Engineering Simplicity

5G Fronthaul Class of Service—Juniper Validated Design (JVD)

Published
2025-06-05

Table of Contents

5G Fronthaul Class of Service—Juniper Validated Design (JVD)

Juniper Networks Validated Designs provide customers with a comprehensive, end-to-end blueprint for deploying Juniper solutions in their network. These designs are created by Juniper's expert engineers and tested to ensure they meet the customer's requirements. Using a validated design, customers can reduce the risk of costly mistakes, save time and money, and ensure that their network is optimized for maximum performance.

About this Document

This document provides a Juniper Validated Design (JVD) for implementing a Class of Service (CoS) architecture in a 5G Fronthaul network using the Juniper ACX7100 and Juniper ACX7509 Cloud Metro Routers. It also shares details on how the solution performs and provides key performance indicators (KPIs) to track its success.

After conducting thorough testing, we validated that the Juniper ACX7100 and ACX7509 provide a scalable and flexible platform for implementing CoS in 5G Fronthaul networks. These routers offer robust performance, high port density, and advanced traffic management capabilities, allowing for seamless expansion and adaptation to evolving network requirements.

This JVD is based on 5G reference architectures, but the CoS modeling, best practices, and performance results can be applied to many implementations.

Solution Benefits

Robust QoS management is an integral part of any 5G Fronthaul network. The Class of Service (CoS) functionality of the ACX7100 and ACX7509 routers offers numerous benefits, including:

- **Priority Handling**: The ACX7100 and ACX7509 have built-in features that enable priority handling for network traffic. These features allow for the effective implementation of Quality of Service (QoS) and CoS functionalities. CoS allows for prioritizing specific traffic types or applications, ensuring that critical services receive the necessary bandwidth and low latency. This prioritization is vital in 5G architectures, where ultra-reliable low latency communication (URLLC) workloads must coexist with high bandwidth applications.

- **Quality of Service (QoS)**: The ACX7100 and ACX7509 support various QoS techniques, including priority queuing, weighted fair queuing (WFQ), and weighted round-robin (WRR) scheduling. These techniques allow the routers to prioritize and allocate resources to different traffic classes based on their defined priorities. By assigning appropriate priority levels, the routers can ensure that time-sensitive or high-priority traffic, such as voice or video data, receives preferential treatment over other traffic types.
- **Traffic Shaping and Policing**: The ACX7100 and ACX7509 supports advanced traffic shaping and policing capabilities. Traffic shaping allows for controlling the flow of traffic to match the available bandwidth and manage congestion. Policing, on the other hand, enforces traffic limits by dropping or marking packets that exceed the defined thresholds.
- **Service Differentiation**: With CoS, service providers can offer differentiated service levels to their customers. They can prioritize premium services, guaranteeing a better user experience and meeting specific service-level agreements (SLAs). This capability helps service providers attract and retain customers by offering tailored services to meet their individual needs.
- **Scalability and Flexibility**: The Juniper ACX7100 and ACX7509 provide scalable and flexible platforms for implementing CoS in 5G Fronthaul networks. These routers offer robust performance, high port density, and advanced traffic management capabilities, allowing for seamless expansion and adaptation to evolving network requirements.

Use Case and Reference Architecture

IN THIS SECTION

5G xHaul architectures must support stringent CoS to meet critical demands of delay sensitive traffic characteristics. The reference architecture encompasses: Fronthaul, Midhaul, Backhaul, and Core segments, and follows recommendations of Open Radio Access Network (O-RAN) Alliance transport network architecture (ORAN.WG9.XPSAAS-v02.00).

Figure 1 on page 3 illustrates the physical segments of the 5G xHaul reference network, including the latency requirements. To confirm the effectiveness of the 5G xHaul and CoS reference architecture, we focused specifically on the Fronthaul segment, which is crucial for 5G technology.

Figure 1: 5G xHaul Reference Network and Latency Requirements

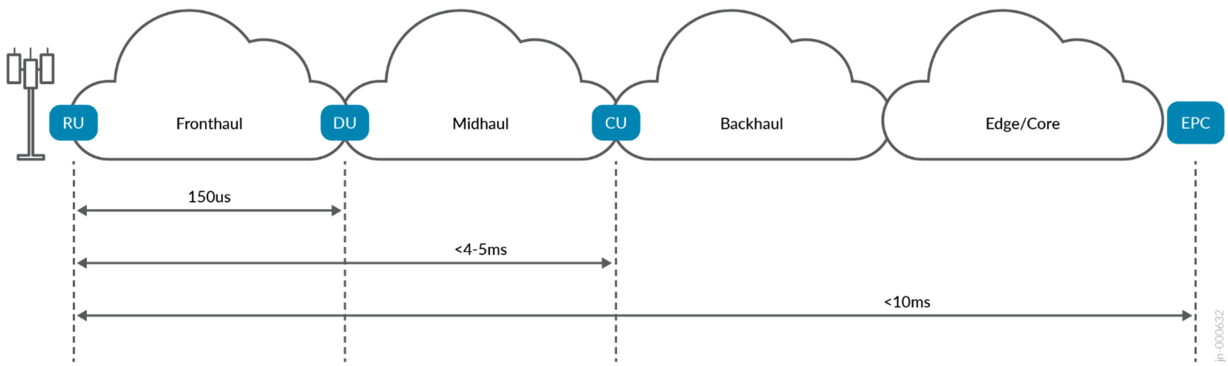

In this architecture, the Fronthaul portion in the radio access network (RAN), connects the Open Radio Unit (O-RU) and Open Distributed Unit (O-DU), (shown as RU and DU in Figure 1 on page 3), allowing them to communicate for control, data, and management traffic. It also ensures time and frequency synchronization between RAN elements. Because low latency is crucial (must be below 150µs from RU to DU), the Fronthaul segment has very few network elements, typically limited to one or two hops. Fronthaul devices must be able to deliver sub-10µs latency requirements.

Fronthaul Network Topology

The Fronthaul network deployment scenarios were carefully designed to support both the traditional 4G Mobile Backhaul and the evolution into the 5G network infrastructure over the same physical network. This approach allows MSOs to make a smooth transition from 4G to 5G without disrupting their existing services. They can gradually introduce the necessary changes and upgrades to accommodate the new requirements of 5G networks.

The Fronthaul network consists of ACX7000 series routers interconnected by high-capacity links. The ACX7100-48L serves as the Cell Site Router (CSR), providing connectivity between O-RU and Hub Site Router (HSR). The ACX7100-48L supports a range of port densities, including 47 ports of 10/25/50G + 1x10/25G and 6 x 400G (24x100G). The ability to support 400G access topologies is a key building block to the solution.

The ACX7100-32C and ACX7509 serve as the HSRs. The HSRs provide connectivity from the CSR to the O-DU, including pre-aggregation for Midhaul and Backhaul network segments. The ACX7100-32C supports 32 x 100G and 4x400G, and the ACX7509 supports three FPC variants (16x400GE, 4x200/400GE, and 20x1/10/25/50GE).

Solution QoS Architecture and Design

IN THIS SECTION

The test bed, illustrated in Figure 2 on page 5, emulates a Seamless SR-MPLS architecture for a MBH with four layers of transport infrastructure: Access, Pre-Aggregation, Aggregation, and Core. In this setup, the 5G Fronthaul segment uses a spine-leaf access topology with redundant HSR (AG1.1/1.2) to support it. This configuration combines the roles of 4G Pre-Aggregation and 5G HSR. The Pre-Aggregation AG1 nodes provide O-DU connectivity and include additional emulated access insertion points (RT) to scale the environment. The Midhaul and Backhaul segments are represented by ring topologies, which include Aggregation and Core roles but are not the focus of the JVD.

Figure 2: 5G Fronthaul Services Topology

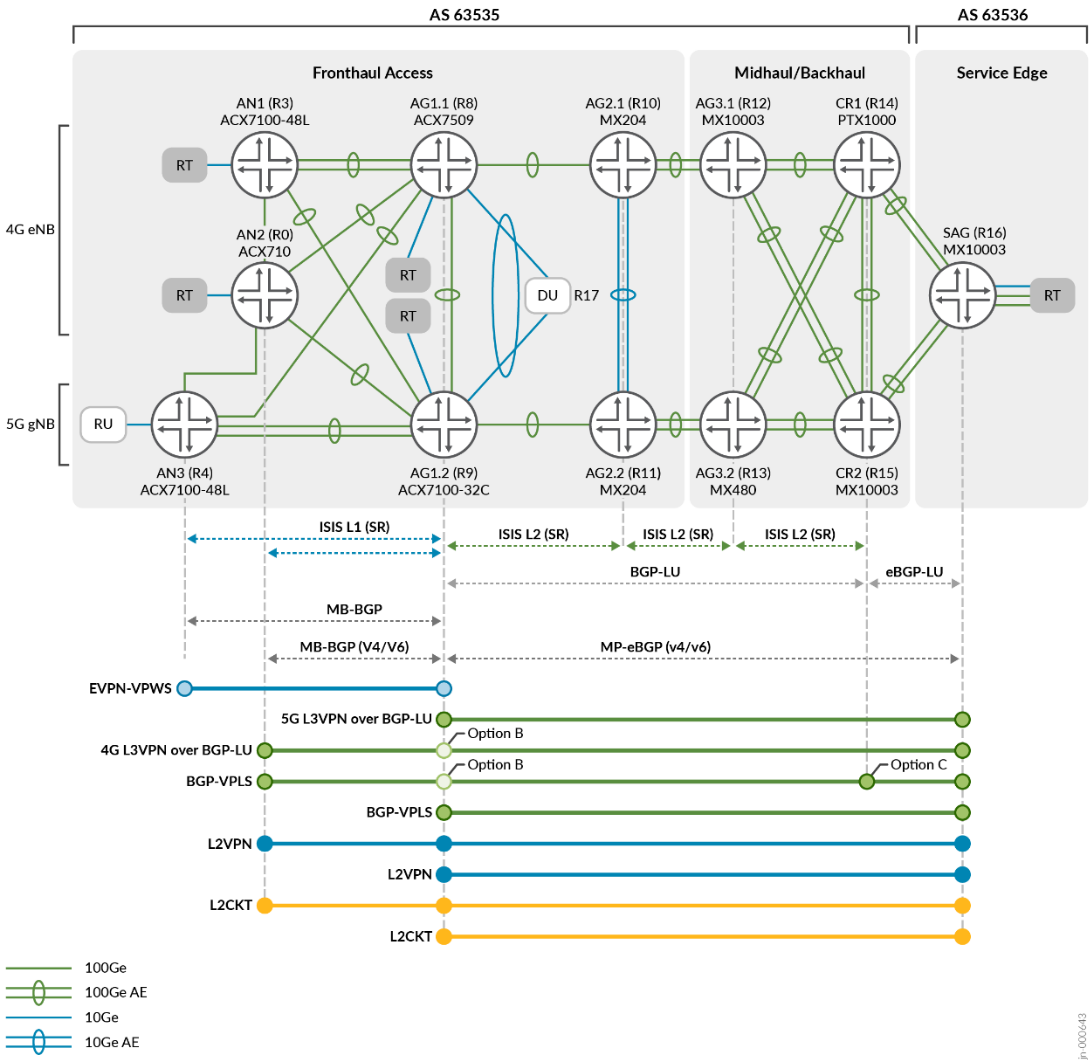

Traffic Flows

The test network includes the following types of traffic flows:

- Layer 2 eCPRI between O-RU to O-DU traffic flows – 5G Fronthaul
- Layer 3 IP packet flows between 5G O-DU and CU/EPC – 5G Midhaul and Backhaul
- Layer 3 IP packet flows between 4G CSR and EPC (SAG) – 4G L3-MBH
- Layer 2 flows between CSR (AN) to EPC (SAG) – 4G L2-MBH

- Layer 2 Midhaul flows emulating additional attachment segments – 4G Midhaul and Backhaul

Segment Routing

Segment Routing MPLS is used as the underlying technology, which spans across multiple ISIS domains and inter-AS connections. The access nodes are placed within an ISIS L1 domain, establishing adjacencies with L1/L2 HSR nodes. The L2 domain extends from the aggregation to the core segments of the network. To achieve seamless SR-MPLS connectivity, BGP Labeled Unicast (BGP-LU) is enabled at the border nodes.

Route Reflectors

To handle the increased network scale and facilitate continued network growth, two sets of route reflectors are used at CR1 and CR2, primarily serving the westward HSR (AG1) clients. AG1.1/AG1.2 act as redundant route reflectors specifically for the access Fronthaul segment. Inter-AS Option-B solutions are supported through Multi-Protocol BGP peering between the Services Aggregation Gateway router (SAG) and the HSR (AG1).

Overlay Services

The overlay services utilize different combinations of VLAN operations, as shown in Table 3 on page 13. These operations are applied to each Layer 2 service type defined in this JVD, including EVPN, L2Circuit, VPLS, and L2VPN. Flow Aware Transport Pseudowire Label (FAT-PW) is supported, starting with the Junos OS Evolved 22.3R1, for L2Circuit and L2VPN services, and is included in this JVD. Ethernet OAM with performance monitoring is enabled for EVPN Fronthaul and VPLS MBH services, ensuring effective monitoring of the network performance. Additionally, L3VPN services incorporate IPv6 tunnelling to validate 6vPE functionality.

5G QoS Identifier (5QI) Model

In 4G LTE networks, the QoS model is based on bearers. A bearer represents a logical connection between the user equipment (UE) and the network, and each bearer is associated with a specific QoS Class Identifier (QCI). The QCI determines the characteristics and priority of traffic flows within that bearer. The 4G LTE QCI model focusses on providing different levels of service based on predefined QCI values, which are standardized and mapped to specific QoS parameters such as packet delay, packet loss, and throughput.

In contrast, 5G networks introduce a flow-based QoS model using the concept of QoS Identifiers (5QIs). In this model, QoS is associated with individual data flows within the network. A flow represents a specific communication path or data stream between the UE and the network. Each flow is assigned a unique 5QI value that indicates the required QoS parameters for that flow, such as latency, reliability, and throughput.

The flow-based 5G QoS model provides more granular control and flexibility compared to the bearer-based 4G LTE QCI model. Instead of assigning a single QoS value to an entire bearer, 5G allows for different flows within the same bearer to have their own individual QoS requirements. This enables more fine-grained QoS management, allowing different applications and services to receive tailored treatment based on their specific needs. It also enables the network to dynamically adjust QoS parameters for individual flows, optimizing resource allocation and providing a better user experience.

When transitioning from the 4G LTE QCI model to the 5G QoS Identifier (5QI) model, there are some similarities in how we define and handle traffic. However, in 5G, we introduce new categories specifically for flows that require very low delay and high bandwidth. These are crucial for user and control traffic streams between O-RU and O-DU in the 5G Fronthaul segment.

5G incorporates several new delay-critical Guaranteed Bit Rate (GBR) flow categories. In the 5G Fronthaul segment, eCPRI based flows define the user and control traffic streams between O-RU and O-DU. These traffic characteristics require high bandwidth and extremely low delay. All devices in the access topology must assign the highest priority to this traffic type. This means that Service Providers may need to rethink their existing Class of Service designs where the highest priority queues were typically reserved for voice or video traffic.

The O-RAN 5QI/QCI Exemplary Grouping model, shown in Figure 3 on page 8 and defined in O-RAN [O-RAN.WG9.XPSAAS-v02.00], groups common QCI and 5CI QoS flow characteristics into four exemplary groups based on delay budget.

Figure 3: O-RAN 5QI/QCI Exemplary Grouping

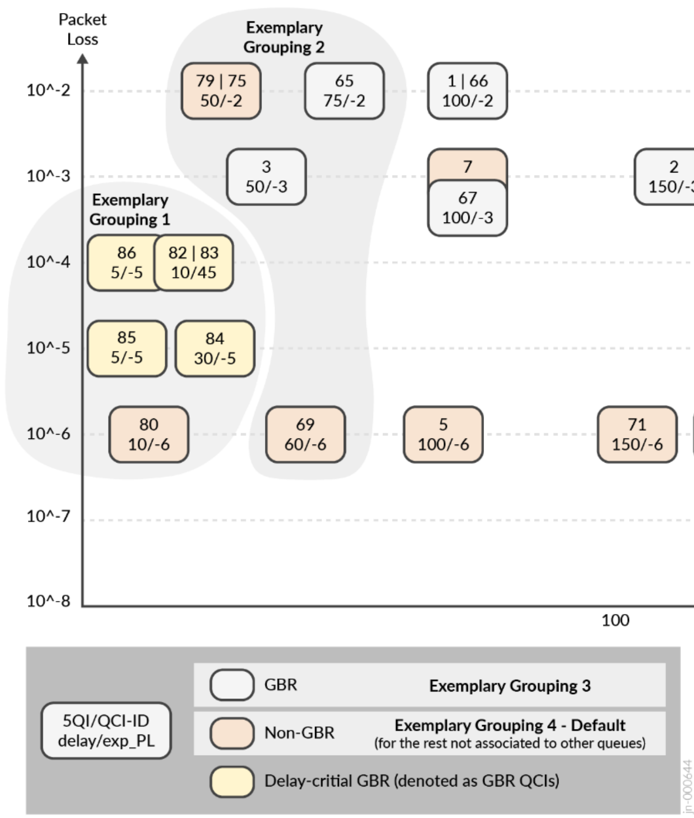

QoS schemas vary between mobile operators. This JVD doesn't endorse any specific design as the recommended approach.

A major objective is to examine predictable behaviors in terms of how critical and non-critical traffic flows are handled across services in the 5G xHaul network. The implementation of QoS should exhibit deterministic functionality. The transport architecture must prove capable to support the adaptation of existing and emerging mobile applications, preserving delay budget integrity while guaranteeing traffic priorities.

For more information about the specific latency and delay budgets considered for this validation, reach out to your Juniper Networks representative.

O-RAN/3GPP proposes two common QoS profiles to accommodate transport network requirements. In Profile A (Figure 4 on page 9), a single priority queue is defined for handling ultra-low latency flows such as PTP and eCPRI. This queue is serviced ahead of all other queues. Lower priority queues are then serviced as weighted fair queuing (WFQ) round robin.

The Profile B model uses a hierarchy of queue priorities: high, medium, and low. These priority queues support preemption to minimize packet delay variations (PDV) and prioritize critical flows that require low latency. Specifically, the queue assigned for eCPRI traffic needs to have the ability to interrupt or take priority over other queues.

Figure 4: Single Priority Queue (Profile A)

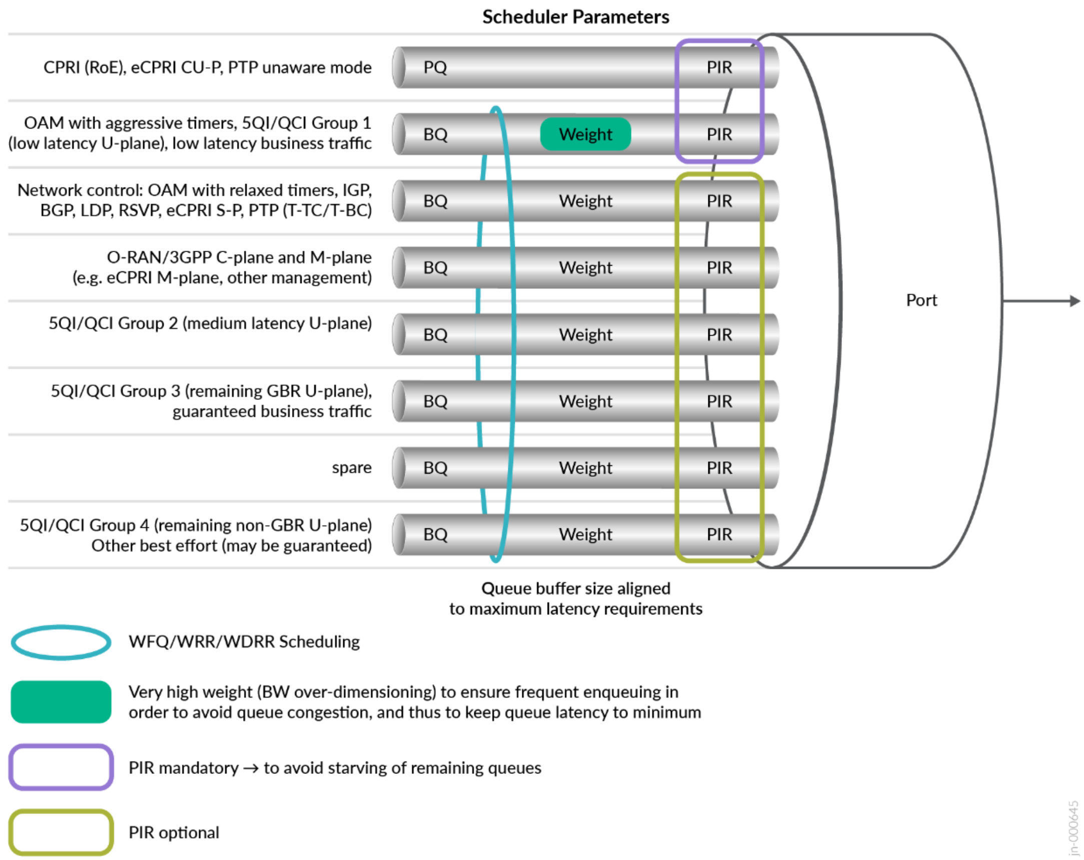

ACX Metro Routers support multiple strict-high (SH) or low priority queues. Strict-high queues are serviced as round-robin without the ability to preempt another priority queue. Profile A is selected for this JVD, reserving a strict-high queue for ultra-low latency between RU and DU.

Class of Service Building Blocks

Quality of Service (QoS) governs how traffic is forwarded, stored, or dropped in conjunction with mechanisms to manage and avoid congestion, which includes the basic building blocks:

- Classification
- Scheduling and Queuing
- Rewriting
- Shaping & Rate Limiting

QoS models differ between operators based on unique traffic profiles and characteristics. For the solution architecture, a realistic CoS model is created based on a pseudo customer, shown in Table 1 on page 10. For more details on the CoS model leveraged in this JVD, contact your Juniper Networks representative.

Table 1: Validated Scheduling Profiles

Forwarding Classes	Scheduling Parameters				Classification & Rewrite				Traffic Profile		
	Queue	Queue Priority	Transmit rate	Buffer size	802.1p	DSCP	MPLS EXP	Packet Loss Priority	Resource Type	Traffic Type	QCI/5CI Mapping
Business	5	Low	20%	20%	4	CS4, AF4x	4	Low	GBR	Guaranteed U-Plane Business Conversational Real Time Gaming/ Video	QCI1-4, 6 QCI65-67
Network Control	4	Low	5%	2%	7 6	CS7 CS6	7 6	Low	GBR	Protocol, Timing	QCI82-90
Real Time	2	Strict High	40% Shaped	30%	5	CS5 EF	5	Low	Delay-Critical GBR	eCPRI	CPRI QCI82-90
Signaling & OAM	3	Low	5%	2%	3	CS3, AF3x	3	Low	Non-GBR	Signaling & OAM	QCI5
Medium	1	Low	20%	20%	2	CS2, AF2x	2	High	Non-GBR	Streaming Interactive	QCI4, 6-8

Table 1: Validated Scheduling Profiles *(Continued)*

Forwarding Classes	Scheduling Parameters				Classification & Rewrite				Traffic Profile		
	Queue	Queue Priority	Transmit rate	Buffer size	802.1p	DSCP	MPLS EXP	Packet Loss Priority	Resource Type	Traffic Type	QCI/5CI Mapping
Best Effort	0	Low	Remainder	Remainder	1 0	CS1, AF1x RF	1 0	Low High	Non-GBR	Background	QCI9

The solution architecture includes two styles of ingress classification: fixed and behavior aggregate. Fixed classification is context-based where all traffic arriving on a specific interface is mapped into one forwarding class. Behavior Aggregate is packet-based where flows were pre-marked with Layer 3 DSCP, Layer 2 802.1Q Priority Code Points (PCP) or MPLS EXP.

O-RAN/3GPP proposes a minimum of six queues and maximum of eight queues per interface. All platforms in this solution architecture support eight queues. Six queues and associated forwarding classes are used to accommodate the traffic scheme requirements. In supporting Profile-A, only one strict-high queue is used, which is shaped (PIR) to prevent starving low priority queues. All other queues are configured as low priority and serviced as weighted fair queuing (WFQ) round robin, based on the designated transmit-rate.

At egress, DSCP, 802.1p or EXP codepoints and loss priorities (PLP) are rewritten based on the assigned forwarding class and rewrite-rule instruction. ACX Metro Routers support rewriting the outer tag by default. Generally, providers are required to preserve inner (C-TAG) 802.1p bits and transmitted transparently, without any changes or modifications.

Service Carve Out

As a best practice, ultra-low latency services (eCPRI) are assigned the highest priority. The treatment of MBH applications varies. The solution architecture uses priority mappings, grouped by service type. Refer to Table 2 on page 12.

Table 2: Service Definitions

Service	Traffic Type	Forwarding Class	Classifier Type	Priority
EVPN-VPWS	Delay-Critical GBR (eCPRI)	Realtime	Fixed	Strict High
L2Circuit	Non-GBR wholesale user plane	Best Effort	Fixed	Low
L2VPN	4G/5G medium user plane	Best Effort/Medium	Behavior Aggregate	Low
BGP-VPLS	Non-GBR/GBR user plane	Best Effort/Business	Behavior Aggregate	Low
L3VPN	C/M/U-plane GBR/non-GBR	BE/MED/SIG-OAM/ Business	Behavior Aggregate	Low

VLAN Operations

ACX7000 series supports a more comprehensive set of VLAN manipulation operations compared to previous generation ACX routers. This JVD doesn't include all possible operations, but does validate 60 VLAN combinations across L2Circuit, L2VPN, and EVPN services.

The test scenarios include the following VLAN operations:

- Untagged (UT) / Native VLAN
- Single-tag (ST) operations (pop, swap, push)
- Dual-tag (DT) operations (swap-swap, pop-swap/swap-push, pop-pop/push-push, swap-push/pop-swap)
- Rewrite PCP bits
- Preservation of PCP bits
- Classification of PCP bits and FC mapping

Table 3 on page 13 shows the VLAN operations that we validated for EVPN-VPWS, L2Circuit and L2VPN services. Contact your Juniper Networks representative for comprehensive details about the VLAN operations that we validated, including configuration of each operation.

Table 3: Validated VLAN Operations

VLAN Type	Outer Tag	Inner Tag	Input Operation	Output Operation	Classification	Rewrite
dual	101	2201	None	none	fixed	exp rewrite
dual	102	2202	Pop	push	fixed	exp rewrite
dual	103	2203	Swap	swap	fixed	exp rewrite
dual	104	2204	swap-swap	swap-swap	fixed	exp rewrite
dual	105	2205	pop-swap	swap-push	fixed	exp rewrite
dual	106	2206	pop-pop	push-push	fixed	exp rewrite
single	107	--	Push	pop	fixed	exp rewrite
single	108	--	Swap	swap	fixed	exp rewrite
single	109	--	Pop	push	fixed	exp rewrite
single	110	--	swap-push	pop-swap	fixed	exp rewrite
dual	101	2201	None	none	BA (exp)	802.1p rewrite
dual	102	2202	Pop	push	BA (exp)	802.1p rewrite
dual	103	2203	Swap	swap	BA (exp)	802.1p rewrite
dual	104	2204	swap-swap	swap-swap	BA (exp)	802.1p rewrite
dual	105	2205	pop-swap	swap-push	BA (exp)	802.1p rewrite
dual	106	2206	pop-pop	push-push	BA (exp)	802.1p rewrite
single	107	--	Push	pop	BA (exp)	802.1p rewrite
single	108	--	Swap	swap	BA (exp)	802.1p rewrite
single	109	--	Pop	push	BA (exp)	802.1p rewrite

Table 3: Validated VLAN Operations *(Continued)*

VLAN Type	Outer Tag	Inner Tag	Input Operation	Output Operation	Classification	Rewrite
single	110	--	swap-push	pop-swap	BA (exp)	802.1p rewrite

O-RAN and eCPRI Emulation

The O-RAN and eCPRI test scenarios included several functional permutations with emulated O-DU and Remote Radio Unit (RRU). Innovative steps are taken to validate not only performance, but assurance of 5G eCPRI communication, U-Plane message type behaviors, and O-RAN Conformance. These steps ensure the featured DUTs are correctly and consistently transporting these critical 5G services.

Results are further evolved across multiple topologies to capture a range of DUT performance characteristics in non-congested and congested scenarios. Note that this part of validation was not integrated into regression cycles and not reported in this JVD. If you are interested in details of test methodology and results for eCPRI emulation test series, please contact your Juniper Networks representative.

Here is the summary of ORAN and eCPRI test scenarios:

- eCPRI O-RAN Emulation: Leverages a standards IQ Sample file to generate flows.
- O-RAN Conformance: Analyzes messages for conformance to O-RAN specification.
- Crafted eCPRI O-RAN: Produces variable eCPRI payload for comparing latency performance.
- eCPRI Services Validation: Emulates User Plane messages, performing functional and integrity analysis.
- eCPRI Remote Memory Access (Type 4) : Performs read or write from/to a specific memory address on the opposite eCPRI node and validates expected success/failure conditions.
- eCPRI Delay Measurement Message (Type 5): Estimates one-way delay between two eCPRI ports.
- eCPRI Remote Reset Message (Type 6): One eCPRI node requests a reset of another node. Validates expected sender/receiver operations.
- eCPRI Event Indication Message (Type 7): Either side of the protocol indicates to the other an event has occurred. An event is either a raised or ceased fault or a notification. Validation confirms events raised as expected.

Figure 5: eCPRI User Plane Common Header Format

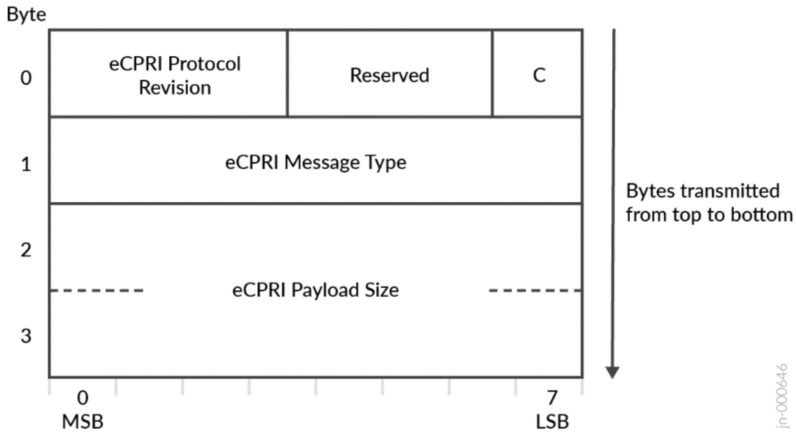

Contact your Juniper Networks representative for comprehensive test results.

Solution and Validation Key Parameters

IN THIS SECTION

The test network consists of Fronthaul, Midhaul, Backhaul, and Service Edge segments. Refer to .

- **Fronthaul segment**: Uses a spine-leaf access topology, connecting to redundant HSR (AG1.1/1.2) nodes, which also handle 4G pre-aggregation and 5G HSR functions. The pre-aggregation AG1 nodes provide connectivity for O-DUs and include additional emulated access insertion points (RT) for scalability.
- **Midhaul and Backhaul segments:** These are represented by ring topologies and serve aggregation and core roles, but they are not the focus of this discussion.

Figure 6: 5G Fronthaul Lab Topology

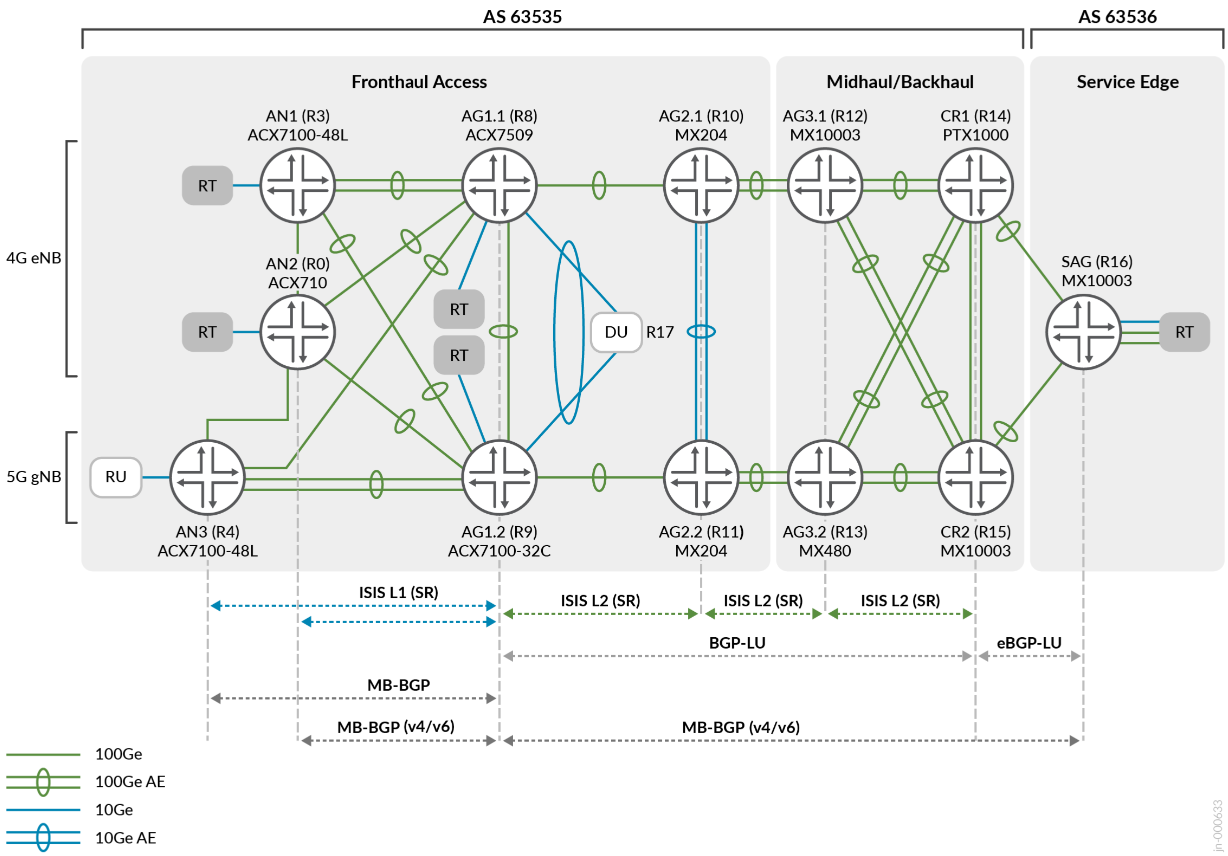

Supported Platforms and Positioning

To review the software versions and platforms on which this JVD was validated by Juniper Networks, see the Validated Platforms and Software section in this document.

Solution Validation Goals

The main goal is to validate that the ACX7100 and ACX7509 platforms can meet the CoS performance requirements for maintaining the reliability and quality of critical 5G Fronthaul traffic between the RU and the DU. Our secondary objective is to verify that the CoS behaviors remain consistent across the mobile backhaul (MBH) topology and other devices at different capture points.

CoS Functionality Goals

For this JVD, we validated CoS operations and functionalities to determine if the topology can deliver the following goals:

- Classify traffic based on DSCP, 802.1p, and EXP with Packet Loss Priority (PLP) high and low.
- Preserve QoS codepoints end to end for inner and outer tags.
- Support ingress classification using Fixed and Behavior Aggregate styles.
- Create at least six forwarding classes and six queues (all featured platforms support eight)
- Create and configure two priority queues, one with strict-high priority and the other with low priority. Schedule these queues based on a specified percentage of bandwidth allocation (transmit rate) and buffers. The strict-high priority queue is prioritized ahead of the low priority queue.
- The port shaper inherits and scales configured scheduling characteristics.
- Rewrite operations, based on queue assignment, support 802.1p, DSCP, and EXP.
- For dual-tagged frames, the rewrite operation only affects the outermost tag, leaving the inner tag intact.
- Latency budgets in scenarios where the network is not congested and is offering a line rate below 100%, the strict-high queue, which has a higher priority, remains within the following latency budget:
  - O-RU-to-O-DU latency equates to averaging ≤10µ per device (≤6µs single DUT).
  - RU-to-SAG latency ≤10ms (expected ≤150µ).
  - Enhanced Common Public Radio Interface (eCPRI) Type 5 One Way Delay Measurement ≤6µ per device.

We also validated how the CoS network manages congestion scenarios including the following key objectives for 5G architectures:

- Preserve critical eCPRI Fronthaul traffic
- Maintain traffic priorities across shared links
- Maintain traffic priorities within and between VPN services that share common links
- Detect eCPRI Type 7 event indication errors when the network experiences congestion
- Maintain consistent CoS and resiliency across negative stress conditions (enabled/disable control and data plane daemons, add/delete configurations, and so on.)

Another goal was to identify any limitations, anomalies, and open problem reports (PRs) that were discovered during the validation stages. Whenever possible, we addressed and fixed these PRs, and then validated the fixes as part of the test execution to ensure that everything was functioning as intended.

Solution Validation Non-Goals

Non-goals represent protocols outside the scope of this JVD:

- Latency validation under congestion scenarios.
- Multifield classification to forwarding class mapping.
- Custom drop profiles weighed random early detection (WRED).
- Temporal transmit rate or buffer (elastic buffer is used).
- Underlay/Overlay validation other than those specified in the Solution Goals section.
- Failover and convergence scenarios.
- End-to-End Timing and Synchronization Distribution: Synchronous Ethernet, IEEE1588v2.
- SLA Monitoring: RFC 2544, Y.1564, TWAMP, Active Assurance.
- Telemetry, Management, and Automation.

Failure Scenarios

- The validation covers the following failure scenarios:
- Injected failure events (including eCPRI type)
- Link congestion
- Queue congestion
- Process restart

Results Summary and Analysis

IN THIS SECTION

ACX7100 and ACX7509 Functions and Performance

The CoS validation focused on demonstrating ACX7100 and ACX7509 classification, scheduling, shaping and rewriting behaviors across services leveraging a 5G xHaul infrastructure. The test scenarios included measuring latency for eCPRI Fronthaul traffic types with emulated O-RU/O-DU and O-RAN test sequences. The JVD is executed over a Seamless SR-MPLS architecture with a focus on CoS.

The topology generates reasonable multi-vector scale of Layer 2/Layer 3 connectivity services as compared with mobile network operator (MNO) and metro area network (MAN) operator expectations for real network deployments, while satisfying stringent SLA requirements.

The ACX7100 and ACX7509 meet the CoS requirements mentioned here for 5G xHaul solutions. They are capable of preserving and respecting priority forwarding throughout the network topology. These platforms are particularly well-suited for access aggregation purposes.

Figure 7: Class of Service Functional Diagram

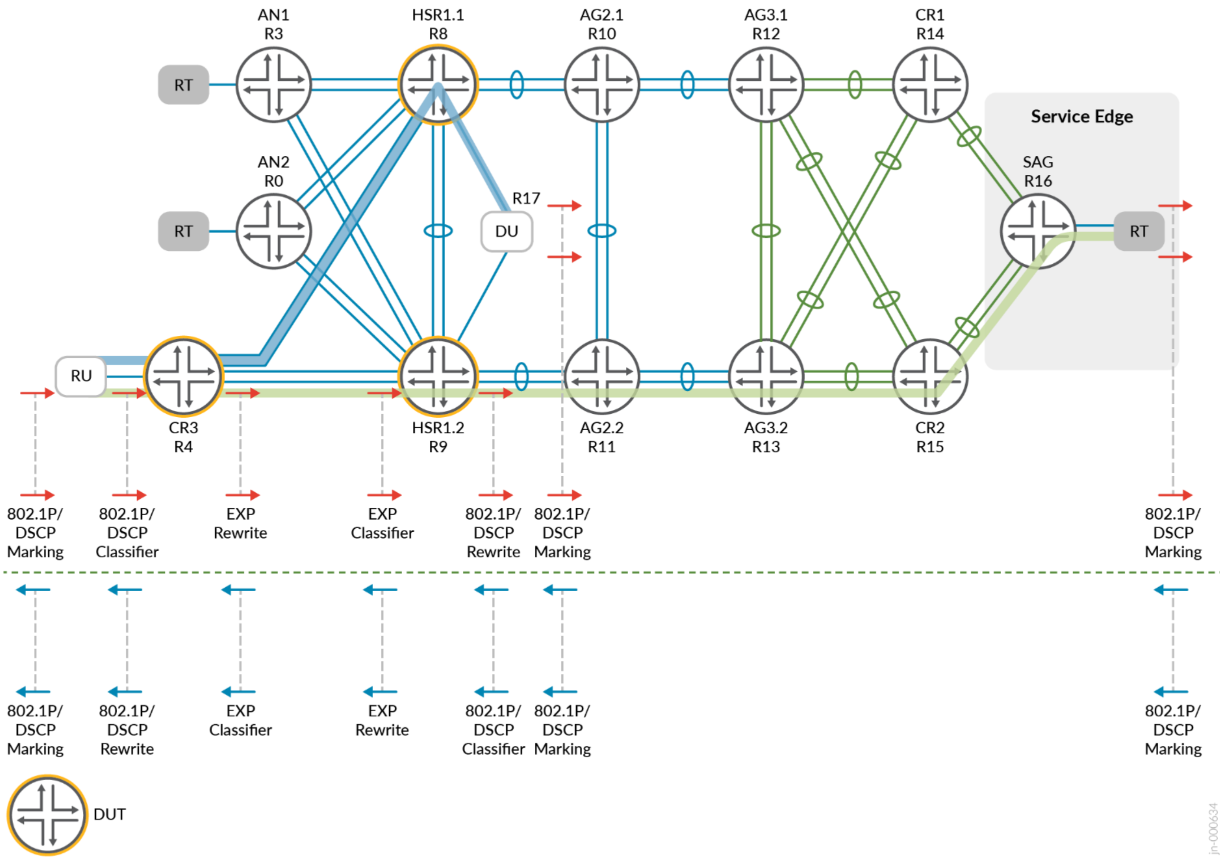

Across the end-to-end topology, classification and rewrite was performed on 802.1p, DSCP and EXP, as outlined in Figure 7 on page 20. Table 5 on page 21 summarizes these results for the included services and classification types. In dual-tag scenarios, the outer service tag is used for classification and rewrite. CoS bits can be preserved end-to-end, including for inner or outer tags.

When a port shaper is defined, applicable class of service functions adjusted to the new port speed and performed equivalently. For example, a 1G port shaper was used and transmit-rate percentages were correctly shown to be based on a 1G port speed.

Table 4 on page 21 provides a summary of the executed VLAN operation scenarios and their respective outcomes. The scenarios include Untagged (UT), Single-Tagged (ST), and Dual-Tagged (DT) VLAN tag representations. Table 4 on page 21 presents a condensed overview of these scenarios and their results.

Table 4: CoS Summarized Results

Traffic Scenario	VLAN	Ingress Classification Mapped to FC			Scheduler Honored	Rates	Codepoints Rewritten			Bits Preserved
Fixed Classifier	TAG	802.1p	DSCP	EXP	SH	LOW	802.1p	DSCP	EXP	E2E
EVPN-VPWS	UT	√	--	√	√	--	√	--	√	√
L2Circuit	UT	√	--	√	√	--	√	--	√	√
pop / push	DT	√	--	√	√	--	√	--	√	√
swap / swap	DT	√	--	√	√	--	√	--	√	√
swap-swap / swap-swap	DT	√	--	√	√	--	√	--	√	√
pop-swap / swap-push	DT	√	--	√	√	--	√	--	√	√
pop-pop / push-push	DT	√	--	√	√	--	√	--	√	NA
push / pop	ST	√	--	√	√	--	√	--	√	√
swap / swap	ST	√	--	√	√	--	√	--	√	√
pop / push	ST	√	--	√	√	--	√	--	√	√
swap-push / pop-swap	ST	√	--	√	√	--	√	--	√	√
BA Classifier	TAG	802.1p	DSCP	EXP	SH	LOW	802.1p	DSCP	EXP	E2E

Table 4: CoS Summarized Results *(Continued)*

Traffic Scenario	VLAN	Ingress Classification Mapped to FC			Scheduler Honored	Rates	Codepoints Rewritten			Bits Preserved
L3VPN	UT	--	√	√	--	√	--	√	*	√
L2VPN	UT	√	--	√	--	√	√	--	*	√
BGP-VPLS	UT	√	--	√	--	√	√	--	*	√
pop / push	DT	√	--	√	--	√	√	--	√	√
swap / swap	DT	√	--	√	--	√	√	--	√	√
swap-swap / swap-swap	DT	√	--	√	--	√	√	--	√	√
pop-swap / swap-push	DT	√	--	√	--	√	√	--	√	√
pop-pop / push-push	DT	√	--	√	--	√	√	--	√	NA
push / pop	ST	√	--	√	--	√	√	--	√	√
swap / swap	ST	√	--	√	--	√	√	--	√	√
pop / push	ST	√	--	√	--	√	√	--	√	√
swap-push / pop-swap	ST	√	--	√	--	√	√	--	√	√

*Contact your Juniper Networks representative for details.

All listed input/output VLAN mapping operations were validated across L2Circuit, L2VPN, and EVPN services. Contact your Juniper Networks representative for an analysis explaining the results for each function.

Latency Budgets

The 5G xHaul infrastructure has specific guidelines for latency, especially in the Fronthaul segment where extremely low latency flows are crucial. The overall latency budget considers factors such as fiber length, connected devices, and transport design. O-RAN sets a maximum limit of 100µs for one-way latency from O-RU to O-DU, with each device having a latency of around ≤10µs. However, there is a growing demand from operations for device latency to be even lower, around 5-6µs. This represents a significant shift in requirements compared to earlier 4G architectures.

Table 5 on page 23 provides a snapshot of the performance comparison in Fronthaul, Midhaul, and MBH for various service types on ACX7100 and ACX7509. The total latency factors in the number of hops, such as EVPN-VPWS with three hops resulting in a measurement of 15.1µs, approximately 5µs per hop in the Fronthaul segment. For a more in-depth analysis of per-device latency for ultra-low latency traffic types, refer to the Latency measurements in Table 5 on page 23. If you need additional information, including specific data outputs and reports, contact your Juniper Networks representative.

Table 5: Latency Measurements

Traffic Type	Latency (µs)	Queue Type	Port Speed	Segment	DUT	Hops
EVPN-VPWS (eCPRI)	15.1µs	strict-high	10G	Fronthaul	ACX7100	3
EVPN-VPWS (eCPRI)	16.1µs	strict-high	10G	Fronthaul	ACX7509	2
EVPN-VPWS (eCPRI)	10.6µs	strict-high	1G Shaper	Fronthaul	ACX7100	3
EVPN-VPWS (eCPRI)	8.4µs	strict-high	1G Shaper	Fronthaul	ACX7509	2
L2Circuit	62.3µs	Low	100G	Midhaul	ACX7509	6
L2Circuit	74.6µs	Low	10G	Midhaul	ACX7509	5
L2Circuit	103.5µs	Low	1G Shaper	Midhaul	ACX7509	5
L2Circuit	68.2µs	Low	100G	MBH	ACX7100	6
L2Circuit	81.8µs	Low	10G	MBH	ACX7100	6
L2Circuit	127.4µs	Low	1G Shaper	MBH	ACX7100	6

Table 5: Latency Measurements *(Continued)*

Traffic Type	Latency (µs)	Queue Type	Port Speed	Segment	DUT	Hops
L3VPN (CRIT)	61.6µs	strict-high	100G	Midhaul	ACX7509	6
L3VPN	126µs	Low	10G	Midhaul	ACX7509	5
L3VPN	117.7µs	Low	1G Shaper	Midhaul	ACX7509	5
L3VPN	145.9µs	Low	10G	MBH	ACX7100	6
L3VPN	137.1µs	Low	1G Shaper	MBH	ACX7100	6

In non-congestion scenarios, the priority queue provides latency performance improvement over low priority queues.

Recommendations

ACX7100 and ACX7509 platforms support deterministic and effective QoS, performing within expectations for the 5G CSR and HSR in the Fronthaul segment.

Layer 3 (DSCP), MPLS (EXP), and Layer 2 (802.1p) single-tagged or dual-tagged traffic were correctly classified based on appropriate codepoints. Priority hierarchies were honored within expected guaranteed and excess regions, defined by transmit rate. Shaping the strict-high queue prevents starving low priority queues, which were guaranteed the configured committed information rates (CIR) using transmit-rate percentages. When a port is shaped, the configured CoS scheduling parameters are correctly adjusted. Codepoint preservation was achieved across all featured VLAN manipulation sequences (where expected). The platforms met latency performance requirements in non-congested scenarios, with 64b frame size achieving the best results. All stated goals were passed, with a few deviations explained below.

During the validation, several functional differences were observed with ACX7000 series compared to ACX5448/ACX710 platforms. Contact your Juniper Networks representative for the complete test report.

- ACX7100/ACX7509 supports significantly more VLAN manipulation operations compared to earlier ACX platforms, to bring relative parity with MX platforms.
- 802.1p bits are preserved without incurring a default rewrite, where ACX5448/ACX710 performs default rewrite.

- EXP bits are preserved only when default exp classification is used.
- Multiple rewrite rules are allowed with support for both multiple classification and rewrite.
- The ACX7000 series does not include simultaneous use of transmit-rate and shaping-rate for a queue. When only shaping-rate is utilized, the behavior is such that the strict-high queue deducts the configured rate from the total port speed, and the remaining bandwidth is allocated to the other queues based on their configured percentages. For example, if a 10G port has a 40% shaping-rate queue (SH), it deducts 4G from the total port speed. The remaining 6G is allocated to the other queues. For instance, a low priority queue with a 50% transmit-rate receives 3G.