Engineering
Simplicity

Junos Multi-Access User Plane User Guide

Published

2025-12-15

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Junos Multi-Access User Plane User Guide

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

Table of Contents

About This Guide

Use this guide to understand the Junos Multi-Access User Plane and how to configure an MX Series router as an SAEGW-U, SGW-U, PGW-U, and UPF to provide high-throughput 4G and 5G mobility and fixed wireless services.

1

CHAPTER

Understanding Junos Multi-Access User Plane

IN THIS CHAPTER

Junos Multi-Access User Plane Overview

IN THIS SECTION

Introduction

The 3rd Generation Partnership Project (3GPP) introduced the Evolved Packet Core (EPC) for core network architecture. As Figure 1 on page 2 shows, the four main EPC network elements are:

- Serving Gateway
- Packet Data Network (PDN) Gateway
- Mobility Management Entity (MME)
- Home Subscriber Server (HSS)

Figure 1: 3GPP Evolved Packet Core Architecture

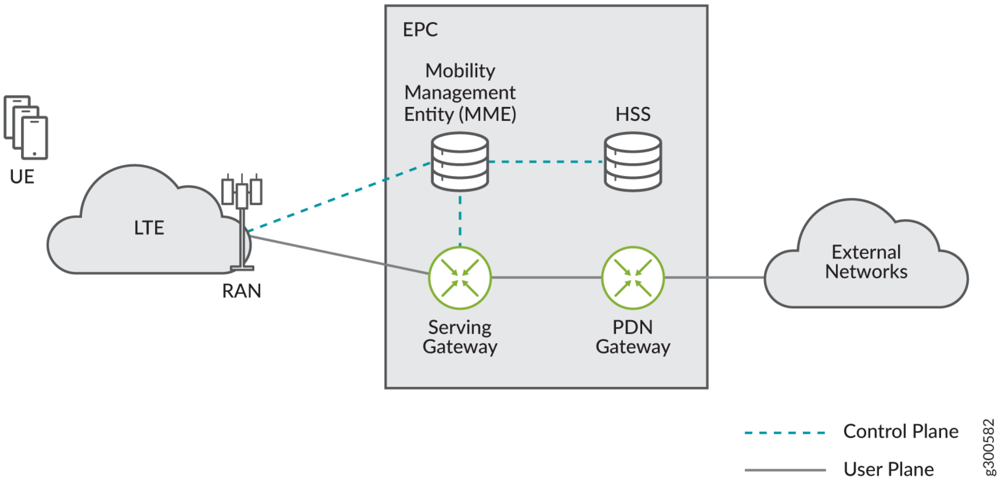

User Equipment (UE) has control path connectivity and data path connectivity to the EPC network elements over eNodeB base stations. The EPC provides data connectivity to external networks such as the Internet.

3GPP TS 29.244 Release 14 introduced CUPS, which stands for Control and User Plane Separation. CUPS provides the architecture enhancements for the separation of functionality in the EPC's serving gateway (SGW) and PDN gateway (PGW). As Figure 2 on page 3 shows, both the SGW and the PGW of the EPC can be separated into their control plane and user plane functions. CUPS introduces new interfaces, Sxa and Sxb, between the control plane and user plane functions of the SGW and PGW, respectively. CUPS enables control plane and user plane functions to be deployed, scaled and operated separately while integrated over a standard reference interface.

Figure 2: 3GPP Release 14 CUPS Architecture

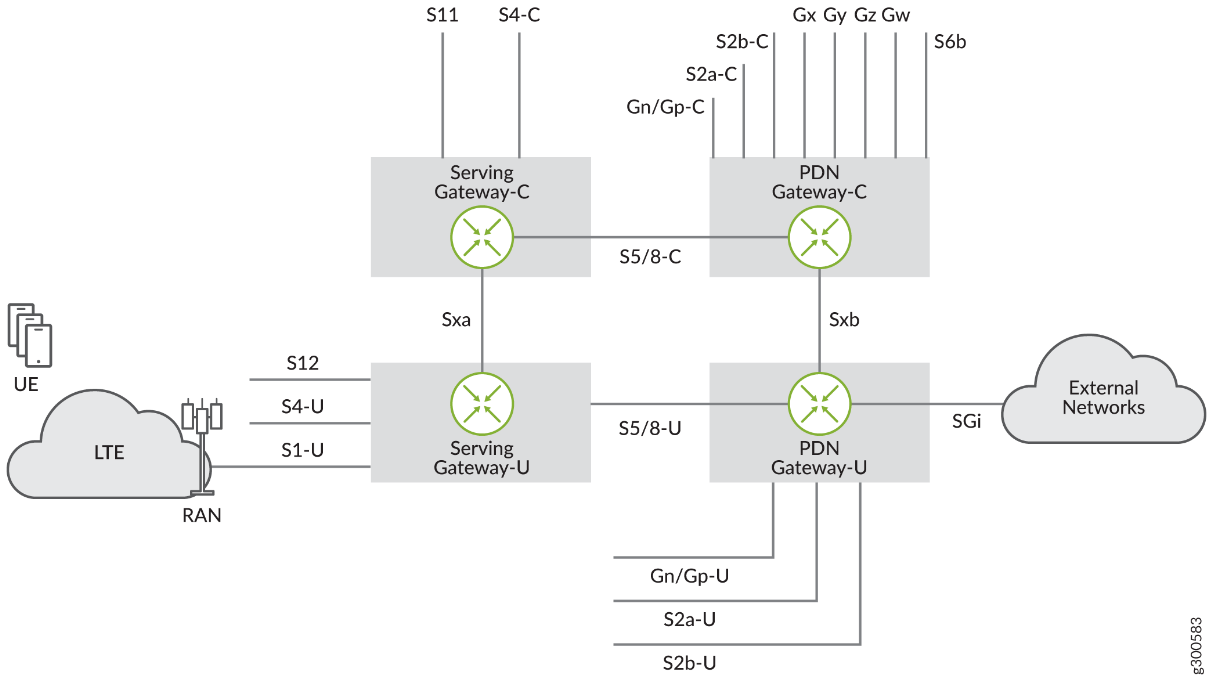

The control plane provides the following functionality:

- Receives traffic rules and actions
- Triggers accounting
- Makes session level announcements
- Receives usage information
- Receives user plane status information

- Northbound integration with the signaling plane
- Configures and enables Lawful Intercept sessions

The user plane provides the following functionality:

- Subscriber tunnel encapsulations (GTP-U)
- Packet routing and forwarding
- QoS and buffering
- Policy enforcement
- Statistics gathering and reporting
- Enacts Lawful Intercept requests
- Optional advanced services

With this functional separation, the control plane and the user plane have very distinct deployment requirements and can be in different physical locations. While the control plane function is very complex, the user plane function requires high packet processing capability and rich policy enforcement. You can distribute the user plane more than the control plane and locate the user plane closer to end-user access points. This distribution enables higher bandwidth per user while delivering lower latency. Control plane and user plane separation provides the following benefits:

- Independent scaling of the user plane and the control plane
- Network architecture flexibility including:
  - Ability to deploy from the edge to the core.
  - Ability to segregate different traffic types and services across different user planes while maintaining a common or single control plane.
- Operational flexibility
- Easier migration path from 4G to 5G services. CUPS is optional for 4G, but is an integral part of the 5G network architecture.

Junos Multi-Access User Plane supports a combined SGW user plane (SGW-U) and PGW user plane (PGW-U) in a single MX series router (see Figure 3 on page 5). The combined SGW-U/PGW-U is referred to as a SAEGW-U (System Architecture Evolution Gateway-User Plane). Juniper’s MX SAEGW-U can interoperate with a third-party combined SGW-C/PGW-C, referred to as a SAEGW-C, through a combined Sxab interface.

NOTE: Juniper's MX SAEGW-U communicates with the third-party SAEGW-C over the Sxab interface through Packet Forwarding Control Protocol (PFCP) as specified in 3GPP TS 29.244.

Figure 3: Junos Multi- Access User Plane SAEGW-U

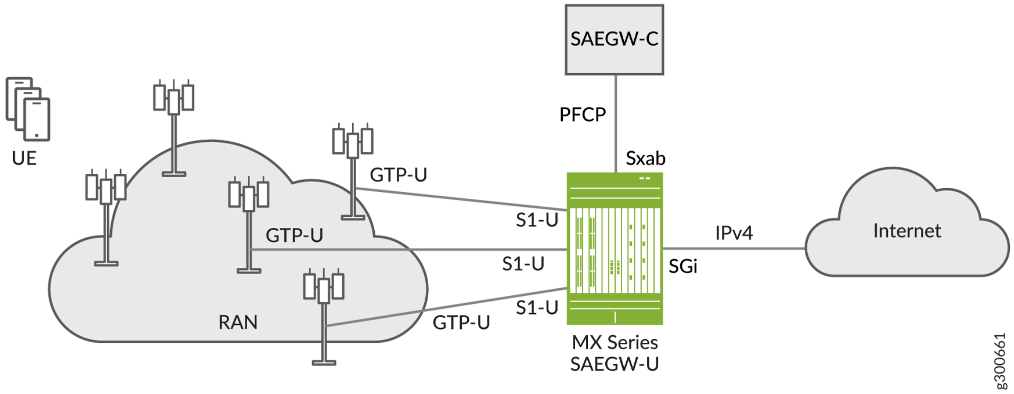

Junos Multi-Access User Plane also supports running an MX router as either a standalone SGW-U or a standalone PGW-U. A standalone SGW-U enables high-throughput 4G **mobility service** (relocation of a UE to a new eNodeB, new SGW-U, or new SAEGW-U). Junos Multi-Access User Plane support GTP-U based S5-U and S8-U interfaces, which are links between SGW-U and PGW-U devices. Junos Multi-Access User Plane also provides tunnel relay functionality to forward user plane traffic between S1-U and S5-U/S8-U interfaces and between S5-U/S8-U and SGi interfaces.

Figure 4 on page 6 shows the basic topology of running MX routers separately as and SGW-U and a PGW-U to enable mobility.

Figure 4: Junos Multi-Access User Plane SGW-U and PGW-U

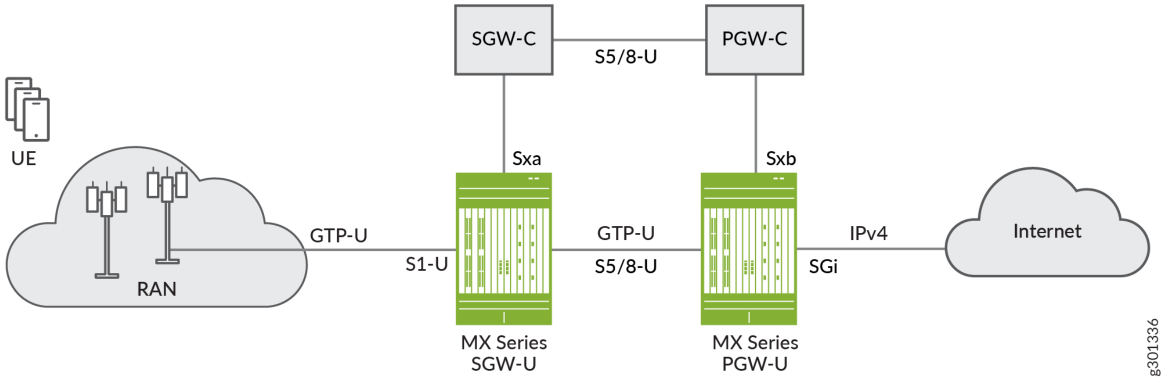

SGW-Cs and PGW-Cs handle logistics of UE handover, including SGW & PGW selection. The SGW-C and PGW-C participate in control protocol exchanges and update their SGW-U/PGW-U counterparts with any new or changed attributes of the UE session and bearers.

We support the following mobility scenarios:

- Handover with eNodeB and no SGW change
- Handover with SGW change (direct forwarding)
- Handover with SGW change (indirect forwarding)

Junos Multi-Access User Plane supports 5G user plane function (UPF) in addition to the SAEGW-U/SGW-U/PGW-U functions (see Figure 5 on page 7). Junos Multi-Access User Plane supports seamless transition from 4G to 5G services by supporting both networks on the same MX Series router with the same configuration. Junose Multi-Access User Plane supports 4G sessions and 5G sessions simultaneously.

Figure 5: Support for both 4G/LTE and 5G User Plane Functionality

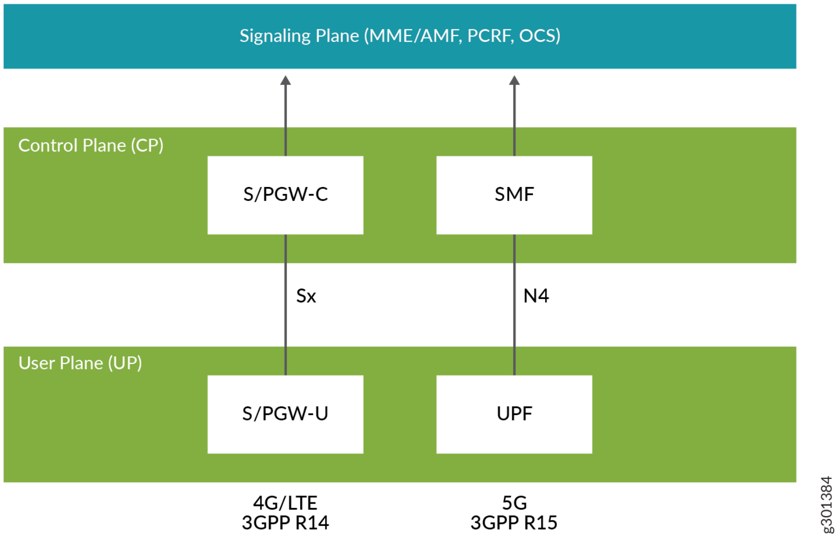

Junos Multi-Access User Plane supports MX routers functioning as user plane functions (UPFs) in accordance with 3GPP Release 15 CUPS architecture. The UPF provides high-throughput 5G fixed wireless and mobile wireless service in non-standalone (NSA) mode.

Figure 6 on page 7 shows the basic topology of running an MX router as a UPF to enable 5G services.

Figure 6: Junos Multi-Access UPF in 5G CUPS Architecture

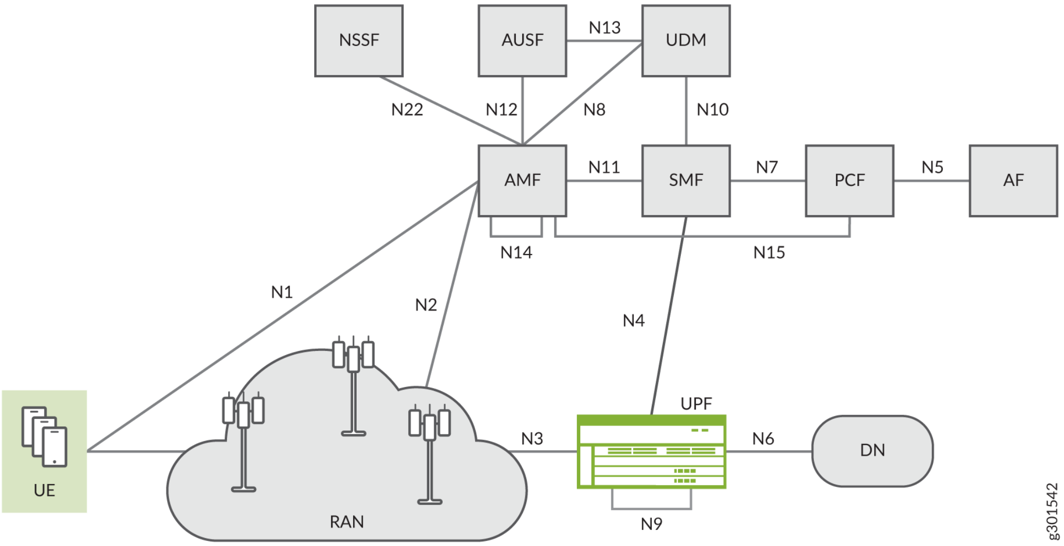

The 5G system architecture consists of the following network functions:

- Authentication Server Function (AUSF)
- Access and Mobility Management Function (AMF)
- Data Network (DN), e.g. operator services, Internet access or 3rd party services
- Network Slice Selection Function (NSSF)
- Policy Control Function (PCF)
- Session Management Function (SMF)
- Unified Data Management (UDM)
- User Plane Function (UPF)
- Application Function (AF)
- User Equipment (UE)
- (Radio) Access Network ((R)AN)

The session management function (SMF) includes the following functionality. A single instance of an SMF can support some or all of the SMF functionalities.

- Session management, e.g. session establishment, modify and release, including tunnel maintain between the UPF and a RAN node
- UE IP address allocation and management (including optional authorization)
- DHCPv4 (server and client) and DHCPv6 (server and client) functions
- Selection and control of the UPF
- Configure traffic steering at the UPF to route traffic to the proper destination
- Termination of interfaces towards policy control functions
- Charging data collection and support of charging interfaces
- Control and coordination of charging data collection at UPF
- Termination of SM parts of NAS messages
- Downlink data notification
- Initiator of RAN-specific SM information, sent via AMF over N2 to AN
- Determine SSC mode of a session

- Roaming functionality:
  - Handle local enforcement to apply QoS SLAs (VPLMN)
  - Charging data collection and charging interface (VPLMN)
  - Support for interaction with external DN for transport of signaling for PDU session authentication/authorization by external DN

The user plane function (UPF) includes the following functionality. A single instance of a UPF can support some or all of the UPF functionalities.

- Anchor point for Intra-/Inter-RAT mobility (when applicable)
- External PDU session point of interconnect to the data network
- Packet routing and forwarding
- Packet inspection
- User plane part of policy rule enforcement, e.g., gating, redirection, traffic steering)
- Traffic usage reporting
- QoS handling for user plane, e.g. uplink/downlink rate enforcement, reflective QoS marking in the downlink direction
- Uplink traffic verification (SDF to QoS Flow mapping)
- Transport level packet marking in the uplink and the downlink directions
- Downlink packet buffering and downlink data notification triggering
- Sending and forwarding of one or more end marker messages to the source RAN node

Junos Multi-Access User Plane acts as the UPF in the 5G CUPs architecture and includes support for the following:

- N3, N4, N6, and N9 interface support
- Roaming through the N9 interface
- GPRS tunneling protocol, user plane (GTP-U) tunneling to the control plane
- QoS Flow ID (QFI) support for 5G QoS flows

N3, N4, and N6 interfaces are similar to S1-U, Sx, and SGi interfaces in the 4G CUPs architecture, respectively. The N9 interface is similar to the S5/8-U interface. The N9 interface carries GTP-U encapsulated traffic and only connects from one UPF to another. In home-routed roaming scenarios, N9 reference points carry the user plane traffic back to an anchor UPF in the Home Public Land Mobile

Network (HPLMN). Junos Multi-Access User Plane supports either a single N9 reference point or a single N6 reference point per PDU session.

QoS in 4G networks is bearer-based where the mapping is one to one between a bearer and a radio bearer. QoS in 5G networks is flow-based where a QFI (QoS Flow Identifier) classifies and marks packets. Multiple QoS flows map to a radio bearer. Each QoS flow is associated with two parameters, a 5G QoS Identifier (5QI) and an allocation and retention priority (ARP).

In summary, starting with Junos OS Release 21.2R1, Junos Multi-Access User Plane supports four different modes of operation on a single MX router:

- **SGW-U**, where the MX router acts as an SGW-U for all sessions and connects to a third-party SGW-C over a single Sxa interface and Juniper or third party PGW-Us over multiple S5/8-U interfaces.
- **PGW-U**, where the MX router acts as a PGW-U for all sessions and connects to a third-party PGW-C over a single Sxb interface and Juniper or third-party SGW-Us over multiple S5/8-U interfaces.
- **Combined SGW/PGW-U (SAEGW-U)**, where depending on the UE location, the MX router acts as an SGW-U for some sessions, a PGW-U for another set of sessions and SAEGW-U for the remaining sessions. In this mode, the SAEGW-U connects to an SAEGW-C over a single Sxab interface and to other Juniper or third-party SGW-Us and PGW-Us over multiple S5/8-U interfaces.
- **UPF**, where the MX router acts as a UPF for all sessions and connects to a third-party SMF over a single N4 interface and to other Juniper or third-party UPFs over multiple N9 interfaces.

3GPP TS 29.244 Release 15 Support

Junos Multi-Access User Plane supports elements of 3GPP TS 29.244 Release 15, including support for the following functionality:

- **PDI Optimization Support**—Packet Detection Information (PDI) optimization is an optional feature that enables the control plane function (CPF) to optimize the signaling towards the UPF by combining the information that is common to multiple Packet Detection Rules (PDRs) as a Traffic Endpoint with a Traffic Endpoint ID (TEID) and then referring to this Traffic Endpoint in messaging. The Traffic Endpoint ID is unique within a PFCP session.
- **GTP Path Management**—GTP path management provides heartbeat and error indication over GTP-U interfaces. A GTP-U peer can send an echo request on a path to a GTP-U peer to find out if it is alive. Junos Multi-Access User Plane devices support responding to echo requests.
- **User ID Support**—The user ID is an information element (IE) that can be present in a PFCP Session Establishment Request. This IE is useful for troubleshooting problems in the UPF affecting a subscriber. The IE is visible in the output for the `show services mobile-edge sessions extensive` command. The user ID is an optional, noncritical IE that can be any length up to 16 digits or 8 characters.

- **Transport Level Marking**—For EPC, the SGW and PGW perform transport level marking on a per EPS bearer basis. Transport level marking is the process of marking traffic with a DSCP value based on the locally configured mapping from the QCI and optionally the ARP level. The CPF can change the transport level marking by changing the Transport Level Marking IE in the related Forwarding Action Rule (FAR).

  > **NOTE**: Juniper Multi-Access User Plane supports transport level marking per bearer for downlink data only.

  Transport Level Marking—For 5GC (5G core), transport level marking occurs on a per QoS flow basis. Transport level marking is the process of marking traffic at the UPF with a DSCP value based on the mapping from the 5QI, the priority level (if explicitly signaled) and, optionally, the ARP priority level configured at the SMF.

- **DDOS Support**—DDOS support is provided for PFCP and GTP path management. To configure DDOS for these protocols, see *protocols (DDoS)* .

- **QoS control/enforcement at the bearer level**—For QoS control/enforcement at the bearer level, the CPF must create the necessary PRDs to represent the service data flow, bearer, or session. The CPF must also create QERs for the QoS enforcement of the aggregate of the SDFs with the same bearer.

  Junos Multi-Access User Plane supports QoS enforcement at either the service data flow (SDF) or the bearer level. If the MX router as UPF receives more than one QER for a bearer, it enforces QoS at the SDF level. If the MX router as UPF receives one QER for a bearer, it enforces QoS at the bearer level.

Hardware and Software Requirements

This section lists the MX Series hardware and software requirements needed to implement Junos Multi-Access User Plane.

Table 1 on page 12 describes the hardware and software requirements for the Junos Multi-Access User Plane solution.

Table 1: Junos Multi-Access User Plane Platform Support

Junos OS Release	Supported Platforms	Line Cards Supporting Anchor PFE Interfaces	Line Cards Supporting Signaling, Ingress, and Egress Interfaces	Supported Routing Engines
Starting in Junos OS Release 19.4R1	• MX240 • MX480 • MX960	• MPC7	• MPC2 • MPC3 • MPC4 • MPC5 • MPC7	• RE-S-1800X4-32G-S • RE-S-X6-64G-S • RE-S-X6-128G
Starting in Junos OS Release 20.2R1	• MX204 • MX10003	• MX10003-LC2103	• MX10003-LC2103	
Starting in Junos OS Release 22.3R1	• MX10004	• MX10004-LC2101 • MX10004-LC480	• MX10004-LC2101 • MX10004-LC480	
Starting in Junos OS Release 23.2R1	• MX10008	• MX10008-LC480	• MX10008-LC480 • MPC10 series • LC9600	

NOTE: One MPC7 line card contains up to two anchor PFE interfaces.

NOTE: MX204 routers do not support GRES or APFE redundancy.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
23.2R1	Starting in Junos OS Release 23.2R1, Junos Multi-Access User Plane supports MX10008 platforms using LC480 line cards, including GRES and Anchor Packet Forwarding Engine redundancy in both 4G and 5G use cases.
22.3R1	Starting with Junos OS Release 22.3R1, Junos Multi-Access User Plane supports MX10004 platforms using LC2101 and LC480 line cards.
21.3R1	Starting in Junos OS Release 21.3R1, Junos Multi-Access User Plane provides a long-route implementation as a replacement for a filter-based implementation to steer traffic to the anchor Packet Forwarding Engine removing the need for a firewall filter to route GTP packets.
21.2R1	Starting in Junos OS Release 21.2R1, Junos Multi-Access User Plane supports routers functioning as user plane functions (UPFs) in accordance with 3GPP Release 15 CUPS architecture.
20.4R1	Starting with Junos OS Release 20.4R1, Junos Multi-Access User Plane supports running an MX router as either a standalone SGW-U or a standalone PGW-U.
20.4R1	Starting with Junos OS Release 20.4R1, Junos Multi-Access User Plane supports elements of 3GPP TS 29.244 Release 15.

CUPS Session Creation and Data Flow with Junos Multi-Access User Plane

IN THIS SECTION

With the introduction of CUPS, it's useful to illustrate how an end-user session is created, how data flows during the session, and how the session is terminated with Junos Multi-Access User Plane.

CUPS Session Creation

NOTE: Before a CUPS session can be created, the control plane function (SAEGW-C, SGW-C, PGW-C) must create an Sx association with the user plane function (SAEGW-U, SGW-U, PGW–U). The control plane sends an Sx Association Setup Request message and the user plane responds with an Sx Association Setup Response message to create the association. Once this is done, the control plane can create Sx sessions on the user plane.

When an end user wants to access the network, a CUPS session must be created. Figure 7 on page 15 illustrates this process once an Sx association is established between an SAEGW-C and an SAEGW-U.

Figure 7: CUPS Session Creation for SAEGW-C and SAEGW-U

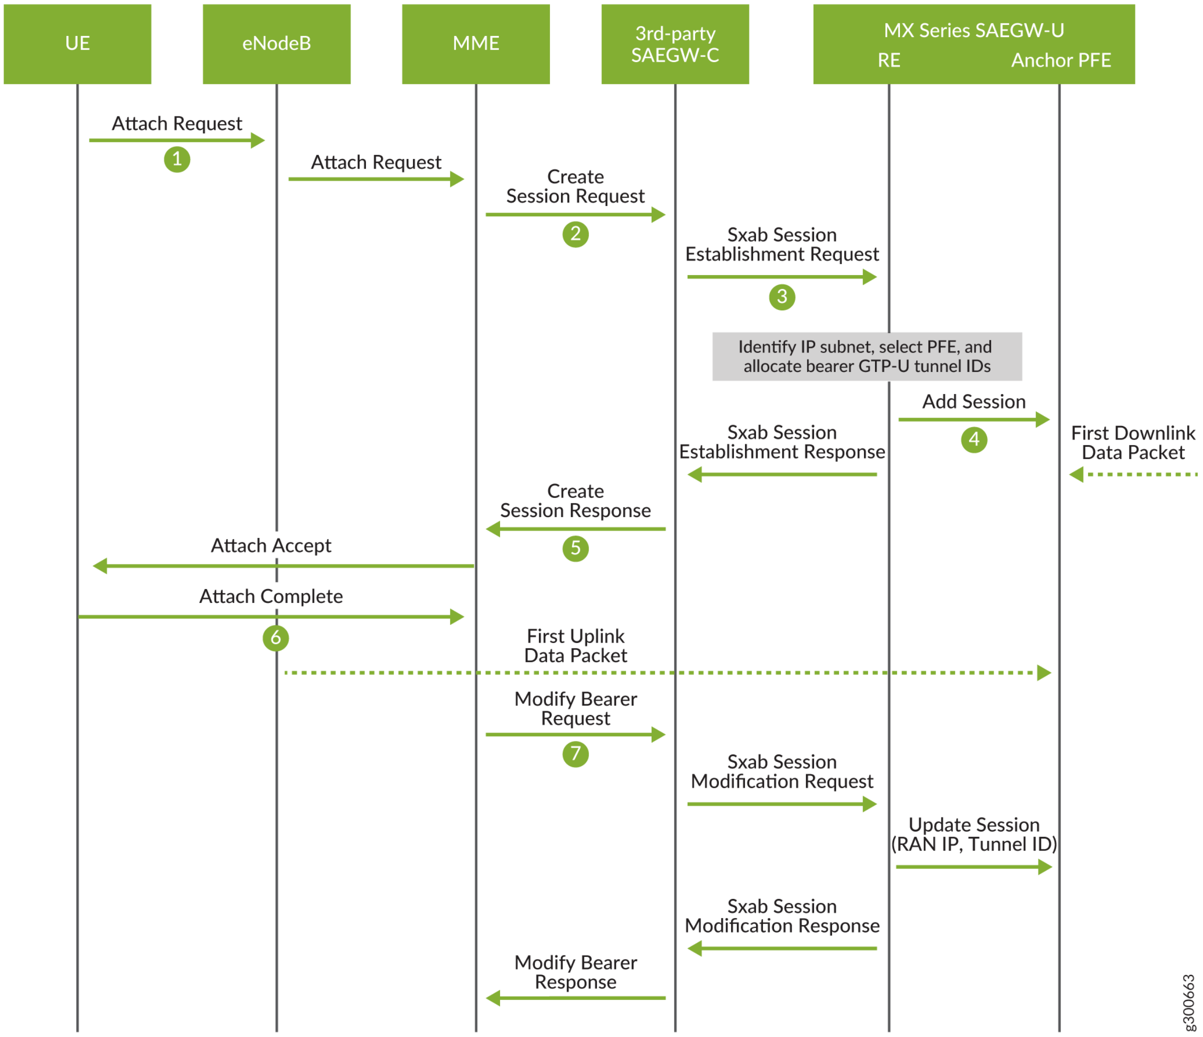
UE
eNodeB
MME
3rd-party SAEGW-C
MX Series SAEGW-U
RE
Anchor PFE
Attach Request
1
Attach Request
Create Session Request
2
Sxab Session Establishment Request
3
Identify IP subnet, select PFE, and allocate bearer GTP-U tunnel IDs
Add Session
4
First Downlink Data Packet
Sxab Session Establishment Response
Create Session Response
5
Attach Accept
Attach Complete
6
First Uplink Data Packet
Modify Bearer Request
7
Sxab Session Modification Request
Update Session (RAN IP, Tunnel ID)
Sxab Session Modification Response
Modify Bearer Response
g300663

1. The user equipment (UE) sends an Attach Request to the eNodeB, which forwards the message to the mobility management entity (MME). The request includes the *APN*.
2. The MME sends a Create Session Request to the SAEGW-C.
3. The SAEGW-C performs the following actions:
   - Validates information elements received in the request.
   - Validates the APN requested by the subscriber.
   - Sends a Sxab Session Establishment Request to the routing engine (RE) of the MX SAEGW-U.

NOTE: Sx session establishment is the SAEGW-C messaging the SAEGW-U control parameters on how to behave when the SAEGW-U encounters certain traffic. The minimum control parameters for Sx session establishment are one packet detection rule (PDR) and one forwarding action rule (FAR) The Sx session establishment effectively logs in the subscriber.

4. The RE of the SAEGW-U performs the following actions:
   - Identifies the IP address for the session.
   - Selects and anchor PFE to use for the session.
   - Allocates the bearer GTP-U tunnel IDs.
   - Adds the session to the anchor PFE.
   - Sends a Sxab Session Establishment Response back to the SAEGW-C.
5. The SAEGW-C sends a Create Session Response back to the MME.
6. The MME sends an Attach Accept message to the UE, which responds with an Attach Complete message.
7. The MME sends a Modify Bearer request to the SAEGW-C, which sends an Sxab Session Modification Request to the RE on the SAEGW-U. The RE updates the session IP address and tunnel ID of the eNodeB. Finally, a Modify Bearer Response is routed back to the MME.

NOTE: Sx Session Modification Request is the SAEGW-C messaging the SAEGW-U to modify any of the following four rules:

- Packet Detection Rule (PDR): contains information describing which packets should receive which treatment (for example, forwarding and other types of enforcement)
- Forwarding Action Rule (FAR): contains information on whether forwarding, dropping, or buffering is applied to a packet
- Usage Reporting Rule (URR): contains information that defines a certain measurement to make on user traffic and how that measurement shall be reported
- Quality Enforcement Rule (QER): contains information related to QoS enforcement of traffic

Junos Multi-Access User Plane does not support Buffering Action Rules (BARs).

8. The default bearer is now active and subscriber data traffic can pass back and forth between the UE through the eNodeB to the SAEGW-U and then the core network.

CUPS Session Data Flow

Once the session is established, the SAEGW-C is no longer directly involved for data flow. Data flows directly back and forth from the UE through the eNodeB to the SAEGW-U and then the core network. See Figure 8 on page 17.

Figure 8: CUPS Session Data Flow

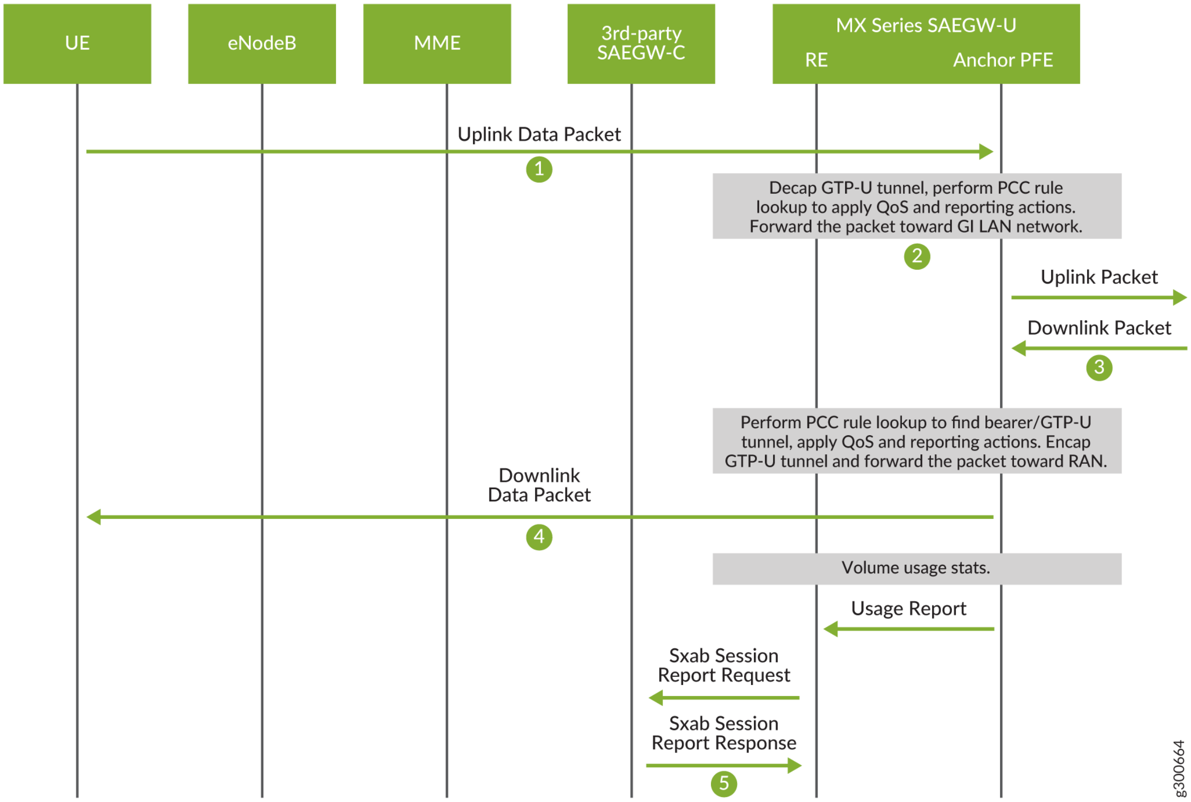
UE
eNodeB
MME
3rd-party SAEGW-C
MX Series SAEGW-U
RE
Anchor PFE
Uplink Data Packet
1
Decap GTP-U tunnel, perform PCC rule lookup to apply QoS and reporting actions. Forward the packet toward GI LAN network.
2
Uplink Packet
Downlink Packet
3
Perform PCC rule lookup to find bearer/GTP-U tunnel, apply QoS and reporting actions. Encap GTP-U tunnel and forward the packet toward RAN.
Downlink Data Packet
4
Volume usage stats.
Usage Report
Sxab Session Report Request
Sxab Session Report Response
5
g300664

1. The UE sends data to the eNodeB, which encodes the data as a GTP-U packet and forwards that packet to the anchor PFE on the SAEGW-U by way of the S1-U interface.
2. The anchor PFE of the SAEGW-U performs the following actions:
   - De-encapsulates the GTP-U packet.
   - Performs PCC rule lookup to apply QoS and reporting actions.
   - Forwards the de-encapsulated IPv4 packet to the core network over the SGi interface.
3. The SAEGW-U receives a downlink IPv4 packet from the core network.
4. The anchor PFE performs the following actions:

- Performs PCC rule lookup to determine the bearer GTP-U tunnel.
- Applies QoS and reporting actions.
- Encapsulates the IPv4 packet in GTP-U.
- Forwards the GTP-U packet to the eNodeB, which de-encapsulates the packet and forwards the data to the UE.

5. The SAEGW-U also creates a usage report for the session and sends the report to the SAEGW-C over the Sxab interface.

Charging and Usage Reports

Junos Multi Access User Plane supports charging and usage reports according to 3GPP TS 23.203, Policy and charging control architecture. Junos Multi Access User Plane supports the following usage reports:

- Volume threshold only
- Volume quota only
- Volume threshold and volume quota

Junos Multi Access User Plane uses the following process to generate usage reports:

1. The SAEGW-U creates a rating group for each bearer (default or dedicated). Routing groups can be created per session data flow (SDF) or for an entire bearer consisting of many SDFs.
2. The SAEGW-C associates a Usage Reporting Rule (URR) ID with a PDR and sends the URR ID over the Sx interface.
3. The SAEGW-U associates the URR ID with a rating group.
4. The SAEGW-C also messages what type of report needs to be generated for the URR ID (volume threshold only, volume quota only, volume threshold and quota).
5. The default action when the volume quota is reached is to drop all traffic for the session data flow.
6. When the subscriber session ends, the SAEGW-U generates and sends a final usage report to the SAEGW-C.

NOTE: The SAEGW-U supports pausing charging measurements for any URR ID where the SAEGW-C sets the Inactive Measurement flag of the Measurement Information IE of the URR. The SAEGW-U also supports sending immediate reports to the SAEGW-C on a URR query or remove request from the SAEGW-C; the SAEGW-U sends the usage report in the Modify Response.

Handover between eNodeBs and no SGW or SAEGW Change

Starting with Junos OS 20.4R1, Junos Multi Access User Plane supports UE mobility.

Figure 9 on page 20 shows the entire handover process when a UE switches from one eNodeB to another without requiring a SGW or SAEGW change (i.e., both eNodeBs are served by the same SGW). This is the simplest version of mobility handover.

Figure 9: Handover between eNodeBs

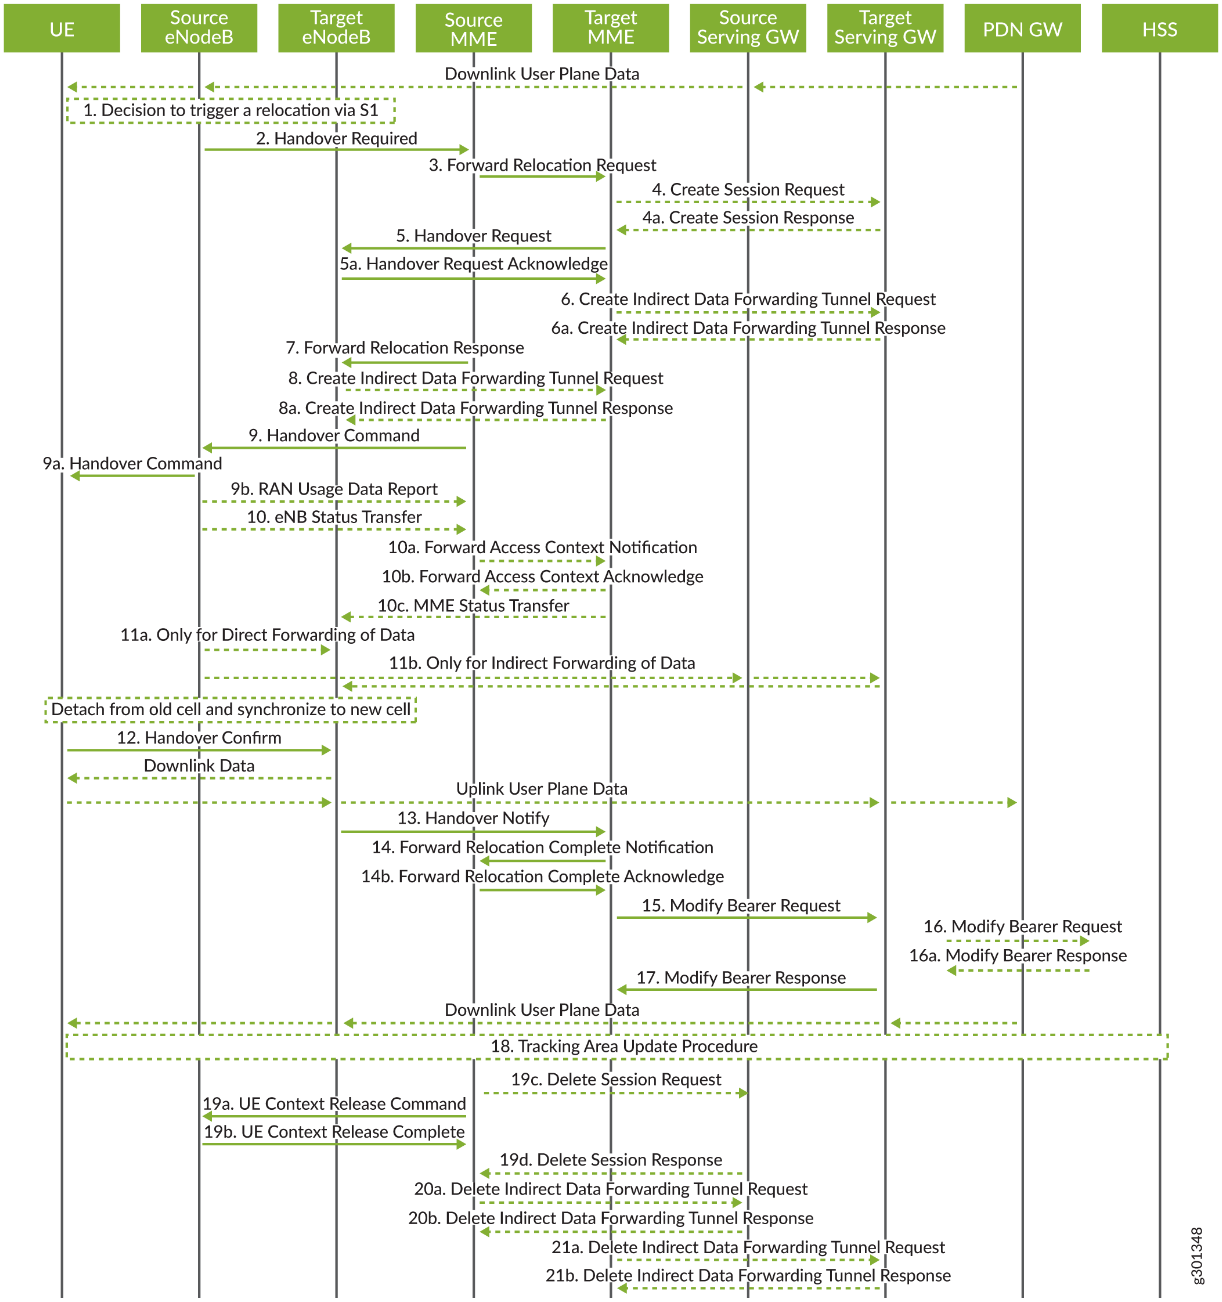
UE
Source eNodeB
Target eNodeB
Source MME
Target MME
Source Serving GW
Target Serving GW
PDN GW
HSS
Downlink User Plane Data
1. Decision to trigger a relocation via S1
2. Handover Required
3. Forward Relocation Request
4. Create Session Request
4a. Create Session Response
5. Handover Request
5a. Handover Request Acknowledge
6. Create Indirect Data Forwarding Tunnel Request
6a. Create Indirect Data Forwarding Tunnel Response
7. Forward Relocation Response
8. Create Indirect Data Forwarding Tunnel Request
8a. Create Indirect Data Forwarding Tunnel Response
9. Handover Command
9a. Handover Command
9b. RAN Usage Data Report
10. eNB Status Transfer
10a. Forward Access Context Notification
10b. Forward Access Context Acknowledge
10c. MME Status Transfer
11a. Only for Direct Forwarding of Data
11b. Only for Indirect Forwarding of Data
Detach from old cell and synchronize to new cell
12. Handover Confirm
Downlink Data
Uplink User Plane Data
13. Handover Notify
14. Forward Relocation Complete Notification
14b. Forward Relocation Complete Acknowledge
15. Modify Bearer Request
16. Modify Bearer Request
16a. Modify Bearer Response
17. Modify Bearer Response
Downlink User Plane Data
18. Tracking Area Update Procedure
19c. Delete Session Request
19a. UE Context Release Command
19b. UE Context Release Complete
19d. Delete Session Response
20a. Delete Indirect Data Forwarding Tunnel Request
20b. Delete Indirect Data Forwarding Tunnel Response
21a. Delete Indirect Data Forwarding Tunnel Request
21b. Delete Indirect Data Forwarding Tunnel Response
g301348

The following steps describe just the interactions between the control plane and user plane functions of the SGW and PGW (steps 15-17 in Figure 9 on page 20).

Step 15: Target MME to Target SGW Modify Bearer Request

1. SGW-C sends Sx Session Modification Request to MX SGW-U. The message includes F-TEIDu(s) corresponding to the new eNodeB. It may also instruct MX SGW-U to send “end marker” message towards the new eNodeB.

2. If requested to do so, MX SGW-U sends “end marker” message on S1-U interface towards the old eNodeB for all bearers referred to by Sx Session Modification Message.
3. MX SGW-U updates downlink peer F-TEID in the bearer(s) to F-TEIDu(s) received in the Sx Session Modification Request.
4. MX SGW-U sends Sx Session Modification Response to SGW-C

Step 16: Target SGW to PGW Modify Bearer Request

> **NOTE**: This step doesn’t affect any F-TEIDu assignments on any of the bearers. It may however update other forwarding & charging parameters based on the new location of the UE.

1. PGW-C sends Sx Session Modification Request to MX PGW-U.

2. MX PGW-U updates corresponding bearers and sends Sx Session Modification Response to PGW-C.

Handover with SGW Change

Considering the CUPS model, there are two types of procedures involving SGW change:

- **Type 1**: Only Create Session Request message is sent from MME/SGSN to SGW-C during SGW change.

- **Type 2**: Create Session Request message followed by Modify Bearer Request message is sent from MME/SGSN to SGW-C during SGW change.

For the MX SGW-U, the main difference between these two types is that in the first, the new SGW-C is provided with both eNodeB and PGW F-TEIDu(s) within Create Session Request, while in the second, the eNodeB’s F-TEIDu(s) are provided in the Modify Bearer Request, which translates to one extra Sx Session Modify Request/Response exchange between SGW-C and SGW-U. Because Type 1 can be considered a subset of Type 2, we present here the process for Type 2 handover.

Figure 9 on page 20 shows the entire handover process when a UE switches from one eNodeB to another with requiring a SGW change. The following steps describe just the interactions between the control plane and user plane functions of the SGW and PGW (steps 4,4a, 15-17 and 19 in Figure 9 on page 20).

Step 4: Target MME to Target SGW Create Session Request

1. The target SGW-C sends Sx Session Establishment Request to the target MX SGW-U. If PGW-U is a different physical node than the target SGW-U, the message includes F-TEIDu(s) of the PGW-U for every bearer of the session. It does not include local F-TEIDu(s) since MX SGW-U only supports UP function allocated local F-TEIDu.

2. The target MX SGW-U creates a new session and allocates local F-TEIDu(s) for all bearers indicated in Sx Session Establishment Request. If the message included PGW-U's F-TEIDs, we use them to set uplink peer F-TEIDu(s) for all referenced bearers.
3. The target MX SGW-U sends Sx Session Establishment Response message to the target SGW-C.

Step 15: Target MME to Target SGW Modify Bearer Request

1. The target SGW-C sends Sx Session Modification Request to the target MX SGW-U. The message includes F-TEIDu(s) for all bearers corresponding to the new eNodeB.

2. The target MX SGW-U updates downlink peer F-TEID in the bearer(s) to F-TEIDu(s) received in the Sx Session Modification Request.

3. MX SGW-U sends Sx Session Modification Response to SGW-C.

Step 16: Target SGW to PGW Modify Bearer Request

1. PGW-C sends Sx Session Modification Request to MX PGW-U. The message includes F-TEIDu(s) of the target SGW-U for all bearers. It may also instruct MX PGW-U to send “end marker” message.

2. If instructed to do so, MX PGW-U sends “end marker” message towards old SGW-U.

3. MX PGW-U updates downlink peer F-TEID for all the referenced bearers to F-TEIDu(s) received in the Sx Modification Request Message

4. MX PGW-U sends Sx Session Modification Response to the target SGW-C.

Step 19: Source MME to Source SGW Delete Session Request

1. Source SGW-C sends Sx Session Delete Request to the source MX SGW-U.

2. Source MX SGW-U deletes all bearers and the session.

3. Source MX SGW-U sends Sx Session Delete Response to the source SGW-C.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
20.4R1	Starting with Junos OS 20.4R1, Junos Multi Access User Plane supports UE mobility.

GRES on Junos Multi-Access User Plane

Graceful Routing Engine switchover (GRES) in Junos OS enables a router with redundant Routing Engines to continue forwarding packets even if one Routing Engine fails. GRES preserves interface and kernel information. Traffic is not interrupted.

NOTE: MX204 routers do not support GRES.

To preserve routing during a switchover, GRES must be combined with either:

- Graceful restart protocol extensions
- *Nonstop active routing* (NSR)

For Junos Multi-Access User Plane, GRES switchover protects the PFCP KeepAlive protocol. The new primary Routing Engine starts answering peer keepalives.

Any updates to the primary Routing Engine are replicated to the backup Routing Engine as soon as they occur.

Primary Role switches to the backup Routing Engine if:

- The primary Routing Engine kernel stops operating.
- The primary Routing Engine experiences a hardware failure.
- The administrator initiates a manual switchover.

NOTE: To quickly restore or to preserve routing protocol state information during a switchover, GRES must be combined with either graceful restart or nonstop active routing, respectively. For more information about graceful restart, see Graceful Restart Concepts. For more information about nonstop active routing, see Nonstop Active Routing Concepts.

If the backup Routing Engine does not receive a keepalive from the primary Routing Engine after 2 seconds, it determines that the primary Routing Engine has failed; and assumes primary role.

The Packet Forwarding Engine:

- Seamlessly disconnects from the old primary Routing Engine
- Reconnects to the new primary Routing Engine
- Does not reboot

- Does not interrupt traffic

The new primary Routing Engine and the Packet Forwarding Engine then become synchronized. If the new primary Routing Engine detects that the Packet Forwarding Engine state is not up to date, it resends state update messages.

NOTE: Successive Routing Engine switchover events must be a minimum of 240 seconds (4 minutes) apart after both Routing Engines have come up.

If the router or switch displays a warning message similar to `Standby Routing Engine is not ready for graceful switchover. Packet Forwarding Engines that are not ready for graceful switchover might be reset` , do not attempt switchover. If you choose to proceed with switchover, only the Packet Forwarding Engines that were not ready for graceful switchover are reset. None of the FPCs should spontaneously restart. We recommend that you wait until the warning no longer appears and then proceed with the switchover.

NOTE:

- We do *not* recommend performing a commit operation on the backup Routing Engine when GRES is enabled on the router or switch.
- We do *not* recommend enabling GRES on the backup Routing Engine in *any* scenario.

Figure 10 on page 25 shows the system architecture of graceful Routing Engine switchover and the process a routing platform follows to prepare for a switchover.

Figure 10: Preparing for a Graceful Routing Engine Switchover

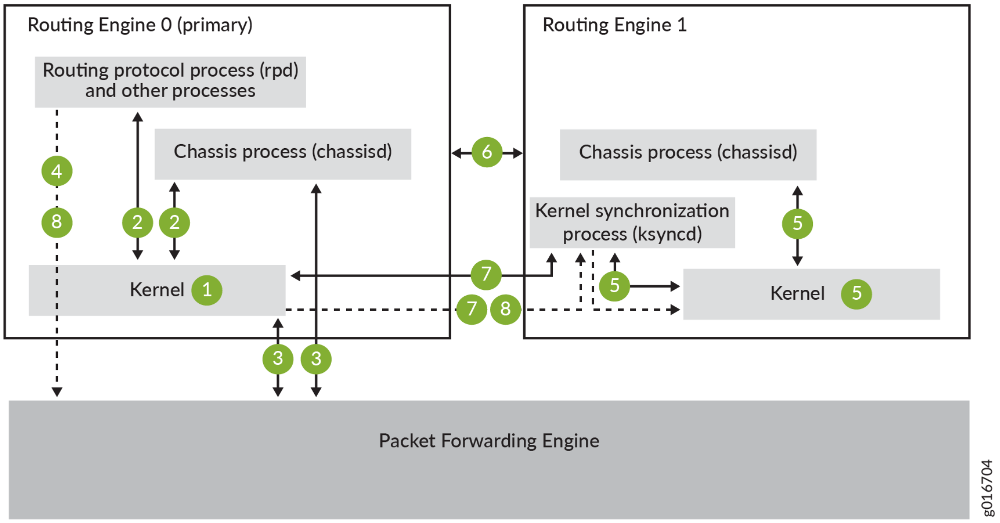
Routing Engine 0 (primary)
Routing protocol process (rpd) and other processes
Chassis process (chassisd)
Kernel
1
2
2
3
3
4
8
Routing Engine 1
Chassis process (chassisd)
Kernel synchronization process (ksyncd)
Kernel
5
5
5
6
7
7
8
Packet Forwarding Engine
g016704

NOTE: Check GRES readiness by executing both:

- The `request chassis routing-engine master switch check` command from the primary Routing Engine
- The `show system switchover` command from the Backup Routing Engine

The switchover preparation process for GRES is as follows:

1. The primary Routing Engine starts.
2. The routing platform processes (such as the chassis process [chassisd]) start.
3. The Packet Forwarding Engine starts and connects to the primary Routing Engine.
4. All state information is updated in the system.
5. The backup Routing Engine starts.
6. The system determines whether GRES has been enabled.
7. The kernel synchronization process (ksyncd) synchronizes the backup Routing Engine with the primary Routing Engine.
8. After ksyncd completes the synchronization, all state information and the forwarding table are updated.

Figure 11 on page 26 shows the effects of a switchover on the routing (or switching)platform.

Figure 11: Graceful Routing Engine Switchover Process

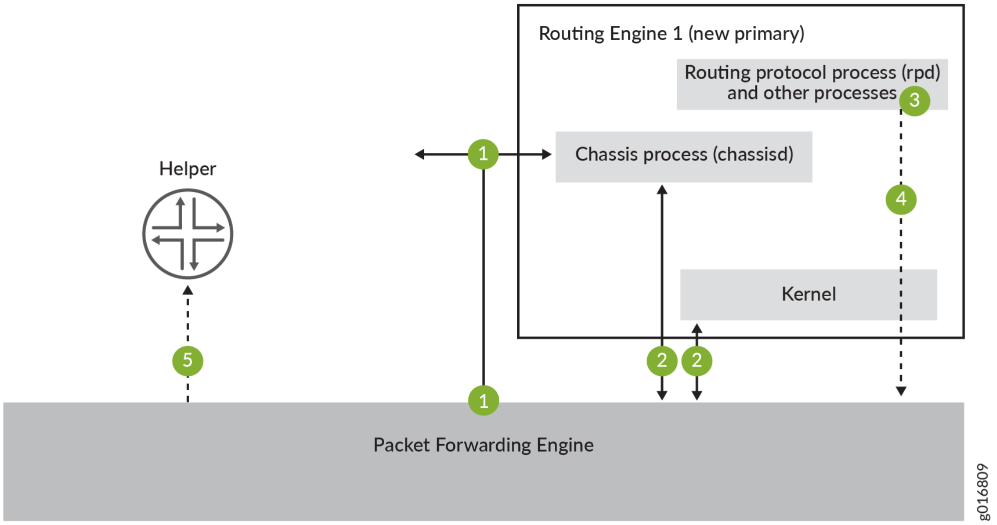
Routing Engine 1 (new primary)
Routing protocol process (rpd) and other processes
3
1
Chassis process (chassisd)
4
Helper
Kernel
5
2
2
1
Packet Forwarding Engine
g016809

A switchover process consists of the following steps:

1. When keepalives from the primary Routing Engine are lost, the system switches over gracefully to the backup Routing Engine.

2. The Packet Forwarding Engine connects to the backup Routing Engine, which becomes the new primary.

3. Routing platform processes that are not part of GRES (such as the routing protocol process rpd) restart.

4. State information learned from the point of the switchover is updated in the system.

5. If configured, graceful restart protocol extensions collect and restore routing information from neighboring peer *helper* routers.

NOTE: For MX Series routers using enhanced subscriber management, the new backup Routing Engine (the former primary Routing Engine) will reboot when a graceful Routing Engine switchover is performed. This cold restart resynchronizes the backup Routing Engine state with that of the new primary Routing Engine, preventing discrepancies in state that might have occurred during the switchover.

NOTE: In Junos Multi-Access User Plane configuration, if the mobile-edge configuration is committed and then GRES needs to be enabled or disabled, a reboot of the entire chassis is required.

NOTE: In Junos Multi-Access User Plane, any subscriber session whose Session State is *not* `ESTABLISHED`, a graceful restart logs out that subscriber and cleans up any state. The SAEGW-C will need to reestablish this session

Table 2: Effects of a Routing Engine Switchover

Feature	Benefits	Considerations
Dual Routing Engines only (no features enabled)	• When the switchover to the new primary Routing Engine is complete, routing convergence takes place and traffic is resumed.	• All physical interfaces are taken offline. • Packet Forwarding Engines restart. • The backup Routing Engine restarts the routing protocol process (rpd). • All hardware and interfaces are discovered by the new primary Routing Engine. • The switchover takes several minutes. • All of the router's adjacencies are aware of the physical (interface alarms) and routing (topology) changes.

Table 2: Effects of a Routing Engine Switchover *(Continued)*

Feature	Benefits	Considerations
GRES enabled	• During the switchover, interface, mobile-edge subscriber information, and kernel information is preserved. • The switchover is faster because the Packet Forwarding Engines are not restarted.	• The new primary Routing Engine restarts the routing protocol process (rpd). • All hardware and interfaces are acquired by a process that is similar to a warm restart. • All adjacencies are aware of the router's change in state. • Mobile-edge PFCP peer is not aware that GRES happened.
GRES *and* NSR enabled	• Traffic is not interrupted during the switchover. • Interface, mobile-edge subscriber information, and kernel information are preserved.	• Unsupported protocols must be refreshed using the normal recovery mechanisms inherent in each protocol. • Mobile-edge PFCP peer is not aware that GRES happened.

Table 2: Effects of a Routing Engine Switchover *(Continued)*

Feature	Benefits	Considerations
GRES *and* graceful restart enabled	- Traffic is not interrupted during the switchover. - Interface, mobile-edge subscriber information, and kernel information are preserved. - Graceful restart protocol extensions quickly collect and restore routing information from the neighboring routers.	- Neighbors are required to support graceful restart, and a wait interval is required. - The routing protocol process (rpd) restarts. - For certain protocols, a significant change in the network can cause graceful restart to stop. - Starting with Junos OS Release 12.2, if adjacencies between the restarting router and the neighboring peer 'helper' routers time out, graceful restart can stop and cause interruptions in traffic. - Mobile-edge PFCP peer is not aware that GRES happened.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
12.2	Starting with Junos OS Release 12.2, if adjacencies between the restarting router and the neighboring peer 'helper' routers time out, graceful restart can stop and cause interruptions in traffic.

RELATED DOCUMENTATION

Understanding Graceful Routing Engine Switchover

Configuring Graceful Routing Engine Switchover

Lawful Intercept on Junos Multi-Access User Plane

IN THIS SECTION

Lawful intercept is a process for obtaining communications network data related to a target individual or organization, as authorized by a judicial or administrative order.

Junos Multi-Access User Plane supports 4G and 5G compliance for lawful interception. You can lawfully intercept up to two percent of the maximum subscribers that the platform supports.

Lawful intercept on the Junos Multi-Access User Plane supports graceful Routing Engine switchover (GRES). It does not currently support unified ISSU.

Lawful Intercept on 4G Networks

The Junos Multi-Access User Plane supports 4G-compliant lawful intercept only over the System Architecture Evolution Gateway-User Plane (SAEGW-U).

Figure 12 on page 31 shows the entities involved in lawful interception on 4G networks and the interfaces between them.

Figure 12: Lawful intercept on 4G networks

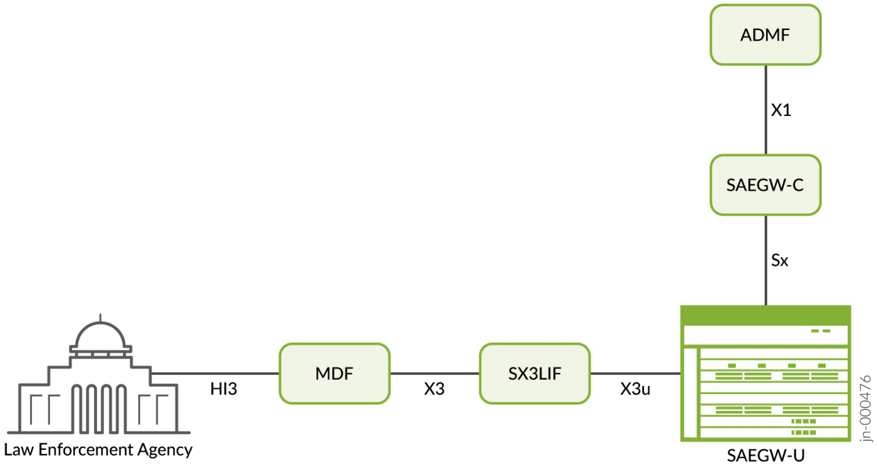

The process for lawful interception on 4G networks is as follows:

1. The Administrative Function (ADMF) performs the overall management of the lawful intercept system. It sends signals containing targeting information to the control plane over the X1 interface. The ADMF also informs the control plane which subscribers' traffic to duplicate and to send to the law enforcement agency.

2. The control plane (SAEGW-C) triggers the user plane over the Sx interface to start duplicating the subscriber traffic. It also sends the address of the Mediation and Delivery Function (MDF) to which the data is to be forwarded.

3. The Junos Multi-Access User Plane (SAEGW-U) sends the duplicated traffic to the split X3 lawful intercept internetworking function (SX3LIF) over the X3u interface. The SX3LIF then forwards it to the MDF over the X3 interface.

4. The MDF derives the intercept-related information from the LI_X3 headers sent along with the duplicated traffic. It then forwards the intercepted traffic to the law enforcement agency over the HI3 interface.

Lawful Intercept on 5G Networks

The Junos Multi-Access User Plane supports 5G-compliant lawful intercept only over the user plane function (UPF) or content of communication point of interception (CC-POI).

Figure 13 on page 32 shows the entities involved in lawful interception on 5G networks and the interfaces between them.

Figure 13: Lawful intercept on 5G networks

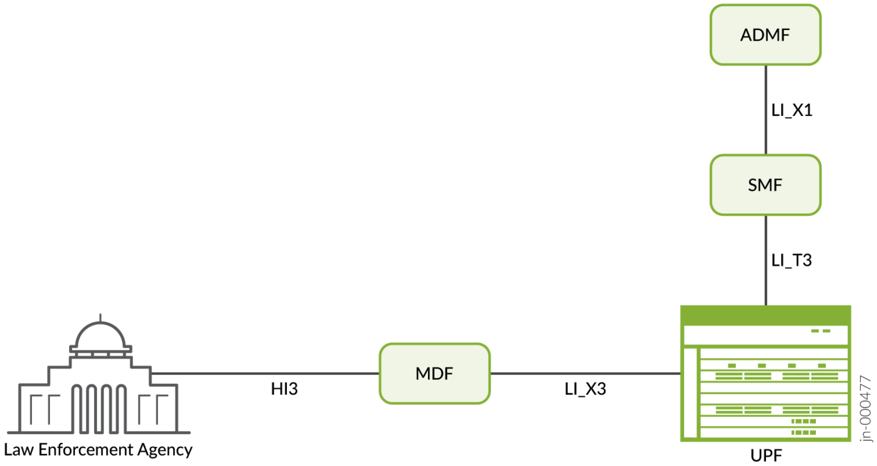

The process for lawful interception is as follows:

1. The Administrative Function (ADMF) performs the overall management of the lawful intercept system. It sends signals containing targeting information to the control plane over the LI_X1 interface. It also informs the control plane which subscribers' traffic to duplicate and to send to the law enforcement agency.

2. The control plane (SMF) triggers the user plane over the LI_T3 interface to start duplicating the subscriber traffic. It also sends the address of the Mediation and Delivery Function (MDF) to which the data is to be forwarded.

3. The Junos Multi-Access User Plane (UPF) sends the duplicated traffic to the MDF over the LI_X3 interface.

4. The MDF derives the intercept-related information from the LI_X3 headers sent along with the duplicated traffic. It then forwards the intercepted traffic to the law enforcement agency over the HI3 interface.

Requirements for Lawful Intercept

To configure lawful intercept for 4G networks, you must:

- Configure a loopback address on the Junos Multi-Access User Plane. This address is used as the source address for the lawfully intercepted traffic.

  ```
  [edit interfaces lo0 unit 0 family inet]
  user@host# set address loopback-address
  ```

To configure lawful intercept for 5G networks, you must:

- Set the loopback address to 127.0.0.1/32 on the Junos Multi-Access User Plane. This address is used as the source address for the lawfully intercepted traffic.

  ```
  [edit interfaces lo0 unit 0 family inet]
  user@host# set address 127.0.0.1/32
  ```

- Configure the HTTPS server on the Junos Multi-Access User Plane using the REST-API under the `[edit system services rest https-5g]` hierarchy.

  ```
  [edit system services rest https-5g]
  user@host# set port port-number
  user@host# set addresses ip-address
  user@host# set server-certificate local-certificate-identifier
  user@host# set mutual-authentication certificate-authority certificate-authority-profile-name
  ```

  For more information on the REST API, see *Configuring the REST API*.

  For information on certificate authority and public key infrastructure, see *PKI in Junos OS*.

- (Optional) Configure the connection limit for the REST-API. We recommend this configuration to enhance the performance of lawful intercept.

  ```
  [edit system services rest control]
  user@host# set connection-limit limit
  ```

Security Considerations for Lawful Intercept

Due to its nature, lawfully intercepted data must be transmitted in a secure manner. The 3GPP standards require that lawful intercept traffic on 5G networks must be sent to the Mediation and Delivery Function (MDF) using Transport Layer Security (TLS).

The Junos Multi-Access User Plane does not support TLS, but can send lawful intercept traffic to the MDF over IPsec tunnels. When using IPsec on the Junos Multi-Access User Plane, make sure that you:

- Configure the virtual multiservices interfaces and the next-hop-style service set.

  Use the following commands to configure the virtual multiservices interfaces:

  ```
  [edit interfaces]
  user@host# set interface-name unit logical-unit-number family inet
  user@host# set interface-name unit logical-unit-number service-domain inside
  user@host# set interface-name unit logical-unit-number family inet
  user@host# set interface-name unit logical-unit-number service-domain outside
  ```

  Use the following commands to configure a next-hop-style service set:

  ```
  [edit services]
  user@host# set service-set service-set-name
  [edit services service-set service-set-name]
  user@host# set next-hop-service inside-service-interface interface-name.unit-number outside-
  service-interface interface-name.unit-number
  ```

  The inside-service-interface must be a service interface logical unit that is configured with `service-domain inside` The outside-service-interface must be a service interface logical unit that is configured with `service-domain outside`.

- Configure the Internet Key Exchange (IKE).

  ```
  [edit security ike]
  user@host# set proposal proposal-name
  [edit security ike proposal proposal-name]
  user@host# set authentication-method pre-shared-keys
  user@host# set dh-group dh-group-number
  user@host# set authentication-algorithm algorithm
  user@host# set encryption-algorithm algorithm

  [edit security ike]
  user@host# set policy policy-name
  [edit security ike policy policy-name]
  user@host# set proposals proposal-name
  user@host# set pre-shared-key ascii-text key

  [edit security ike]
  ```

```
user@host# set gateway gateway-name
[edit security ike gateway gateway-name]
user@host# set ike-policy policy-name
user@host# set address ip-address
user@host# set external-interface interface-name
user@host# set local-address address
```

For more information on IKE, see *Internet Key Exchange (IKE) for IPsec VPN*.

- Use the IPsec tunnel mode.

NOTE: To use IPsec, you must have an MX-SPC3 services card installed on your device.

```
[edit security ipsec]
user@host# set proposal proposal-name
[edit security ipsec proposal proposal-name]
user@host# set protocol (ah | esp)
user@host# set authentication-algorithm algorithm
user@host# set encryption-algorithm algorithm

[edit security ipsec]
user@host# set policy policy-name
[edit security ipsec policy policy-name]
user@host# set proposals proposal-name

[edit security ipsec]
user@host# set vpn vpn-name
[edit security ipsec vpn vpn-name]
user@host# set bind-interface interface-name
user@host# set ike gateway gateway-name
user@host# set ike ipsec-policy policy-name
user@host# set establish-tunnels immediately
```

For more information on configuring IPsec tunnels, see *IPsec VPN Configuration Overview*.

If IPsec cannot be used, we recommend using a content of communication point of aggregation (CC-PAG) device, which encrypts the data using TLS before sending it to the MDF.

NOTE: You can still lawfully intercept traffic on 5G networks without configuring IPsec. However, we recommend that you use IPsec to comply with 3GPP standards.

The LI_T3 interface between the SMF and UPF supports TLS encryption. You can enable this by configuring a server certificate and mutual authentication on the local HTTPS server under the `[edit system services rest https-5g]` hierarchy.

Load and Overload Control on Junos Multi-Access User Plane

IN THIS SECTION

Load and Overload Control are a set of optional features within the 5G Core network (5GC). You can use these features to manage traffic loads on your network and maintain service to user equipment (UE) during extreme situations and peak traffic.

Load Control Information (LCI) and Overload Control Information (OCI) reports provide you with the information needed to manage your control plane traffic. LCI and OCI use the information collected from the user plane (UPF).

Load Control Information

Load Control allows the User Plane Function (UPF) to send specific traffic load information to the Control Plane Function (CPF) through the Sx/N4 interface. With this information, you are able to adaptively balance your session loads across various UPFs.

Load Control Information (LCI) reports gather specific information from the user plane:

- **CPU Usage:** CPU utilization on the current UPF.
- **Session Capacity:** The maximum number of sessions supported based on your system profile.

- **Memory Usage:** Maximum memory in use. Statistics for shared and heap memory, as well as the maximum usage.
- **Bandwidth Usage:** The amount of the available bandwidth in use on downlink and uplink paths.
- **Metric Calculation:** LCI is reported as a metric percentage from 0 to 100. If the metric reaches 90 percent, additional LCI reports generate for any change. If the metric reaches 95 percent, a full LCI and telemetry report generate each time the timer expires (every 15 seconds).

Overload Control Information

Overload Control allows the UPF to gracefully reduce its traffic load. It instructs peer control plane functions to reduce the amount of sessions being sent to the UPF. In this manner, overload control is able to avoid spreading an overload to other nodes in the network.

Like LCI, Overload Control Information (OCI) reports gather specific information from the UP:

- **UE Registration Surges:** A large number of UEs registering on the network will cause a surge of login sessions. OCI monitors the inflight number and reports an overload if the number of inflight sessions reach 80percent of the maximum.
- **UE Mobility and Application Signal:** A large number of UEs using specific applications can cause a lot of hand-over and QoS flow control signaling. This results in a surge of session modify messages. If the number of "session modify" messages exceeds a preset threshold, the UPF generates an overload report.
- **Packet Forwarding Engine Congestion Signal:** If the Packet Fowarding Engine receives more messages than it can handle, the Packet Forwarding Engine will send traffic in the Waiting-To-Send (WTS) queue. If any items are present in the WTS queue, an OCI report generates.
- **Anchor Packet Forwarding Engine Failover:** When an anchor Packet Forwarding Engine fails over and the secondary Packet Forwarding Engine boots, a 30 percent overload report is sent.
- **Overload Metric Report:** The OCI includes the metric of the load that needs to be reduced and how long the load reduction will last.

Maintenance Mode

Maintenance mode allows you to resolve an overload issue and perform other back-end network management tasks - such as a system upgrade. When you enter maintenance mode, the UPF sends a metric value of 100 percent to the CPF. Maintenance mode prevents new sessions from starting on the

UPF, but allows existing sessions and traffic to continue. When you exit maintenance mode, the true load metric is sent to the CPF and normal session activity resumes.

To enter and exit maintenance mode, you will use `service-mode` statements.

See *service-mode (mobile-edge SAEGW)* for additional information about using maintenance mode.

Mobile-Edge Configuration Commit Check

Whenever you modify or delete a configuration that supports mobile-edge features, Junos OS will check to see if any active sessions exist. If any active sessions exist, Junos OS provides an error message and rejects your modifications.

QoS in Junos Multi-Access User Plane

Quality of service (QoS) is a set of Junos OS performance features. You (the network administrator) can use these features to improve network performance by differentiating traffic into classes and then applying different behaviors to different types of traffic. These features also enable you to guarantee a certain level of network performance.

Usually, routers forward traffic using best-effort service without any control on throughput, packet loss, jitter, or delay. The QoS features provide differentiated services depending on the type of traffic. This feature improves performance in cases where the best-effort delivery is insufficient, such as in transmission of real-time audio and video.

Junos Multi-Access User Plane provides the following QoS features:

- **GBR forwarding queues**—Ensure that traffic on these queues has a minimum bit rate or guaranteed bit rate (GBR). We support eight preconfigured forwarding queues on the virtual routing and forwarding (VRF) loopback interface of the anchor Packet Forwarding Engine. Of the eight queues, four are guaranteed bit rate queues with three priorities each (high, medium, and low), and four are excess or best-effort queues with two priorities each (high and low).

> **NOTE**: The user plane sends all downlink traffic over the VRF interface using these eight forwarding queues. In the case of an intermediate UPF, which uses the VRF interface for both uplink and downlink traffic, the uplink traffic uses the best-effort queues. Otherwise, the uplink traffic does not use the VRF interface.

- **Forwarding queue mapping**—Uses the DiffServ code point (DSCP) value that the control plane provides in the Transport Level Marking to map the bearer or flow to the appropriate forwarding queue. For 4G, the user plane determines the DSCP value based on the QoS Class Identifier (QCI) and, optionally, the Admission Retention and Pre-emption Policy level. For 5G, the user plane determines the DSCP value based on the 5G QoS Identifier (5QI), the priority level, and, optionally, the Admission Retention and Pre-emption Policy level.

- **Bandwidth reservation**—Supports reserving a fraction of the bandwidth for queues with express traffic and guaranteed bit rate traffic in the downlink direction. The express queue handles high-priority, highly policed traffic that has a higher priority than guaranteed bit rate traffic. For example, the express queue handles IP Multimedia subsystem (IMS) signals, which have a higher priority than guaranteed bit rate traffic. Three queues—`gbr-high`, `gbr-medium`, and `gbr-low`—handle the guaranteed bit rate traffic.

- **Call admission control**—Tracks the total guaranteed bit rate bandwidth in the downlink direction across all bearers and flows assigned to each anchor Packet Forwarding Engine. This feature ensures that bearers and flows do not oversubscribe the reserved guaranteed bit rate bandwidth.

  The user plane admits a new session with guaranteed bit rate flow only if the least utilized anchor Packet Forwarding Engine can accommodate the requested guaranteed bit rate bandwidth. If an existing session is modified such that the new total guaranteed bandwidth exceeds the reserved bandwidth, the user plane accepts the modification and logs an oversubscription error message.

  The user plane always admits non-guaranteed bit rate flows as they are not offered any guaranteed bandwidth.

- **Policing**—Enables you to limit traffic in the uplink and downlink directions to a specified bandwidth. Guaranteed bit rate flows support policers for both guaranteed bit rate and maximum bit rate. Express flows and best-effort flows support only maximum bit rate policers.

  The user plane always transmits green packets (conforming to the guaranteed bit rate). Yellow packets exceed the guaranteed bit rate but conform to the maximum bit rate. The user plane drops these packets immediately for express queues, and drops them based on the overall load for guaranteed bit rate and best-effort queues. The user plane always drops red packets (exceeding the maximum bit rate).

RELATED DOCUMENTATION

Example: Configure GBR Forwarding on Junos Multi-Access User Plane | 74

2

CHAPTER

Configuring Junos Multi-Access User Plane

IN THIS CHAPTER

MX Series Router As Junos Multi-Access User Plane

IN THIS SECTION

Overview

Junos Multi-Access User Plane on a single MX Series router can function in four modes:

- As a combined SGW user plane (SGW-U) and PGW user plane (PGW-U) in a single MX series router. The combined SGW-U/PGW-U is referred to as a SAEGW-U (System Architecture Evolution Gateway-User Plane). The SAEGW-U interoperates with a third-party SAEGW-C through a combined Sxab interface.

- As a standalone SGW user plane (SGW-U). The SGW-U interoperates with a third-party SGW-C through a the Sxa interface and one or more Juniper or third-party PGW-Us over one or more S5/8-U interfaces.

- As a standalone PGW user plane (PGW-U) in a single MX router. The PGW-U interoperates with a third-party PGW-C through a the Sxb interface and one or more Juniper or third-party SGW-Us over one or more S5/8-U interfaces.

- As a standalone user plane function (UPF) for carrying 5G traffic. The UPF interoperates with a third-party session management function (SMF).

Configuring Junos Multi-Access User Plane on an MX Series router is essentially the same for each of these functions. This topic describes this configuration process.

As Figure 14 on page 42 shows, Juniper's MX SAEGW-U interoperates with a third-party SAEGW-C through a combined Sxab interface.

Figure 14: MX Series SAEGW-U in the CUPS Wireless Network Architecture

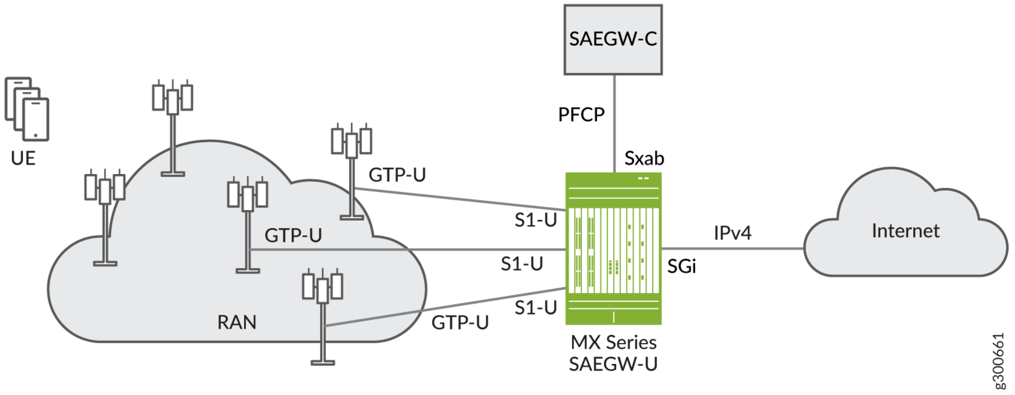

As Figure 15 on page 42 shows, the SGW-U interoperates with a third-party SGW-C through a the Sxa interface and one or more Juniper or third-party PGW-Us over one or more S5/8-U interfaces. The PGW-U interoperates with a third-party PGW-C through a the Sxb interface and one or more Juniper or third-party SGW-Us over one or more S5/8-U interfaces.

Figure 15: MX Series SGW-U and PGW-U in the CUPS Wireless Network Architecture

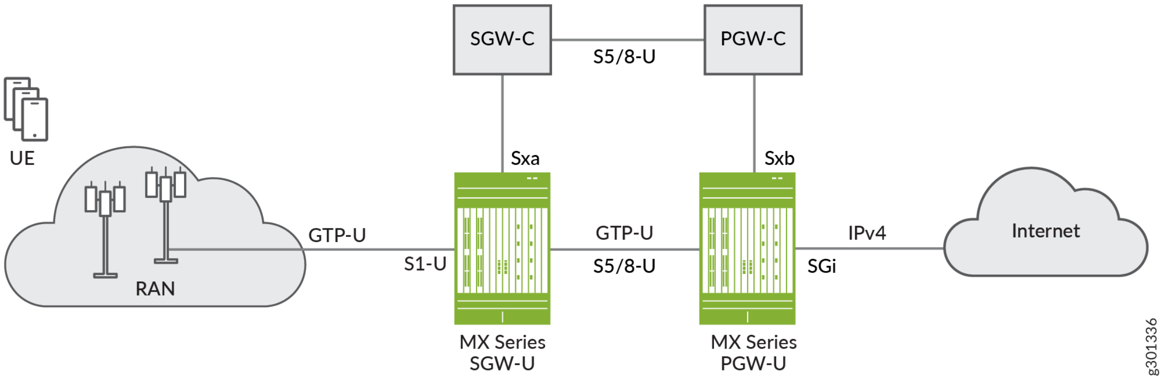

As Figure 16 on page 43 shows, Juniper's MX UPF interoperates with a third-party SMF through the N4 interface.

Figure 16: MX Series UPF in the 5G CUPS Wireless Network Architecture

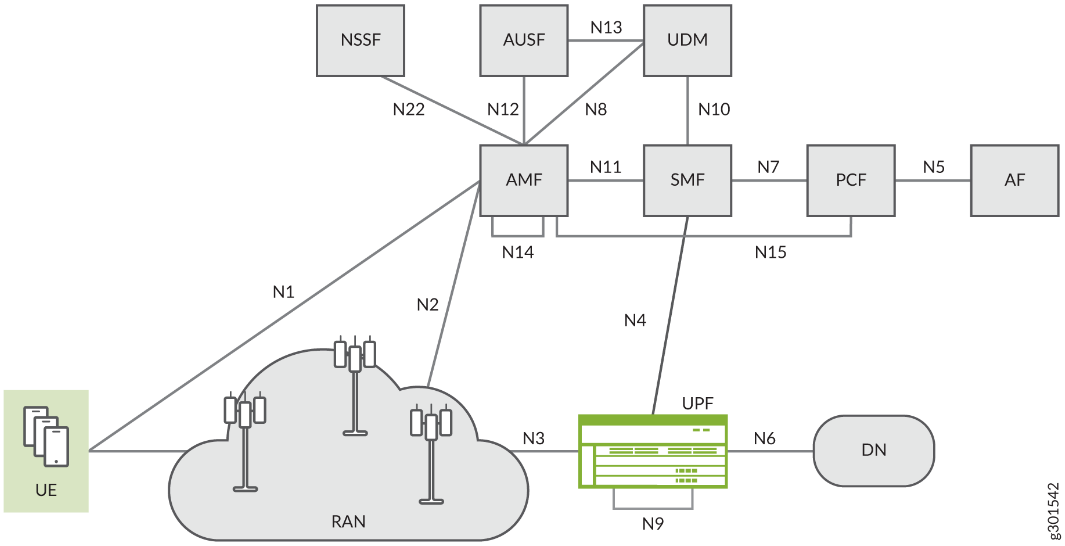

The MX as SAEGW-U, SGW-U, PGW-U or UPF supports the following CUPS interfaces:

- **Sxab/Sxa/Sxb/N4**—Packet Forwarding Control Protocol (PFCP) enables communication between the Junos Multi-Access User Plane and the control plane. PFCP encodes TLV messages for transport over UDP/IP. This interface can also transport user data packets (GTP-U based) between the user plane and control plane. Junos Multi-Access User Plane runs PFCP as the control protocol with the third-party control plane to set up data paths for wireless subscribers.

- **S1-U/N3**—This interface is the data path between an eNodeB and the Junos Multi-Access User Plane. Application data packets from end-user equipment are encapsulated over GTP. For upstream packets, Junos Multi-Access User Plane is responsible for GTP tunnel termination and forwarding the user packets to the core. For downstream packets from core, Junos Multi-Access User Plane adds the GTP header and forwards to eNodeB(s). The data plane handles IP packets encapsulated in GTP-U from/to eNodeBs that arrive for the mobile subscribers and performs routing to/from the external Internet.

- **S5/8-U**—The S5/8-U interface is the data path between an SGW-U and a PGW-U.

- **N9**—Interface between two UPFs.

- **SGi/N6**—Interface to the core Internet, supporting IPv4.

Junos Multi-Access User Plane provides purely the user plane function in the form of an MX router that interacts with a third-party control plane function. The MX router receives instructions from the control plane through the Sxab/Sxa/Sxb/N4 interface using PFCP. Based on those instructions, the MX routing engine manages user plane sessions and programs data paths in the anchor PFEs. For the MX router to provide the user plane functionality, it must contain the following minimum elements:

- **At least one anchor PFE interface**–An anchor PFE interface is a line card interface that has no physical interface connection, but rather provides the core processing of data traffic by doing the following:
    - Encoding/decoding of GTP-U packets. The anchor PFE interface decodes GTP-U packets from eNodeBs and forwards them to the core network and encodes IPv4 packets from the core network and forwards them to eNodeBs.
    - Enforces class of service and firewall filter rules on subscriber sessions
    - Collects statistics on data usage for charging/accounting purpose
- **At least one signalling/control interface**-This is the Sxab/Sxa/Sxb/N4 interface in the CUPS architecture. The signalling/control interface is a physical interface that does the following:
    - Sends/receives PFCP packets to/from the control plane
- **At least one ingress interface**-This is the S1-U or N3 or the S5/8-U or N9 interface in the CUPS architecture, depending on where the device is in the data stream. The ingress interface is a physical interface that does the following:
    - As the S1-U or N3 interface, forwards GTP-U packets between eNodeBs and the anchor PFE.
    - As the S5/8-U or N9 interface, receives GTP-U packets from the downstream UPF.
- **At least one egress interface**-This is the SGi or N6 or the S5/8-U or N9 interface in the CUPS architecture, depending on where the device is in the data stream. The egress interface is a physical interface that does the following:
    - As the SGi or N6 interface, forwards IPv4 packets between the anchor PFE and the core network.
    - As the S5/8-U or N9 interface, sends GTP-U packets to the upstream UPF.

NOTE: You can configure all interface types on the same line card, as long as that line card supports all of the interface types. See Table 1 on page 12 for a list of line card support by interface type.

Configuring Junos Multi-Access User Plane on an MX Router

IN THIS SECTION

As Figure 17 on page 45 shows, a standard setup of an MX router as either an SAEGW-U, an SGW-U, a PGW-U, or a UPF includes an ingress line card, and egress line card, and a recommended two anchor PFE line cards operating redundantly.

Figure 17: Standard setup for MX router as Junos Multi-Access User Plane

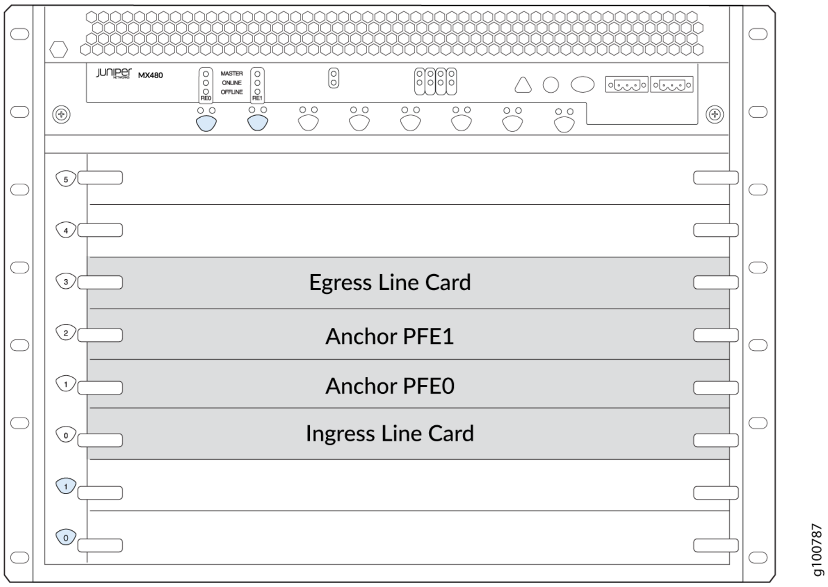

- The ingress line card provides the S1-U interface (SAEGW-U, SGW-U), the S5/8-U interface (PGW-U), or the N3 interface (UPF). , The ingress line card also provides the Sxab interface (SAEGW-U), the Sxa interface (SGW-U), the Sxb interface (PGW-U), or the N4 interface (UPF).

- The anchor PFE line cards provide the core processing of data traffic through internal `pfe-` interfaces. At least one anchor PFE card is required, but two are recommended to provide redundancy.

- The egress line card provides the SGi interface (SAEGW-U, PGW-U), the S5/8-U interface (SGW-U), or the N6 interface (UPF).

- You can configure all of this functionality on a single line card as long as that line card supports all of the Junos Multi-Access User Plane functionality. We show separate line cards here for simplicity and recommended setup.

To configure Junos Multi-Access User Plane on an MX router, perform the following configuration procedures in the listed order:

DDoS Attack Protection Configuration

Define DDoS attack protection for PFCP protocol traffic.

1. Configure protection for the PFCP protocol.

```
[edit system ddos-protection protocols]
user@host# set pfcp aggregate bandwidth packets-per-second
user@host# set pfcp aggregate burst size
user@host# set pfcp aggregate recover-time seconds
```

2. Configure GTP path management protection.

```
[edit system ddos-protection protocols]
user@host# set gtp-path-mgmt aggregate bandwidth packets-per-second
user@host# set gtp-path-mgmt aggregate burst size
user@host# set gtp-path-mgmt aggregate recover-time seconds
user@host# commit
```

GRES Configuration

The graceful Routing Engine switchover (GRES) feature in Junos OS enables a router with redundant Routing Engines to continue forwarding packets, even if one Routing Engine fails. GRES preserves interface and kernel information. Traffic is not interrupted.

Configure Graceful Restart (GRES).

```
[edit chassis]
user@host# set redundancy graceful-switchover
user@host# commit
```

SEE ALSO

GRES on Junos Multi-Access User Plane | 23

Chassis Configuration for the Anchor PFE Line Cards

Define each Packet Forwarding Engine (PFE) on each anchor PFE line card as an anchor interface.

1. Enable enhanced IP network services.

```
[edit chassis]
user@host# set network-services enhanced-ip
```

2. Configure slots for anchor PFE processing.

```
[edit chassis]
user@host# set fpc anchor-pfe0-slot pfe 0 forwarding-packages mobility user-plane
user@host# set fpc anchor-pfe0-slot pfe 1 forwarding-packages mobility user-plane
user@host# set fpc anchor-pfe1-slot pfe 0 forwarding-packages mobility user-plane
user@host# set fpc anchor-pfe1-slot pfe 1 forwarding-packages mobility user-plane
user@host# commit
```

Interface Configuration

Configure the interfaces needed.

1. Define the egress interface (SGi, S5/8-U, or N6). This interface is on the egress line card.

```
[edit interfaces]
user@host# set egress-interface-name unit 0 family inet address interface-address
```

2. Define the interface that connects to the control plane function (Sxab, Sxa, Sxb, or N4). This interface is on the ingress line card.

```
[edit interfaces]
user@host# set cpf-interface-name unit 0 family inet address interface-address
```

3. Define the ingress interface (S1-U, S5/8-U, N3). This interface is on the ingress line card.

```
[edit interfaces]
user@host# set ingress-interface-name unit 0 family inet address interface-address
```

4. Define the UPF local address and Mobile Edge interface.

```
[edit interfaces]
user@host# set lo0 unit 0 family inet address UPF-local-address
user@host# set mif unit 0 family inet
```

> **NOTE**: If you are connecting to multiple control planes, define the local address under the *control-plane-peers* stanza for each control plane function (CPF) rather than define a single loopback address.

> **NOTE**: `mif.0` is used in the default `inet.0` routing instance. Junos OS creates a default APN with `inet.0` as the routing instance. If you want to configure other routing instances, you must create `mif` interfaces with unit numbers other than `0`.

5. Assuming two anchor PFE linecards, each with two PFEs, define the anchor PFE interfaces.

```
[edit interfaces]
user@host# set apfe0 anchoring-options primary-list pfe-pfe0-slot/0/0
user@host# set apfe0 anchoring-options secondary pfe-pfe1-slot/0/0
user@host# set apfe1 anchoring-options primary-list pfe-pfe0-slot/1/0
user@host# set apfe1 anchoring-options secondary pfe-pfe1-slot/1/0
user@host# commit
```

NOTE: You cannot mix primary and secondary anchor PFEs on the same MPC. An MPC can have only either primary anchor PFEs or secondary anchor PFEs.

CAUTION: Changing the anchor PFE redundancy configuration once sessions are active kills all active sessions.

Mobile Edge Configuration

Once you’ve configured all of the necessary interfaces, you can configure the MX router to be a UPF.

1. Configure the connection to the control plane.

```
[edit services mobile-edge gateways saegw gateway-name control-plane-peers]
user@host# set local-address local-address
user@host# set apn-services apns apn-name mobile-interface mif.0
user@host# set peer-groups group-name path-management enable
user@host# set peer-groups group-name heartbeat-interval seconds
user@host# set peer-groups group-name n3-requests n3-requests
user@host# set peer-groups group-name t3-response seconds
user@host# set peer-groups group-name peer-address remote-peer-address
user@host# set peer-groups group-name peer-hostname remote-peer-hostname
```

NOTE: If you are connecting to multiple control planes, define the local address under the `control-plane-peers` stanza for each CPF. The loopback address, however, is still required for Lawful Intercept to function.

2. Configure the connection to the access network through the S1-U or N3 interface. if applicable.

```
[edit services mobile-edge gateways saegw gateway-name access-network-peers]
user@host# set local-address local-address
user@host# set peer-groups group-name peer-address remote-peer-address
user@host# set peer-groups group-name peer-hostname remote-peer-hostname
```

3. Configure the connection to the core network peers through the S5/8-U or N9 interface, if applicable.

```
[edit services mobile-edge gateways saegw gateway-name core-network-peers]
user@host# set local-address local-address
user@host# set routing-instance routing-instance-name
user@host# set path-management enable
user@host# set t3-response seconds
```

4. Define the interfaces that will provide the anchor PFE functionality.

```
[edit services mobile-edge gateways saegw gateway-name system]
user@host# set anchor-pfes interface apfe0
user@host# set anchor-pfes interface apfe1
user@host# commit
```

RELATED DOCUMENTATION

Anchor PFEs and Redundancy in Junos Multi-Access User Plane | **51**

Example: Configuring an MX Router as an SAEGW-U | **59**

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
23.1R1	Starting in Junos OS Release 23.1R1, Junos Multi-Access User Plane performs a commit check for active sessions when the mobile edge configuration is modified or deleted. If the check finds any active sessions, the commit will be prevented.
21.3R1	Starting in Junos OS Release 21.3R1, Junos Multi-Access User Plane provides a long-route implementation as a replacement for a filter-based implementation to steer traffic to the anchor Packet Forwarding Engine removing the need for a firewall filter to route GTP packets.

Anchor PFEs and Redundancy in Junos Multi-Access User Plane

IN THIS SECTION

Understanding the Anchor PFE

An anchor PFE is the *Packet Forwarding Engine* (PFE) on a standard line card that has no direct interface connections, but rather provides the core processing of data traffic by doing the following:

- Encoding/decoding of GTP-U packets. The anchor PFE decodes GTP-U packets from eNodeBs and forwards them to the core network and encodes IPv4 packets from the core network and forwards them to eNodeBs.
- Enforces class of service and firewall filter rules on subscriber sessions.
- Collects statistics on data usage for charging/accounting purpose.

Following are important points to consider when setting up anchor PFEs:

- You must configure at least one anchor PFE line card. We recommend at least two with 1:1 hot-standby redundancy.
- Each anchor PFE requires a defined `pfe-` interface of the form `pfe-x/y/z`.

Configuring No Redundancy for the Anchor PFEs

When no redundancy is required, all anchor PFE interfaces are equally available. The SAEGW-U uses all anchor PFE logical interfaces to anchor sessions/bearers. The routing engine (RE) of the SAEGW-U steers the GTP-U traffic for sessions and bearers to each of the anchor PFEs. The GTP processing for sessions and bearers and filter processing happens on the respective anchor PFE. The charging data is also maintained, collected and reported from the anchor for each session and its bearer.

When there is no redundancy configured, a failure of the anchor PFE line card is catastrophic for the SAEGW-U sessions/bearers in that control plane sessions corresponding to the failed anchor PFE and its data plane are lost. If supported by the SAEGW-C, the SAEGW-U can send an `Sx Session Set Deletion Request` for the lost sessions through the Sx interface, and the sessions are flushed in the SAEGW-C. All charging and other accounting data is lost for the sessions and bearers. New sessions can come up on the failed anchor PFE interface *only* when all sessions are flushed in the SAEGW-C for the failed anchor PFE, even if the anchor PFE comes up sooner. If other anchor PFE interfaces are available, new sessions can come up instantly on those anchor PFE interfaces.

Following is a typical configuration for two anchor PFEs with no redundancy.

1. Configure slot 1 and slot 2 for anchor processing.

```
[edit chassis]
user@host# set fpc 0 pfe 0 forwarding-packages mobility user-plane
user@host# set fpc 1 pfe 0 forwarding-packages mobility user-plane
```

2. Configure interfaces in slot 1 and slot 2 for PFCP processing.

```
[edit services mobile-edge gateways gateway-name system]
user@host# set anchor-pfes interface pfe-0/0/0
user@host# set anchor-pfes interface pfe-1/0/0
```

Configuring 1:1 Hot-standby Redundancy for the Anchor PFEs

To have 1:1 PFE redundancy, an aggregated anchor PFE group can be formed as below using exactly two PFE logical interfaces from different slots:

- Aggregated Anchor PFE group 1 – pfe-0/0/0 (primary), pfe-1/0/0 (secondary)
- Aggregated Anchor PFE group 2 – pfe-0/1/0 (primary), pfe-1/1/0 (secondary)

You cannot have primary and secondary anchor PFEs on the same line card. For example, the following combination is not supported:

- Aggregated anchor PFE group 1 – pfe-0/0/0 (primary), pfe-1/1/0 (secondary)
- Aggregated anchor PFE group 2 – pfe-1/0/0 (primary), pfe-0/1/0 (secondary)

We also do not recommended configuring anchor PFEs on two separate line cards with their secondary anchor PFEs on just one line card. For examle:

- Aggregated anchor PFE group 1 – pfe-0/0/0 (primary), pfe-2/0/0 (secondary)
- Aggregated anchor PFE group 2 – pfe-1/0/0 (primary), pfe-2/1/0 (secondary)

When aggregated anchor PFE configuration is used, both the primary anchor PFE and secondary anchor PFE have the session state. But the routing engine (RE) steers the GTP-U traffic for sessions and bearers only to the primary anchor PFE. The GTP processing for sessions and bearers and filter processing happens on the primary anchor PFE. The charging data is also maintained, collected and reported from the primary anchor PFE. The secondary is in hot-standby mode and is ready for takeover only in the event of primary anchor PFE failure.

Given the considerable load that a single anchor PFE linecard can need to handle, a single anchor PFE linecard is limited to a maximum of two redundancy groups. You can configure a single anchor PFE for one of the following roles:

- Dedicated primary for one redundancy group
- Dedicated secondary for one redundancy group
- Primary for two redundancy groups
- Secondary for two redundancy groups

When 1:1 redundancy is operational, the redundancy interface process monitors the health of the primary and secondary anchor PFEs.

A secondary anchor PFE failure results in zero data plane traffic loss on the primary anchor PFE. All active sessions remain unaffected. New sessions can come up without any latency. When the secondary anchor PFE is restored, there is a catchup phase to program the already active sessions and bearers in the secondary anchor PFE. After this is completed, new sessions are programmed in the secondary anchor PFE in parallel to the primary anchor PFE. From this point forward, the secondary anchor PFE can take over anytime.

If the primary anchor PFE fails, the secondary anchor PFE starts handling traffic. It might take a few seconds to detect the failure of the primary anchor PFE and for the RE to re-route the GTP-U traffic to the secondary PFE. This delay results in traffic loss during the anchor PFE switchover. Additionally, there is a loss of any charging data not reported by the primary anchor PFE before it failed. Anchor PFE switchover does not affect active sessions/bearers. In-flight changes to sessions and bearers as well as new sessions being created during anchor PFE switchover are rolled back. If supported by the SAEGW-C, the SAEGW-U can send an `Sx Session Set Deletion Request` for the lost sessions through the Sx interface, and the sessions are flushed in the SAEGW-C. After the anchor PFE switchover, the configured primary anchor PFE can be restored, starting as a secondary anchor PFE and going through catch-up similar to secondary APFE failure and restoration described above.

To configure redundancy for two anchor PFE line cards:

1. Configure PFE interfaces in each anchor PFE line card slot in aggregated anchor PFE configuration. For example:

```
[edit interfaces]
user@host# set apfe0 anchoring-options primary-list pfe-0/0/0
user@host# set apfe0 anchoring-options secondary pfe-1/0/0
user@host# set apfe1 anchoring-options primary-list pfe-0/1/0
user@host# set apfe1 anchoring-options secondary pfe-1/1/0
```

2. Reference the aggregated anchor PFE interfaces in the SAEGW-U configuration. For example:

```
[edit services mobile-edge gateways gateway-name system]
user@host# set anchor-pfes interface apfe0
user@host# set anchor-pfes interface apfe1
```

When the configured primary anchor PFE fails, the secondary anchor PFE takes over. When the failed primary anchor PFE recovers, it does not automatically resume primary status. It is now in secondary status until the configured secondary anchor PFE fails.

1. However, you can force the two anchor PFEs to revert to their configured state by setting a `revert-time`, in hours, under the `[edit interfaces aggregated-pfe-group anchoring-options]` hierarchy. For example:

```
[edit interfaces]
user@host# set apfe0 anchoring-options revert-time 2
user@host# set apfe1 anchoring-options revert-time 2
```

RELATED DOCUMENTATION

MX Series Router As Junos Multi-Access User Plane | 41

Downlink Forwarding Queues and Forwarding Classes

IN THIS SECTION

Downlink Forwarding Queues

Each anchor Packet Forwarding Engine is equipped with a set of 8 queues on the VRF Loopback interface (lo0). All mobile subscriber traffic on the anchor Packet Forwarding Engine traveling downlink will use these queues. These queues allow for service differentiation of your mobile subscriber traffic. The Downlink Forwarding Queues are preconfigured with default names and attributes, but you can manage the queues individually with one of these two commands:

```
[edit class-of-service forwarding-classes]
user@host# set queue queue # class-name
```

```
[edit class-of-service forwarding-classes]
user@host# set class class-name queue-num queue #
```

To view your queue details, use the command *show services mobile-edge anchor class-of-service*.

VRF Loopback Queue Assignment

A downlink forwarding queue/class is assigned to each bearer/service data flow (SDF). Assignment depends on whether your network has provided Transport Level Marking (TLM) within the forwarding action request (FAR).

- **TLM Provided:** If TLM is included with the FAR, then the TLM DSCP value maps mapped to a forwarding class that is using a configured classifier.
- **TLM Not Provided:** If TLM is not included with the FAR, the GBR-LOW queue/class is assigned to the Guaranteed Bit Rate (GBR). The BE-LOW-0 queue gets assigned to the non-GBR bearer/SDF.

  See *downlink-dscp-to-forwarding-class* for more information.

WAN Egress Forwarding Class Assignment (5G Only)

You can configure a WAN Egress Classifier for the User Plane Function (UPF). This classifier maps DSCP codes from the FAR's TLM to the forwarding class assigned to the packet. This classifier then selects the specific egress queue to use on the WAN interface. You use the `downlink-dscp-to-egress-forwarding-class` command to configure the WAN egress classifier.

The Egress Forwarding class gets assigned to the SDF. All packets matching that SDF inherit the forwarding class. If the Egress Forwarding Class is not assigned to the SDF, existing forwarding class assignments are unaffected.

See *downlink-dscp-to-egress-forwarding-class* for more information.

RELATED DOCUMENTATION

Understanding CoS Forwarding Classes

Defining CoS Forwarding Classes (CLI Procedure)

forwarding-classes

Session Maintenance and Optimization

IN THIS SECTION

Overview

You can optimize how Junos OS manages and maintains your sessions through the following options:

- Session Scaling Profiles
- Session Scale Monitoring
- Session Summaries and Viewing
- Clear Sessions
- Unused Route Removal
- Peer Group Routing Instances

Session Scaling Profiles

SAEGW allocates memory based on the session scaling profile selected in the configuration. If you do not select a session scale, then Junos OS will select the maximum available scale for your configuration. See Table 1 on page 57 for the available scaling options on your device. You can configure session scaling with the `mobile-edge session-scale` statement.

Table 3: Supported Session Scaling Profiles by Platform

Device	Supported Session Scales
MX204	• 50k • 100k

Table 3: Supported Session Scaling Profiles by Platform *(Continued)*

Device	Supported Session Scales
MX10003	• 100k • 250k • 500k
MX240/480/960/10004/10008	• 250k • 500k • 750k • 1m

Session Scale Monitoring

When your session scaling hits 80%, 90%, and 100% of the scale profile's capacity, Junos OS sends telemetry and ERRMSG notifications to you. When you hit 100% capacity, Junos OS rejects new sessions.

View Sessions and Session Summaries

You can view session entries from the internal table with the `transient-sessions` filter. This filter enables you to view all sessions and session IDs that are in a transient state. If a session stays in a transient state for several minutes, you can consider it a potentially stuck session.

You can view session summaries in both 5-minute and and hourly increments with the `show services mobile-edge sessions histogram` command.

Clear Sessions

You can clear any subsciber session from Junos OS with the session ID. This allows you to remove any subscriber sessions that get stuck in any state. You use the `clear services mobile-edge sessions` command to clear sessions.

Unused Route Removal

Junos OS deletes exact routes when the last session using that route is deleted. Junos OS will not advertise the deleted routes, and the deleted routes aren't visible to Junos OS users via the `show route` command.

Peer Group Routing Instances

You must designate a routing instance in Junos OS for each peer group on the same user plane function (UPF). This isolates control traffic between subscriber management functions (SMFs) that terminate on the same UPF.

RELATED DOCUMENTATION

show services mobile-edge sessions

mobile-edge session-scale

clear services mobile-edge sessions

Example: Configuring an MX Router as an SAEGW-U

IN THIS SECTION

This example shows how to configure an MX Series Router as an SAEGW-U for the Junos Multi-Access User Plane solution.

NOTE: This example is also valid for configuring an MX Series Router as a UPF for 5G sessions. Junos Multi-Access User Plane can support 4G and 5G sessions simultaneously.

Requirements

This example uses the following hardware and software components:

- MX480 (can also be MX240, MX960) router with:
  - Two MPC7s to act as anchor packet forwarding engines (PFEs) to handle GTP-U processing
  - Two MPC2s (can also be MPC3, MPC5, MPC7, MPC10) to act as ingress and egress PFEs
- Junos OS Release 21.3R1 or later

Figure 18 on page 61 below shows the hardware for this example.

Figure 18: Standard setup for MX Series router as SAEGW-U

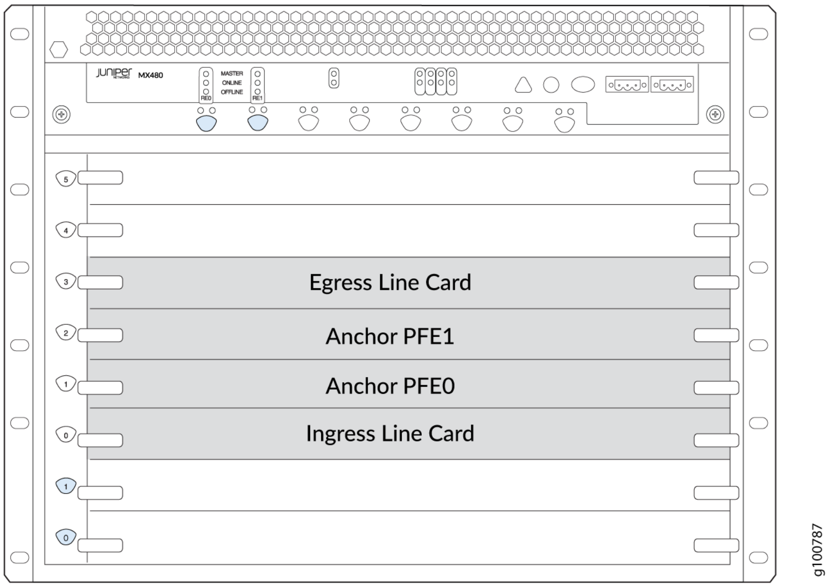

- The ingress line card (slot 0) provides the S1-U interface, connecting to the radio access network (RAN), and the combined Sxa/Sxb interface, connecting to the SAEGW-C.
- The anchor PFE line cards (slots 1 and 2) provide the core processing of data traffic through internal `pfe-` interfaces. At least one anchor PFE card is required, but two are recommended to provide redundancy.
- The egress line card (slot 3) provides the SGi interface, connecting to the core Internet.

Before you configure the MX Series Router as an SAEGW-U for the Junos Multi-Access User Plane solution, be sure you have:

- At least one configured SAEGW-C that you provide
- At lease one eNodeB
- Access to a packet data network (PDN)

Overview

IN THIS SECTION

Topology

In this example (see Figure 19 on page 62):

- An MPC2 is in slot 0 with ge-0/0/0.1 providing the combined Sxa/Sxb interface and ge-0/0/0.2 providing the S1-U interface.
- MPC7s are in slots 1 and 2 to provide the anchor PFE interfaces.
- And MPC2 is in slot 3 with ge-3/0/0.1 and ge-3/0/0.2 providing SGi interfaces.

Figure 19: Configuring an MX Router as an SAEGW-U

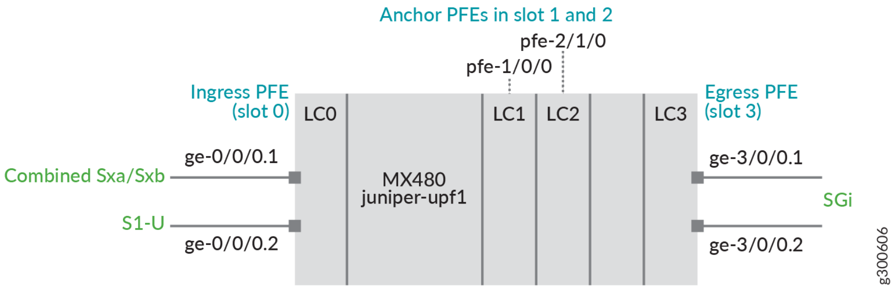

Configuration

IN THIS SECTION

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the `[edit]` hierarchy level.

```
set system ddos-protection protocols pfcp aggregate bandwidth 20000
set system ddos-protection protocols pfcp aggregate burst 9000
set system ddos-protection protocols pfcp aggregate recover-time 30
set system ddos-protection protocols gtp-path-mgmt aggregate bandwidth 8400
set system ddos-protection protocols gtp-path-mgmt aggregate burst 8400
set system ddos-protection protocols gtp-path-mgmt aggregate recover-time 30
set chassis redundancy graceful-switchover
set chassis fpc 1 pfe 0 forwarding-packages mobility user-plane
set chassis fpc 2 pfe 1 forwarding-packages mobility user-plane
set chassis network-services enhanced-ip
set interfaces ge-0/0/0 vlan-tagging
set interfaces ge-0/0/0 unit 1 vlan-id 101
set interfaces ge-0/0/0 unit 2 vlan-id 102
set interfaces ge-0/0/0 unit 1 family inet address 10.0.0.1/24
set interfaces ge-0/0/0 unit 2 family inet address 20.0.0.1/24

set interfaces ge-3/0/0 vlan-tagging
set interfaces ge-3/0/0 unit 1 vlan-id 101
set interfaces ge-3/0/0 unit 2 vlan-id 102
set interfaces ge-3/0/0 unit 1 family inet address 30.0.1.1/24
set interfaces ge-3/0/0 unit 2 family inet address 30.0.2.1/24
set interfaces lo0 unit 0 family inet address 100.0.0.1/32
set interfaces mif unit 0 family inet
set interfaces mif unit 1 family inet
set interfaces apfe0 anchoring-options primary-list pfe-1/0/0
set interfaces apfe0 anchoring-options secondary pfe-2/1/0
set services mobile-edge gateways saegw juniper-upf1 control-plane-peers local-address 10.0.0.1
set services mobile-edge gateways saegw juniper-upf1 control-plane-peers path-management enable
set services mobile-edge gateways saegw juniper-upf1 control-plane-peers heartbeat-interval 60
set services mobile-edge gateways saegw juniper-upf1 control-plane-peers apn-services apns apn-
default mobile-interface mif.0
set services mobile-edge gateways saegw juniper-upf1 control-plane-peers apn-services apns apn-
vrf1 mobile-interface mif.1
set services mobile-edge gateways saegw juniper-upf1 access-network-peers local-address 20.0.0.1
set services mobile-edge gateways saegw juniper-upf1 system anchor-pfes interface apfe0
set routing-instances vrf1 instance-type virtual-router
set routing-instances vrf1 interface mif.1
set routing-instances vrf1 interface ge-3/0/0.2
set routing-instances vrf1 routing-options static route 0.0.0.0/0 next-table inet.0
```

Procedure

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see Using the CLI Editor in Configuration Mode.

To configure the MX Router as an SAEGW-U:

1. Enable DDoS attack protection for PFCP protocol traffic.

```
[edit system ddos-protection protocols]
user@host# set pfcp aggregate bandwidth 20000
user@host# set pfcp aggregate burst 9000
user@host# set pfcp aggregate recover-time 30
user@host# set gtp-path-mgmt aggregate bandwidth 8400
user@host# set gtp-path-mgmt aggregate burst 8400
user@host# set gtp-path-mgmt aggregate recover-time 30
```

2. Configure Graceful Restart (GRES).

```
[edit chassis]
user@host# set redundancy graceful-switchover
```

3. Configure slot 1 & slot 2 for anchor PFE processing.

```
[edit chassis]
user@host# set fpc 1 pfe 0 forwarding-packages mobility user-plane
user@host# set fpc 2 pfe 1 forwarding-packages mobility user-plane
```

4. Enable enhanced IP network services.

```
[edit chassis]
user@host# set network-services enhanced-ip
```

5. Configure the ingress logical interfaces using vlans.

```
[edit interfaces ge-0/0/0]
user@host# set vlan-tagging
user@host# set unit 1 vlan-id 101
user@host# set unit 2 vlan-id 102
user@host# set unit 1 family inet address 10.0.0.1/24
user@host# set unit 2 family inet address 20.0.0.1/24
```

6. Configure the egress PFE for routing to core/ Internet for subscriber in VRF default (apn1).

```
[edit interfaces ge-3/0/0]
user@host# set vlan-tagging
user@host# set unit 1 vlan-id 101
user@host# set unit 2 vlan-id 102
user@host# set unit 1 family inet address 30.0.1.1/24
user@host# set unit 2 family inet address 30.0.2.1/24
```

7. Configure the loopback address and the mobile interface for subscriber VRFs.

```
[edit interfaces lo0]
user@host# set unit 0 family inet address 100.0.0.1/32
[edit interfaces mif]
user@host# set unit 0 family inet
user@host# set unit 1 family inet
```

8. Define the redundancy anchor PFE interfaces.

```
[edit interfaces]
user@host# set apfe0 anchoring-options primary-list pfe-1/0/0
user@host# set apfe0 anchoring-options secondary pfe-2/1/0
```

9. Name the SAEGW-U gateway `juniper-upf1` and configure the address where PFCP peers will connect to the SAEGW-U. Also, configure two APNs for SAEGW-U (`apn-default` to place sessions in the default routing instance and `apn-vrf1` for sessions into `VRF1`).

```
[edit services mobile-edge gateways]
user@host# set saegw juniper-upf1 control-plane-peers local-address 10.0.0.1
[edit services mobile-edge gateways saegw juniper-upf1 control-plane-peers]
user@host# set path-management enable
user@host# set heartbeat-interval 60
user@host# set apn-services apns apn-default mobile-interface mif.0
user@host# set apn-services apns apn-vrf1 mobile-interface mif.1
```

10. Configure the address where GTP-U peers will connect to the SAEGW-U.

NOTE: This is done at a different command hierarchy from the previous step.

```
[edit services mobile-edge gateways saegw juniper-upf1 access-network-peers]
user@host# set local-address 20.0.0.1
```

11. Configure aggregate interface `apfe0` for PFCP processing.

```
[edit services mobile-edge gateways saegw juniper-upf1 system]
user@host# set anchor-pfes interface apfe0
```

12. Configure the egress PFE for routing to core/ Internet for subscriber in VRF vrf1 (apn2).

```
[edit routing-instances vrf1]
user@host# set instance-type virtual-router
user@host# set interface mif.1
```

```
user@host# set interface ge-3/0/0.2
user@host# set routing-options static route 0.0.0.0/0 next-table inet.0
```

Results

From configuration mode, confirm your configuration by entering the `show chassis`, `show interfaces`, `show services`, `show routing-instances`, and `show unified-edge` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show system
ddos-protection {
    protocols {
        gtp-path-mgmt {
            aggregate {
                bandwidth 8400;
                burst 8400;
                recover-time 30;
            }
        }
        pfcp {
            aggregate {
                bandwidth 20000;
                burst 9000;
                recover-time 30;
            }
        }
    }
}
```

```
user@host# show chassis
redundancy {
    graceful-switchover;
}
fpc 1 {
    pfe 0 {
        forwarding-packages {
            mobility {
                user-plane;
            }
        }
    }
}
fpc 2 {
    pfe 1 {
        forwarding-packages {
            mobility {
                user-plane;
            }
        }
    }
}
network-services {
    enhanced-ip;
}
```

```
user@host# show interfaces
ge-0/0/0 {
    vlan-tagging {
        unit 1 {
            vlan-id 101;
        }
        unit 2 {
            vlan-id 102;
        }
    }
    unit 1 {
        family inet {
            address 10.0.0.1/24;
        }
    }
    unit 2 {
        family inet {            address 20.0.0.1/24;
        }
    }
}
ge-3/0/0 {
    vlan-tagging {
        unit 1 {
            vlan-id 101;
        }
        unit 2 {
            vlan-id 102;
        }
    }
    unit 1 {
        family inet {
            address 30.0.1.1/24;
        }
    }
    unit 2 {
        family inet {
            address 30.0.2.1/24;
        }
    }
}
apfe0 {
    anchoring-options {
        primary-list {
            pfe-1/0/0;
        }
        secondary pfe-2/1/0;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 100.0.0.1/32;
        }
    }
}
mif {
    unit 0 {
        family inet;
    }
    unit 1 {
        family inet;
    }
}
```

```
user@host# show services
mobile-edge {
    gateways {
        saegw juniper-upf1 {
            system {
                anchor-pfes {
                    interface apfe0;
                }
            }
            control-plane-peers {
                local-address 10.0.0.1;
                path-management enable;
                heartbeat-interval 60;
                apn-services {
                    apns apn-default {
                        mobile-interface mif.0;
                    }
                    apns apn-vrf1 {
                        mobile-interface mif.1;
                    }
                }
            }
            access-network-peers {
                local-address 20.0.0.1;
            }
        }
    }
}
```

If you are done configuring the device, enter `commit` from configuration mode.

Verification

IN THIS SECTION

Use various `show` commands to verify the SAEGW-U is functioning properly.

Verify SAEGW-U State

Purpose

Verify the SAEGW-U is running and that GRES is enabled.

Action

```
user@host> show services mobile-edge summary
Graceful-Restart        Enabled
Mastership              Master
State                   Running
Bulk Sync               Synchronized
```

Verify SAEGW-U Peers

Purpose

Verify the SAEGW-U has connected and is communicating with the SAEGW-Cs (control peers) and eNodeBs (access peers).

Action

```
user@host> show services mobile-edge peers statistics
Peers Summary:
    Total control peers: 1
    Total access peers:  1
    Total association setup request rejects:  0

Control Peer Statistics:
    IP address:         10.0.0.0
    Hostname:           saegw-c1
    Routing-Instance:   default

    Heartbeat Requests Received:            11
    Heartbeat Responses Sent:               11

    Heartbeat Requests Sent:                2
    Heartbeat Responses Received:           2
    Association Setup Requests Received:     1
    Association Setup Responses Sent:        1

    Association Release Requests Received:   0
    Association Release Responses Sent:      0

    Session Establishment Requests Received: 30000
    Session Establishment Responses Sent (Accepted):    30000
    Session Establishment Responses Sent (Rejected):    0

    Session Modification Requests Received:             30000
    Session Modification Responses Sent (Accepted):     30000
    Session Modification Responses Sent (Rejected):     0

    Session Deletion Requests Received:                 23169
    Session Deletion Responses Sent (Accepted):         22968
    Session Deletion Responses Sent (Rejected):         0

Access Peer Statistics:
    IP address:         20.0.0.0
    Routing-Instance:   default

    Echo Requests Received:                 0
    Echo Responses Sent:                    0
    Echo Requests Sent:                     0
    Echo Responses Received:                0
```

Verify SAEGW-U Sessions

Purpose

Verify the SAEGW-U has active data sessions.

Action

```
user@host> show services mobile-edge sessions summary
Sessions by State:
    SESSION_WAIT: 35
    ESTABLISHED: 18561
    Total: 18596
```

```
Bearers by State:
    BEARER_WAIT: 30
    ESTABLISHED: 18561
    Total: 18591
```

```
user@host> show services mobile-edge sessions
    Session-address: 23.0.21.163 State: ESTABLISHED Num-bearers: 1
        VRF-ID: 0x0 APN: default
        CPF-peer: 10.0.0.2 Access-peer: 20.0.0.2
        Anchor-PFE: apfe0:pfe-1/0/0 Secondary-anchor-PFE: apfe0:pfe-2/1/0
        Local-SEID: 0x20015a2 Remote-SEID: 0x3cb2

    Session-address: 23.0.47.237 State: ESTABLISHED Num-bearers: 1
        VRF-ID: 0x0 APN: default
        CPF-peer: 10.0.0.2 Access-peer: 20.0.0.2
        Anchor-PFE: apfe0:pfe-1/0/0 Secondary-anchor-PFE: apfe0:pfe-2/1/0
        Local-SEID: 0x2fec Remote-SEID: 0x56fc

    Session-address: 23.0.21.49 State: ESTABLISHED Num-bearers: 1
        VRF-ID: 0x0 APN: default
        CPF-peer: 10.0.0.2 Access-peer: 20.0.0.2
        Anchor-PFE: apfe0:pfe-1/0/0 Secondary-anchor-PFE: apfe0:pfe-2/1/0
        Local-SEID: 0x1531 Remote-SEID: 0x3c40

    Session-address: 23.0.29.83 State: ESTABLISHED Num-bearers: 1
        VRF-ID: 0x0 APN: default
        CPF-peer: 10.0.0.2 Access-peer: 20.0.0.2
        Anchor-PFE: apfe0:pfe-1/0/0 Secondary-anchor-PFE: apfe0:pfe-2/1/0
        Local-SEID: 0x2001d53 Remote-SEID: 0x4462

....
```

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
21.3R1	Starting in Junos OS Release 21.3R1, Junos Multi-Access User Plane provides a long-route implementation as a replacement for a filter-based implementation to steer traffic to the anchor Packet Forwarding Engine removing the need for a firewall filter to route GTP packets.

RELATED DOCUMENTATION

Anchor PFEs and Redundancy in Junos Multi-Access User Plane | 51
MX Series Router As Junos Multi-Access User Plane | 41

Example: Configure GBR Forwarding on Junos Multi-Access User Plane

IN THIS SECTION

Overview

This example shows how to configure GBR and best-effort forwarding on the Junos Multi-Access User Plane. GBR queues enable you to prioritize traffic. You can also guarantee a minimum bandwidth for GBR traffic.

Junos OS handles subscriber traffic differently depending on the Transport Level Marking (TLM) assigned to it. The TLM contains a DSCP value, which classifiers use to map the traffic to the appropriate forwarding class. Each forwarding class is associated with an output queue on the VRF loopback interface of the anchor Packet Forwarding Engine.

We support eight queues on the VRF interface, of which four are GBR queues and four are best-effort queues. These queues are preconfigured with a fixed set of attributes. You cannot modify the queue configurations.

Table 4: Forwarding Queues on the VRF Loopback Interface

Queue Number	Queue Name	Description
0	be-low-0	Best effort queue with lowest priority
1	be-low-1	Best effort queue with low to medium priority
2	be-high-0	Best effort queue with medium to high priority
3	be-high-1	Best effort queue with highest priority
4	gbr-low	GBR queue with low priority
5	gbr-med	GBR queue with medium priority
6	gbr-high	GBR queue with high priority
7	express	GBR queue with highest priority

⚠ **CAUTION**: Make sure that you map the forwarding classes to the correct queues. You can use different names for the forwarding classes, but you must map the queues correctly.

If the TLM provided for a new or modified GBR bearer or flow maps to a non-GBR queue, the system rejects the session establishment or modification. Make sure that the TLM configurations for the control plane and user plane are consistent.

In this example, we define eight forwarding classes and map the classes to the VRF queues. We define a classifier, `upf-class-1`, and configure the forwarding classes, their PLP values, and the applicable DSCP values. We define another classifier called `upf-class-2` and configure the default mapping for DSCP. We add these classifiers to peer routing instances associated with the outgoing interfaces on the gateway `upf`. Classifier `upf-class-1` handles traffic for routing instances `4g-access` and `4g-core`. Classifier `upf-class-2` handles traffic for all other routing instances. You can optionally define the fraction of the system bandwidth reserved for GBR queues.

Requirements

This example uses the following hardware and software components:

- MX480 (can also be MX240, MX960) router configured as an SAEGW-U or a UPF
- Junos OS Releases 22.2R1 or later

Configuration

IN THIS SECTION

CLI Quick Configuration

To quickly configure this example, copy the following commands and paste them into a text file. Remove any line breaks, and change any details necessary to match your network configuration. Copy and paste the commands into the CLI at the [edit] hierarchy level.

```
set class-of-service forwarding-classes class be-low-0 queue-num 0
set class-of-service forwarding-classes class be-low-1 queue-num 1
set class-of-service forwarding-classes class be-high-0 queue-num 2
set class-of-service forwarding-classes class be-high-1 queue-num 3
set class-of-service forwarding-classes class gbr-low queue-num 4
set class-of-service forwarding-classes class gbr-med queue-num 5
set class-of-service forwarding-classes class gbr-high queue-num 6
set class-of-service forwarding-classes class express queue-num 7
set class-of-service classifiers dscp upf-class-1 import default
set class-of-service classifiers dscp upf-class-1 forwarding-class be-high-0 loss-priority low
code-points cs1
set class-of-service classifiers dscp upf-class-1 forwarding-class be-high-0 loss-priority low
code-points cs2
set class-of-service classifiers dscp upf-class-1 forwarding-class be-high-0 loss-priority low
code-points cs3
set class-of-service classifiers dscp upf-class-1 forwarding-class be-high-0 loss-priority low
code-points cs4
set class-of-service classifiers dscp upf-class-1 forwarding-class be-high-0 loss-priority low
code-points cs5
set class-of-service classifiers dscp upf-class-1 forwarding-class be-high-1 loss-priority
medium-low code-points nc1
set class-of-service classifiers dscp upf-class-1 forwarding-class be-low-0 loss-priority medium-
high code-points be
set class-of-service classifiers dscp upf-class-1 forwarding-class be-low-1 loss-priority high
code-points nc2
set class-of-service classifiers dscp upf-class-1 forwarding-class express loss-priority high
code-points ef
set class-of-service classifiers dscp upf-class-1 forwarding-class gbr-high loss-priority low
code-points af11
set class-of-service classifiers dscp upf-class-1 forwarding-class gbr-high loss-priority low
code-points af12
set class-of-service classifiers dscp upf-class-1 forwarding-class gbr-high loss-priority low
code-points af13
set class-of-service classifiers dscp upf-class-1 forwarding-class gbr-low loss-priority medium-
high code-points af31
set class-of-service classifiers dscp upf-class-1 forwarding-class gbr-low loss-priority medium-
high code-points af32
set class-of-service classifiers dscp upf-class-1 forwarding-class gbr-low loss-priority medium-
high code-points af33
set class-of-service classifiers dscp upf-class-1 forwarding-class gbr-low loss-priority medium-
high code-points af41
set class-of-service classifiers dscp upf-class-1 forwarding-class gbr-low loss-priority medium-
high code-points af42
set class-of-service classifiers dscp upf-class-1 forwarding-class gbr-low loss-priority medium-
high code-points af43
set class-of-service classifiers dscp upf-class-1 forwarding-class gbr-med loss-priority medium-
low code-points af21
set class-of-service classifiers dscp upf-class-1 forwarding-class gbr-med loss-priority medium-
low code-points af22
```

```
set class-of-service classifiers dscp upf-class-1 forwarding-class gbr-med loss-priority medium-
low code-points af23
set services mobile-edge gateways saegw upf system class-of-service downlink-dscp-to-forwarding-
class classifier upf-class-1 routing-instance 4g-access
set services mobile-edge gateways saegw upf system class-of-service downlink-dscp-to-forwarding-
class classifier upf-class-1 routing-instance 4g-core
set services mobile-edge gateways saegw upf system class-of-service downlink-dscp-to-forwarding-
class classifier upf-class-2
set services mobile-edge gateways saegw upf system class-of-service reserved-bandwidth express 2
set services mobile-edge gateways saegw upf system class-of-service reserved-bandwidth gbr 10
```

Step-By-Step Procedure

To configure GBR and best-effort forwarding, follow the steps below:

1. Define eight forwarding classes and map them to the output queues on the VRF loopback interface.

```
[edit class-of-service forwarding-classes]
user@host# set class be-low-0 queue-num 0
user@host# set class be-low-1 queue-num 1
user@host# set class be-high-0 queue-num 2
user@host# set class be-high-1 queue-num 3
user@host# set class gbr-low queue-num 4
user@host# set class gbr-med queue-num 5
user@host# set class gbr-high queue-num 6
user@host# set class express queue-num 7
```

2. Define and configure the classifiers that will map the DSCP values to the forwarding classes. You can configure the forwarding class, the PLP value, and the DSCP values that the classifier applies to.

```
[edit class-of-service]
user@host# set classifiers dscp upf-class-1
user@host# set classifiers dscp upf-class-2

[edit class-of-service classifiers dscp upf-class-1]
user@host# set import default
user@host# set forwarding-class be-high-0 loss-priority low code-points [ cs1 cs2 cs3 cs4
cs5 ]
user@host# set forwarding-class be-high-1 loss-priority medium-low code-points nc1
user@host# set forwarding-class be-low-0 loss-priority medium-high code-points be
user@host# set forwarding-class be-low-1 loss-priority high code-points nc2
```

```
user@host# set forwarding-class express loss-priority high code-points ef
user@host# set forwarding-class gbr-high loss-priority low code-points [ af11 af12 af13 ]
user@host# set forwarding-class gbr-low loss-priority medium-high code-points [ af31 af32
af33 af41 af42 af43 ]
user@host# set forwarding-class gbr-med loss-priority medium-low code-points [ af21 af22
af23 ]

[edit class-of-service classifiers dscp upf-class-2]
user@host# set import default
```

3. Assign the classifiers to routing instances associated with the outgoing interfaces. If you do not assign any routing instance to a classifier, the system uses the classifier for all routing instances without an assigned classifier.

```
[edit services mobile-edge gateways saegw upf system class-of-service downlink-dscp-to-
forwarding-class]
user@host# set classifier upf-class-1 routing-instance 4g-access
user@host# set classifier upf-class-1 routing-instance 4g-core
user@host# set classifier upf-class-2
```

4. (Optional) Reserve a percentage of the total bandwidth for the express and GBR queues. The system reserves the same percentage of bandwidth across all anchor Packet Forwarding Engines.

```
[edit services mobile-edge gateways saegw upf system class-of-service reserved-bandwidth]
user@host# set express 2
user@host# set gbr 10
```

Results

From configuration mode, confirm your configuration by entering the `show services` and `show class-of-service` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show class-of-service
forwarding-classes {
    class be-high-0 queue-num 2;
    class be-high-1 queue-num 3;
    class be-low-0 queue-num 0;
    class be-low-1 queue-num 1;
    class express queue-num 7;
    class gbr-high queue-num 6;
    class gbr-low queue-num 4;
    class gbr-med queue-num 5;
}
classifiers {
    dscp upf-class-1 {
        import default;
        forwarding-class be-high-0 {
            loss-priority low code-points [ cs1 cs2 cs3 cs4 cs5 ];
        }
        forwarding-class be-high-1 {
            loss-priority medium-low code-points nc1;
        }
        forwarding-class be-low-0 {
            loss-priority medium-high code-points nc2;
        }
        forwarding-class be-low-1 {
            loss-priority high code-points be;
        }
        forwarding-class express {
            loss-priority high code-points ef;
        }
        forwarding-class gbr-high {
            loss-priority low code-points [ af11 af12 af13 ];
        }
        forwarding-class gbr-low {
            loss-priority medium-high code-points [ af31 af32 af33 af41 af42 af43 ];
        }
        forwarding-class gbr-med {
            loss-priority medium-low code-points [ af21 af22 af23 ];
        }
    }
    dscp upf-class-2 {
        import default;
    }
}
```

```
user@host# show services
mobile-edge {
    gateways {
        saegw upf {
            system {
                class-of-service {
                    downlink-dscp-to-forwarding-class {
                        classifier upf-class-2;
                        classifier upf-class-1 {
                            routing-instance 4g-access;
                            routing-instance 4g-core;
                        }
                    }
                    reserved-bandwidth {
                        express 2;
                        gbr 10;
                    }
                }
            }
        }
    }
}
```

Verification

IN THIS SECTION

Use various show commands to verify that the GBR forwarding is functioning properly.

Verify Forwarding Queue Configuration

IN THIS SECTION

Purpose

Display the configuration of the forwarding queues and verify that the system reserves bandwidth for the GBR queues.

Action

At the CLI, enter the `show services mobile-edge anchor class-of-service` command.

```
user@host> show services mobile-edge anchor class-of-service
Downlink forwarding queues/classes of service:
    Queue 0 (BE-LOW-0)
      Guaranteed priority:  low
      Excess priority:      low
      Reserved Bandwidth:   0 %
      Weight:               10

    Queue 1 (BE-LOW-1)
      Guaranteed priority:  low
      Excess priority:      low
      Reserved Bandwidth:   0 %
      Weight:               20

    Queue 2 (BE-HIGH-0)
      Guaranteed priority:  low
      Excess priority:      high
      Reserved Bandwidth:   0 %
      Weight:               20

    Queue 3 (BE-HIGH-1)
      Guaranteed priority:  low
      Excess priority:      high
      Reserved Bandwidth:   0 %
      Weight:               50

    Queue 4 (GBR-LOW)
      Guaranteed priority:  low
      Excess priority:      low
      Reserved Bandwidth:   10 %
      Weight:               0

    Queue 5 (GBR-MEDIUM)
      Guaranteed priority:  medium-high
      Excess priority:      low
      Reserved Bandwidth:   10 %
      Weight:               0

    Queue 6 (GBR-HIGH)
      Guaranteed priority:  high
      Excess priority:      low
      Reserved Bandwidth:   10 %
      Weight:               0

    Queue 7 (EXPRESS)
      Guaranteed priority:  strict-high
      Excess priority:      low
      Reserved Bandwidth:   2 %
      Weight:               0
```

Verify GBR and Non-GBR Flows

IN THIS SECTION

Purpose

Verify that the system classifies traffic into GBR and non-GBR bearers or flows.

Action

At the CLI, enter the `show services mobile-edge sessions summary` command.

```
user@host> show services mobile-edge sessions summary
Sessions by State:
    ESTABLISHED: 2000
    Total: 2000

Bearers by State:
    ESTABLISHED: 2000
    Total: 2000

Bearers by Downlink FAR State:
    FORWARD: 2000
    Total: 2000

Bearers by Resource Type:
    Non-GBR: 2000
    Total:   2000

5G QoS Flows by State:
    ESTABLISHED: 4000
    Total: 4000

5G QoS Flows by Resource Type:
    GBR:     2000
    Non-GBR: 2000
    Total:   4000
```

3

CHAPTER

Configuration Statements and Operational Commands

IN THIS CHAPTER

Junos CLI Reference Overview

We've consolidated all Junos CLI commands and configuration statements in one place. Read this guide to learn about the syntax and options that make up the statements and commands. Also understand the contexts in which you'll use these CLI elements in your network configurations and operations.

- Junos CLI Reference

Click the links to access Junos OS and Junos OS Evolved configuration statement and command summary topics.

- Configuration Statements
- Operational Commands