.1 255.0.0.1
      exit
   exit
   address-pool pool-intershat2
      vrf-name mpls-vrf-2@isp
      tags
         dnn intershat2
```

```
            exit
            ipv4
                address-range 192.0.0.1 255.0.0.1
            exit
        exit
exit
```

At UPF:

It's recommended to configure VRF in UPF before chunk is pushed from SMF. Else, it leads to the failure of complete IP pool transaction (including chunks that don't belong to the VRF), and retry attempt by SMF after some time.

The following is an example of the UPF configurations:

UPF 1:

```
config
  context EPC2
    sx-service sx
      instance-type userplane
      bind ipv4-address 192.0.0.2 ipv6-address bbbb:aaaa::4
    exit
    user-plane-service up
      associate gtpu-service pgw-gtpu pgw-ingress
      associate gtpu-service sgw-ingress-gtpu sgw-ingress
      associate gtpu-service sgw-engress-gtpu sgw-egress
      associate gtpu-service saegw-sxu cp-tunnel
      associate sx-service sx
      associate fast-path service
      associate control-plane-group g1
    exit

  context isp
    ip vrf mpls-vrf-1
    #exit
    ip vrf mpls-vrf-2
    #exit
    apn mpls1.com
      pdp-type ipv4 ipv6
      bearer-control-mode mixed
      selection-mode sent-by-ms
      ip context-name isp
    exit
exit
control-plane-group g1
    peer-node-id ipv4-address 192.0.0.10
  #exit
  user-plane-group default
```

UPF 2:

```
config
  context EPC2
    sx-service sx
      instance-type userplane
      bind ipv4-address 192.0.0.3 ipv6-address bbbb:aaaa::5
    exit
    user-plane-service up
      associate gtpu-service pgw-gtpu pgw-ingress
      associate gtpu-service sgw-ingress-gtpu sgw-ingress
      associate gtpu-service sgw-engress-gtpu sgw-egress
      associate gtpu-service saegw-sxu cp-tunnel
      associate sx-service sx
      associate fast-path service
```

```
        associate control-plane-group g1
      exit
exit

  context isp
    ip vrf mpls-vrf-1
    #exit
    ip vrf mpls-vrf-2
    #exit
    apn mpls2.com
      pdp-type ipv4 ipv6
      bearer-control-mode mixed
      selection-mode sent-by-ms
      ip context-name isp
    exit
exit

control-plane-group g1
    peer-node-id ipv4-address 192.0.0.10
  #exit
  user-plane-group default
```

Monitoring and Troubleshooting

This section provides information regarding the CLI commands available for monitoring and troubleshooting the feature.

Show Commands and/or Outputs

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

show ip chunks

The output of this CLI command displays all the chunks in that context.

With Overlapping IP Pools functionality, VRF option is introduced in the output of **show ip chunks vrf** *vrf_name* CLI command that displays only the chunks under that VRF.

- chunk-id

- chunk-size

- vrf-name

- start-addr

- end-addr

- used-addrs

- Peer Address

show ipv6 chunks

The output of this CLI command displays all the chunks in that context.

With Overlapping IP Pools functionality, VRF option is introduced in the output of **show ipv6 chunks vrf** *vrf_name* CLI command that displays only the chunks under that VRF.

- chunk-id

- chunk-size

- vrf-name

- start-prefix

- end-prefix

- used-prefixes

- Peer Address

C H A P T E R 37

Voice over New Radio

This chapter covers the following topics:

Feature Summary and Revision History

Summary Data

Table 62: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Table 63: Revision History

Revision Details	Release
First Introduced.	2020.02.0

Feature Description

The UPF supports Voice over New Radio (VoNR) with the existing Session Establishment and Modification procedures. In these procedures, the SMF creates the PDR for 5QI=5 Non-GBR flow for IMS signaling and

PDR for 5QI=1 GBR flow for voice traffic. The UPF does not require any special handling to support mobile-originated or mobile-terminated call flows.

How it Works

The following are the steps in the call flow in which the PDRs are created with 5QI value 5 for IMS signaling and 5QI value 1 or Voice Traffic.

VoNR Call Flow for UPF

This section describes the steps for VoNR session and respective PDR Creation on UPF.

Figure 14: VoNR Call Flow

Table 64: VoNR Call Flow

Step	Description
1	UE1 and UE2 are registered on 5G network. They establish the IMS PDU session and both are registered with same IMS. Both UEs are in connected mode.
2	For IMS signaling, a non-GBR QoS PDRs (UL and DL) is created by SMF which has the QFI value that corresponds to 5QI value 5.
3	Similarly, for Conversational Voice traffic, the PDRs for GBR QoS flow is created with the QFI value that corresponds to 5QI value 1.
4	The QFI to 5QI mapping is maintained by SMF, hence the QFI does not have the values same as 5QI.
5	The above steps are valid for both Mobile Originated (MO) or Mobile Terminate (MT) call flows.
6	Refer to 3GPP TS 23.501, Section 5.7.4 for other types of 5QI mappings for GBR or Non-GBR flows.

PART II

Troubleshooting Information

C H A P T E R **38**

UPF Troubleshooting Information

This chapter covers the following topics related to monitoring and troubleshooting the UPF features:

Debug Logging

👉

Important   The debug logging CLIs must be enabled with the help of System Administrator. Enabling debug logging CLIs can be resource intensive.

Use the following debug CLIs as required:

- **logg filter active facility sx level debug**

- **logg filter active facility user-data level debug**

- **logg filter active facility sessmgr level debug**

- **logg filter active facility uplane level debug**

- **logg filter active facility egtpc level debug**

- **logg filter active facility gtpu level debug**

- **logg filter active facility egtpu level debug**

- **logg filter active facility gtpumgr level debug**

- **logg filter active facility sxdemux level debug**

- **logg filter active facility user-l3tunnel level debug**

> - **logg filter active facility aaamgr level debug**
>
> - **logg filter active facility vpp level debug**
>
> - **logg filter active facility dpath level debug**
>
> - **logg active pdu-verbosity 5**

Monitoring CLI

Subscriber Level Message

Use the **mon sub callid** CLI command for subscriber level message.

Resource Tracking

Use the **show task resources facility** *sessmgr* **all** CLI command to track the CPU/Memory for PROCLET.

Service Status

Use the **show service all** CLI command to check the service status.

Sx Peer Status

Use the **show sx peers** CLI command to check the Sx peer status.

Monitoring Protocol

When using the monitor protocol command, enable option 49 for PFCP, and option 26 for GTP-U.

RAT Type-based Statistics

The RAT Type-based Statistics feature equip users to view data statistics segregated by RAT Type in UPF.

RAT Type-based data statistics in UPF maintains separate buckets. These buckets are created at Session Manager instance level. Bucket is assigned to a subscriber at the time of call-setup, based on RAT Type IE received in "Subscriber-Parameters". If the IE is not received, "Unknown" RAT Type bucket is assigned to that subscriber. During the session, if UPF receives a new RAT Type for a subscriber, the bucket is changed accordingly.

👉

Important Data statistics are not checkpointed and lost during Session Recovery/ICSR. Only "Current-Subscriber" statistics are recalculated after recovery (during the time of call-audit).

Show Command and Output

The following CLI command displays node-level RAT statistics for UPF: **show user-plane-service statistics rat { 5g-nr | all | eutran | unknown | wlan }**

NOTES:

- **5g-nr**: Displays the data statistics for 5G NR subscribers.

- **all**: Displays the data statistics for all RAT Type subscribers.

- **eutran**: Displays the data statistics for EUTRAN subscribers.

- **unknown**: Displays the data statistics for subscribers of unknown RAT type.

- **wlan**: Displays the data statistics for WLAN subscribers.

Statistics

The following table provides description of each field.

Table 65: show user-plane-service statistics rat all

Field	Description
Current Subscribers	
5G NR	Specifies the total number of current 5G NR subscribers.
EUTRAN	Specifies the total number of current EUTRAN subscribers.
WLAN	Specifies the total number of current WLAN subscribers.
Unknown	Specifies the total number of current subscribers of unknown RAT type.
Data Statistics	
5G NR	Specifies the data statistics for 5G NR subscribers.
Uplink	Specifies data statistics for 5G NR subscribers in uplink direction.
Total Pkts	Specifies the total number of uplink packets for 5G NR subscribers.
Total Bytes	Specifies the total number of uplink bytes for 5G NR subscribers.
Total Dropped Pkts	Specifies the total number of uplink packets dropped for 5G NR subscribers.
Total Dropped Bytes	Specifies the total number of uplink bytes dropped for 5G NR subscribers.
Downlink	Specifies data statistics for 5G NR subscribers in downlink direction.
Total Pkts	Specifies the total number of downlink packets for 5G NR subscribers.
Total Bytes	Specifies the total number of downlink bytes for 5G NR subscribers.
Total Dropped Pkts	Specifies the total number of downlink packets dropped for 5G NR subscribers.
Total Dropped Bytes	Specifies the total number of downlink bytes dropped for 5G NR subscribers.
EUTRAN	Specifies the data statistics for EUTRAN subscribers.
Uplink	Specifies data statistics for EUTRAN subscribers in uplink direction.

Field	Description
Total Pkts	Specifies the total number of uplink packets for EUTRAN subscribers.
Total Bytes	Specifies the total number of uplink bytes for EUTRAN subscribers.
Total Dropped Pkts	Specifies the total number of uplink packets dropped for EUTRAN subscribers.
Total Dropped Bytes	Specifies the total number of uplink bytes dropped for EUTRAN subscribers.
Downlink	Specifies data statistics for EUTRAN subscribers in downlink direction.
Total Pkts	Specifies the total number of downlink packets for EUTRAN subscribers.
Total Bytes	Specifies the total number of downlink bytes for EUTRAN subscribers.
Total Dropped Pkts	Specifies the total number of downlink packets dropped for EUTRAN subscribers.
Total Dropped Bytes	Specifies the total number of downlink bytes dropped for EUTRAN subscribers.
WLAN	Specifies the data statistics for WLAN subscribers.
Uplink	Specifies data statistics for WLAN subscribers in uplink direction.
Total Pkts	Specifies the total number of uplink packets for WLAN subscribers.
Total Bytes	Specifies the total number of uplink bytes for WLAN subscribers.
Total Dropped Pkts	Specifies the total number of uplink packets dropped for WLAN subscribers.
Total Dropped Bytes	Specifies the total number of uplink bytes dropped for WLAN subscribers.
Downlink	Specifies data statistics for WLAN subscribers in downlink direction.
Total Pkts	Specifies the total number of downlink packets for WLAN subscribers.
Total Bytes	Specifies the total number of downlink bytes for WLAN subscribers.
Total Dropped Pkts	Specifies the total number of downlink packets dropped for WLAN subscribers.
Total Dropped Bytes	Specifies the total number of downlink bytes dropped for WLAN subscribers.
Unknown	Specifies the data statistics for subscribers of unknown RAT type.
Uplink	Specifies data statistics for unknown RAT type subscribers in uplink direction.
Total Pkts	Specifies the total number of uplink packets for unknown RAT type subscribers.
Total Bytes	Specifies the total number of uplink bytes for unknown RAT type subscribers.

Field	Description
Total Dropped Pkts	Specifies the total number of uplink packets dropped for unknown RAT type subscribers.
Total Dropped Bytes	Specifies the total number of uplink bytes dropped for unknown RAT type subscribers.
Downlink	Specifies data statistics for unknown RAT type subscribers in downlink direction.
Total Pkts	Specifies the total number of downlink packets for unknown RAT type subscribers.
Total Bytes	Specifies the total number of downlink bytes for unknown RAT type subscribers.
Total Dropped Pkts	Specifies the total number of downlink packets dropped for unknown RAT type subscribers.
Total Dropped Bytes	Specifies the total number of downlink bytes dropped for unknown RAT type subscribers.

Bulk Statistics

The following bulk statistics are included in the User Plane Service schema to track RAT Type-based data statistics events.

Table 66: show bulkstats variables user-plane-service

Variable Name	Data Type	Key	Counter Type
vpnname	String	1	Info
vpnid	Int32	1	Info
servname	String	1	Info
servid	Int32	1	Info
curr-pdn-rat-eutran	Int64	0	Gauge
curr-pdn-rat-5g-nr	Int64	0	Gauge
curr-pdn-rat-wlan	Int64	0	Gauge
curr-pdn-rat-unknown	Int64	0	Gauge
uplink-total-pkts-pdn-rat-eutran	Int64	0	Counter
uplink-total-bytes-pdn-rat-eutran	Int64	0	Counter
uplink-total-pkts-dropped-pdn-rat-eutran	Int64	0	Counter
uplink-total-bytes-dropped-pdn-rat-eutran	Int64	0	Counter
downlink-total-pkts-pdn-rat-eutran	Int64	0	Counter

Variable Name	Data Type	Key	Counter Type
downlink-total-bytes-pdn-rat-eutran	Int64	0	Counter
downlink-total-pkts-dropped-pdn-rat-eutran	Int64	0	Counter
downlink-total-bytes-dropped-pdn-rat-eutran	Int64	0	Counter
uplink-total-pkts-pdn-rat-5g-nr	Int64	0	Counter
uplink-total-bytes-pdn-rat-5g-nr	Int64	0	Counter
uplink-total-pkts-dropped-pdn-rat-5g-nr	Int64	0	Counter
uplink-total-bytes-dropped-pdn-rat-5g-nr	Int64	0	Counter
downlink-total-pkts-pdn-rat-5g-nr	Int64	0	Counter
downlink-total-bytes-pdn-rat-5g-nr	Int64	0	Counter
downlink-total-pkts-dropped-pdn-rat-5g-nr	Int64	0	Counter
downlink-total-bytes-dropped-pdn-rat-5g-nr	Int64	0	Counter
uplink-total-pkts-pdn-rat-wlan	Int64	0	Counter
uplink-total-bytes-pdn-rat-wlan	Int64	0	Counter
uplink-total-pkts-dropped-pdn-rat-wlan	Int64	0	Counter
uplink-total-bytes-dropped-pdn-rat-wlan	Int64	0	Counter
downlink-total-pkts-pdn-rat-wlan	Int64	0	Counter
downlink-total-bytes-pdn-rat-wlan	Int64	0	Counter
downlink-total-pkts-dropped-pdn-rat-wlan	Int64	0	Counter
downlink-total-bytes-dropped-pdn-rat-wlan	Int64	0	Counter
uplink-total-pkts-pdn-rat-unknown	Int64	0	Counter
uplink-total-bytes-pdn-rat-unknown	Int64	0	Counter
uplink-total-pkts-dropped-pdn-rat-unknown	Int64	0	Counter
uplink-total-bytes-dropped-pdn-rat-unknown	Int64	0	Counter
downlink-total-pkts-pdn-rat-unknown	Int64	0	Counter
downlink-total-bytes-pdn-rat-unknown	Int64	0	Counter
downlink-total-pkts-dropped-pdn-rat-unknown	Int64	0	Counter
downlink-total-bytes-dropped-pdn-rat-unknown	Int64	0	Counter

Subscriber Level CLI

Use the following subscriber level CLIs as required:

- **show subscribers user-plane-only full all**
- **show subscribers user-plane-only { seid** *seid_value* **| callid** *callid_value* **} pdr all**
- **show subscribers user-plane-only { seid** *seid_value* **| callid** *callid_value* **} far all**
- **show subscribers user-plane-only { seid** *seid_value* **| callid** *callid_value* **} qer all**
- **show subscribers user-plane-only { seid** *seid_value* **| callid** *callid_value* **} urr all**
- **show subscribers user-plane-only { seid** *seid_value* **| callid** *callid_value* **} bar all**
- **show subscribers user-plane-only { seid** *seid_value* **| callid** *callid_value* **} pdr full all**
- **show subscribers user-plane-only { seid** *seid_value* **| callid** *callid_value* **} urr full all**
- **show subscribers user-plane-only { seid** *seid_value* **| callid** *callid_value* **} far full all**
- **show subscribers user-plane-only { seid** *seid_value* **| callid** *callid_value* **} qer full all**
- **show subscribers user-plane-only { seid** *seid_value* **| callid** *callid_value* **} bar full all**
- **show subscribers user-plane-only { seid** *seid_value* **| callid** *callid_value* **} pdr id** *id_value*
- **show subscribers user-plane-only { seid** *seid_value* **| callid** *callid_value* **} flows full**
- **show subscribers user-plane-only { seid** *seid_value* **| callid** *callid_value* **} bli full all**

VPP Statistics

To determine if the flows are offloaded to VPP, check for Fastpath statistics in the output of the following CLI commands:

- **show user-plane-service statistics all**
- **show user-plane-service statistics analyzer name ip [verbose]**
- **show user-plane-service statistics analyzer name ipv6 [verbose]**
- **show user-plane-service statistics analyzer name tcp [verbose]**
- **show user-plane-service statistics analyzer name udp [verbose]**
- **show user-plane-service statistics analyzer name http [verbose]**
- **show user-plane-service statistics analyzer name rtp [verbose]**
- **show subscribers user-plane-only full callid** *call_id*

SNMP Support

The system uses the Simple Network Management Protocol (SNMP) to send traps or events to the EMS server or an alarm server on the network. You must configure SNMP settings to communicate with those devices.

The *SNMP MIB Reference* describes the MIBs and SNMP traps supported by UPF and StarOS.

The following SNMP traps are available in support of their respective feature/functionality:

N4 Session/Node Level Reporting Procedure

The following traps are available to track status and conditions GTP-U path failure:

- EGTPUPathFailure: This trap is generated when no response is received for GTP-U ECHO requests and data path failure is detected towards peer EPC Node.

- EGTPUPathFailureClear: This trap is generated when the data path towards the peer node is available.

UP Session Recovery

The following traps are available after session recovery in the User Plane node:

- ManagerFailure: This trap is generated when there is failure in the Software manager.

- TaskFailed: This trap is generated when a noncritical task has failed and the appropriate recovery steps begin.

- TaskRestart: This trap is generated when a noncritical task has restarted after an earlier failure.

- SessMgrRecoveryComplete: This trap is generated when Session Manager recovery completes. This is typically caused by the failure of Session Manager task and successful completion of recovery.

- ManagerRestart: This trap is generated when the identified manager task has been restarted.

Sx Association

The following traps are available to track the status of an Sx Association:

- SxPeerAssociated: This trap is triggered when an Sx association is detected.

- SxPeerAssociationRelease: This trap is triggered when an Sx association release is detected.

URL Blacklisting

The following SNMP trap are available in support of URL Blacklisting feature:

- BLDBError: Specifies the blacklisting OPTBLDB file error displayed with an error code.

- BLDBErrorClear: Specifies the blacklisting OPTBLDB file error removed.

- BLDBUpgradeError: Specifies the blacklisting OPTBLDB file error displayed with an error code.

- BLDBUpgradeErrorClear: Specifies the Blacklisting OPTBLDB file error removed.

Enabling SNMP Traps

Use the following configuration to enable an SNMP trap.

```
configure
  snmp trap enable trap_name
  end
```

For supplemental information about SNMP Support, see *Management Settings* chapter in the *ASR 5500 System Administration Guide*.

Troubleshooting UPF Features

N4 or Datapath

The following CLI commands are available for troubleshooting N4 or datapath related issues:

- **show gtpu statistics**

- **show user-plane-service { all | bandwidth-policy | charging-action | edr-format | group-of-ruledefs | gtpp-group | name | pdn-instance | rulebase | ruledef | statistics | xheader-format }**

  NOTES:

  - **all**: Displays all User Plane services.

  - **bandwidth-policy**: Displays information for bandwidth-policy in User Plane service.

  - **charging-action**: Displays information for Charging actions in User Plane service.

  - **edr-format**: Displays information for EDR format in user Plane service.

  - **group-of-ruledefs**: Displays information on Group of Ruledefs configured in User Plane service.

  - **gtpp-group**: Displays information for bandwidth policy in User Plane service.

  - **name**: Displays information for specific User Plane service name.

  - **pdn-instance**: Displays information for PDN instance.

  - **rulebase**: Displays information for rulebase in User Plane service.

  - **ruledef**: Displays information for ruledef in User Plane service.

  - **statistics**: Displays node-level statistics for User Plane.

    Additionally, you can also use: **show user-plane-service statistics { all | analyzer | charging-action | fapi | rulebase | tethering-detection }**

  - **xheader-format**: Displays information for X-Header format in User Plane service.

- **show user-plane-service content-filtering category policy-id (all | id** *id_value* **}**

  - **content-filtering**: Displays content filtering information.

  - **category**: Displays content filtering category information.

  - **policy-id**: Displays content filtering category Policy-ID and its definition.

  - **all**: Displays definitions of all content filtering category policies.

- **id** *id_value*: Displays content filtering category definition of a particular Policy-ID. *id_value* is an integer ranging from 1 through 4,294,967,295.

- **show sx-service { all | name | statistics }**

  NOTES:

  - **all**: Displays all Sx Services.

  - **name**: Displays information for specific Sx Service name.

  - **statistics**: Displays the total of collected information for specific protocol since last restart or clear command.

Content Filtering

Use the following CLI command for troubleshooting CF related issues: **show user-plane-service inline-services { content-filtering | info | url-blacklisting }**

NOTES:

- **content-filtering**: Displays content filtering information.

- **info**: Displays information of inline services.

- **url-blacklisting**: Displays URL Blacklisting parameters in User Plane service.

URL Blacklisting

Use the following CLI command for troubleshooting URL Blacklisting related issues: **show user-plane-service url-blacklisting database { all | debug-only | facility | url }**

NOTES:

- **all**: Displays all URL Blacklisting database configurations.

- **debug-only**: Displays the URL Blacklisting static database debug information.

- **facility**: Displays URL Blacklisting database configuration per facility.

- **url**: Displays particular database information for URL Blacklisting.

UPF Sample Basic Configuration

Sample UPF Configuration

Sample Configuration

The following is only a sample configuration file provided solely for your reference. You must create and modify your own configuration file according to the specific needs of your deployment.

```
----- snip -----
  active-charging service acs
    bandwidth-policy BWP
      flow limit-for-bandwidth id 1 group-id 2
      flow limit-for-bandwidth id 2 group-id 3
      flow limit-for-bandwidth id 100 group-id 100
      group-id 2 direction uplink peak-data-rate 256000 peak-burst-size 1000 violate-action
 discard
      group-id 2 direction downlink peak-data-rate 200000 peak-burst-size 1000 violate-action
 discard
      group-id 3 direction downlink peak-data-rate 256000 peak-burst-size 1000 violate-action
 discard
      group-id 4 direction uplink peak-data-rate 300000 peak-burst-size 1200 violate-action
 lower-ip-precedence
      group-id 100 direction downlink peak-data-rate-kbps 4294967295 peak-burst-size
4294967295 violate-action discard

    ruledef L3_SERVER
      ip server-ip-address = 2.2.2.2/32
      tcp either-port = 80
    #exit
    ruledef L4_PORT
      tcp either-port = 80
      udp either-port = 80
      multi-line-or all-lines
    #exit
    ruledef L7_HTTP
      http host contains 2.2.2.2
      multi-line-or all-lines
    #exit
    ruledef http-pkts
      http any-match = TRUE
    #exit
    ruledef http-port
      tcp either-port = 80
      rule-application routing
    #exit
```

```
        ruledef ip-any-rule
          ip any-match = TRUE
        #exit
        urr-list urrs
          rating-group 10 urr-id 5
        #exit
        charging-action starent
          content-id 10
          billing-action egcdr
        #exit
        rulebase default
        #exit
        credit-control group default
          pending-traffic-treatment noquota buffer
          pending-traffic-treatment quota-exhausted pass
          usage-reporting quotas-to-report based-on-grant

        rulebase starent
          billing-records egcdr
          dynamic-rule order first-if-tied
          action priority 5 ruledef http-pkts charging-action standard
          action priority 10 ruledef L7_HTTP charging-action starent
          action priority 20 ruledef L4_PORT charging-action starent
          action priority 100 ruledef L3_SERVER charging-action starent
          action priority 10000 ruledef ip-any-rule charging-action starent
          route priority 1 ruledef http-port analyzer http
          egcdr threshold interval 1000
          bandwidth default-policy BWP
        #exit
        traffic-optimization-policy default
        #exit
    #exit
    context ingress
      interface N3_interface
        ip address 1.1.2.1 1.1.2.2
        ipv6 address abc0:0:0:cb::1/64 secondary
      #exit
      interface N3_interface_LOGICAL loopback
        ip address 1.1.2.1 1.1.2.2
      #exit
      interface N3_interface_LOGICAL2 loopback
        ip address 1.1.2.1 1.1.2.2
      #exit
      interface N4U_interface
        ip address 1.1.2.1 1.1.2.2
        ipv6 address abc0:0:0:cd::1/64 secondary
        ipv6 address abc0:0:0:ca::1/64 secondary
      #exit
      interface N4U_interface_LOGICAL loopback
        ip address 1.1.2.1 1.1.2.2
      #exit
     interface N4_interface
        ip address 1.1.2.1 1.1.2.2
        ipv6 address abc0:0:0:cc::1/64 secondary
      #exit
      interface N4_interface_LOGICAL loopback
        ip address 1.1.2.1 1.1.2.2
      #exit
      subscriber default
      exit
      aaa group default
      #exit
      gtpp group default
      #exit
```

```
              gtpu-service N3-GNB1
                bind ipv4-address 1.1.1.1
              exit
              gtpu-service N3-GNB2
                bind ipv4-address 1.1.2.1
              exit
              gtpu-service control_gtpu
                bind ipv4-address 1.1.2.1
              exit
              sx-service N4
                instance-type userplane
                bind ipv4-address 1.1.2.1
                sx-protocol heartbeat interval 3600
                sx-protocol heartbeat max-retransmissions 1
                sx-protocol association reattempt-timeout 30
              exit
              user-plane-service user-plane-service
                associate gtpu-service N3-GNB1 upf-ingress
                associate gtpu-service control_gtpu cp-tunnel
                associate sx-service N4
                associate fast-path service
                associate control-plane-group default
                load-control capacity 900
              exit
              user-plane-service user-plane-service1
              exit
              ip route 1.1.2.1 1.1.2.2 1.1.2.3 N4_interface
              ip route 1.1.2.1 1.1.2.2 1.1.2.3 N4_interface
              ip route 1.1.2.1 1.1.2.2 1.1.2.3 N3_interface
              ip route 1.1.2.1 1.1.2.2 1.1.2.3 N3_interface
              ip route 1.1.2.1 1.1.2.2 1.1.2.3 N4_interface
              ip route 1.1.2.1 1.1.2.2 1.1.2.3 N3_interface
              ip route 1.1.2.1 1.1.2.2 1.1.2.3 N4_interface
              ip route 1.1.2.1 1.1.2.2 1.1.2.3 N3_interface
          #exit
          context egress
            interface N6_interface
              ip address 1.1.2.1 1.1.2.2
              ipv6 address abc0:0:0:cf::1/64 secondary
            #exit
            subscriber default
            exit
            apn starent.com
              pdp-type ipv4 ipv6
              selection-mode subscribed sent-by-ms chosen-by-sgsn
              gtpp group default accounting-context egress
              ip context-name egress
              active-charging rulebase starent
            exit
            aaa group default
            #exit
            gtpp group default
              gtpp attribute local-record-sequence-number
              gtpp dictionary custom24
              gtpp egcdr service-data-flow threshold interval 60
              gtpp egcdr service-data-flow threshold volume downlink 13000
              gtpp egcdr service-data-flow threshold volume uplink 17000
              gtpp egcdr service-data-flow threshold volume total 22222
            #exit
            ipv6 route 2:2:2:2::/64 next-hop abc0::ab:1c:2ff:def9:1ad interface N6_interface
            ip route 1.1.2.1 1.1.2.2 1.1.2.3 N6_interface
            ip route 1.1.2.1 1.1.2.2 1.1.2.3 N6_interface
            ipv6 route 2:2:2:2::/64 next-hop abc0::ab:1c:2ff:def9:1ab interface N6_interface
            ip route 1.1.2.1 1.1.2.2 1.1.2.3 N6_interface
```

```
#exit
control-plane-group default
  sx-association initiated-by-cp
  peer-node-id ipv4-address 1.1.1.1 interface n4
#exit
user-plane-group default
#exit
port ethernet 1/11
  no shutdown
  vlan 203
    no shutdown
    bind interface N4U_interface ingress
  #exit
  vlan 204
    no shutdown
    bind interface N4_interface ingress
  #exit
  vlan 205
    no shutdown
    bind interface N3_interface ingress
  #exit
  vlan 206
    no shutdown
  #exit
  vlan 207
    no shutdown
    bind interface N6_interface egress
  #exit
#exit
end
```