

5G Mobile xHaul with Seamless MPLS Segment Routing—Juniper Validated Design (JVD)

Published
2025-06-05

Table of Contents

About this Document 1
Solution Benefits 1
Use Case and Reference Architecture 2
Solution Design and Architecture 4
Solution and Validation Key Parameters 17
Results Summary and Analysis 29
Recommendations 40

5G Mobile xHaul with Seamless MPLS Segment Routing—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 presents a Juniper Validated Design (JVD) for a 5G xHaul network using the Juniper ACX7000 series, MX series, and PTX series with a seamless MPLS segment routing framework. The JVD extends solutions presented earlier by [5G Fronthaul Network Using Seamless MPLS Segment Routing](#) and [5G Fronthaul Class of Service](#) JVDs and focusses on insertion of the ACX7024 (AN4) with Junos OS Evolved, which serves as the 5G Cell Site Router (CSR). We conducted thorough analysis of both functional and performance aspects, specifically examining Fronthaul services and Class of Service (CoS) operations.

Using the reference network design, we validated that the ACX7024 is a reliable choice for a CSR, offering an enhanced feature-set and improved performance compared to previous ACX platforms in most situations. It is specifically designed for the CSR role, catering to the scale, bandwidth, and performance requirements associated with this function.

For the full test report with all configuration files, test bed details, multidimensional scale and performance data, contact your Juniper Networks representative.

Solution Benefits

Juniper's ACX7000 series are specifically designed for use as CSRs in 4G and 5G networks. They provide the necessary connectivity and routing capabilities at cell sites to enable seamless communication between the radio access network (RAN) and the core network. The ACX Series routers offer advanced features tailored for mobile backhaul (MBH) and xHaul applications. They support high-speed interfaces, such as Ethernet and optical interfaces, to handle the bandwidth requirements of modern cellular

networks. These routers are designed to handle the challenges of high-volume traffic, low latency, and strict quality of service (QoS) requirements associated with 4G and 5G deployments.

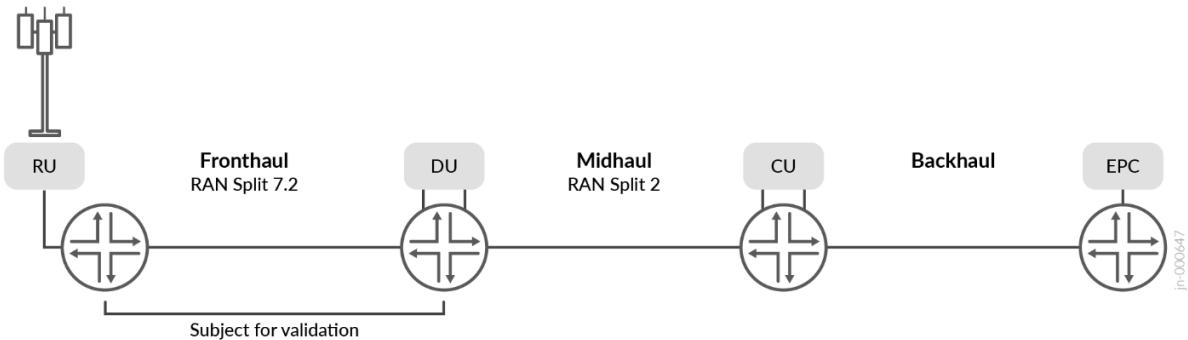
ACX Series routers provide scalability, security features, and advanced traffic management capabilities, making them suitable for diverse deployment scenarios. They can handle a range of services, including Fronthaul, Midhaul, and Backhaul, while ensuring efficient traffic flow, service prioritization, and network resilience.

- **Flexibility:** The seamless MPLS segment routing framework allows for flexible routing and traffic management, ensuring efficient allocation of network resources. This flexibility is crucial for optimizing network performance and accommodating diverse service requirements in a 5G environment.
- **Low Latency:** 5G networks require ultra-low latency to support real-time applications such as autonomous vehicles, remote surgery, and virtual reality. ACX7000 Cloud Metro Routers assure an ultra-low latency forwarding while Segment Routing MPLS (SR-MPLS) provides an optimized latency based path forwarding in the network eliminating the need for complex protocol processing.
- **Scalability:** The 5G network infrastructure needs to support a massive number of devices and provide seamless connectivity. SR-MPLS segment routing enables network scalability by simplifying the forwarding plane and reducing the control plane complexity. This allows for efficient resource utilization and optimized network performance.
- **Traffic Engineering:** Seamless MPLS segment routing provides advanced traffic engineering capabilities, enabling operators to dynamically control and optimize the flow of traffic. This allows for efficient load balancing, congestion avoidance, and QoS management, which are essential for delivering high-performance 5G services.
- **Network Resiliency:** MPLS segment routing offers fast rerouting mechanisms and supports protection and restoration schemes, allowing the network to quickly recover from failures and maintain service continuity.
- **Simplified Operations:** MPLS segment routing simplifies network operations by leveraging source routing. Instead of maintaining complex routing tables at every network node, the forwarding path is explicitly encoded in the packet header. This simplifies network configuration, reduces control plane overhead, and improves overall operational efficiency.

Use Case and Reference Architecture

The 5G xHaul architecture encompasses three physical segments referenced as Fronthaul, Midhaul, and Backhaul. Refer to [Figure 1 on page 3](#).

Figure 1: 5G xHaul Reference Network

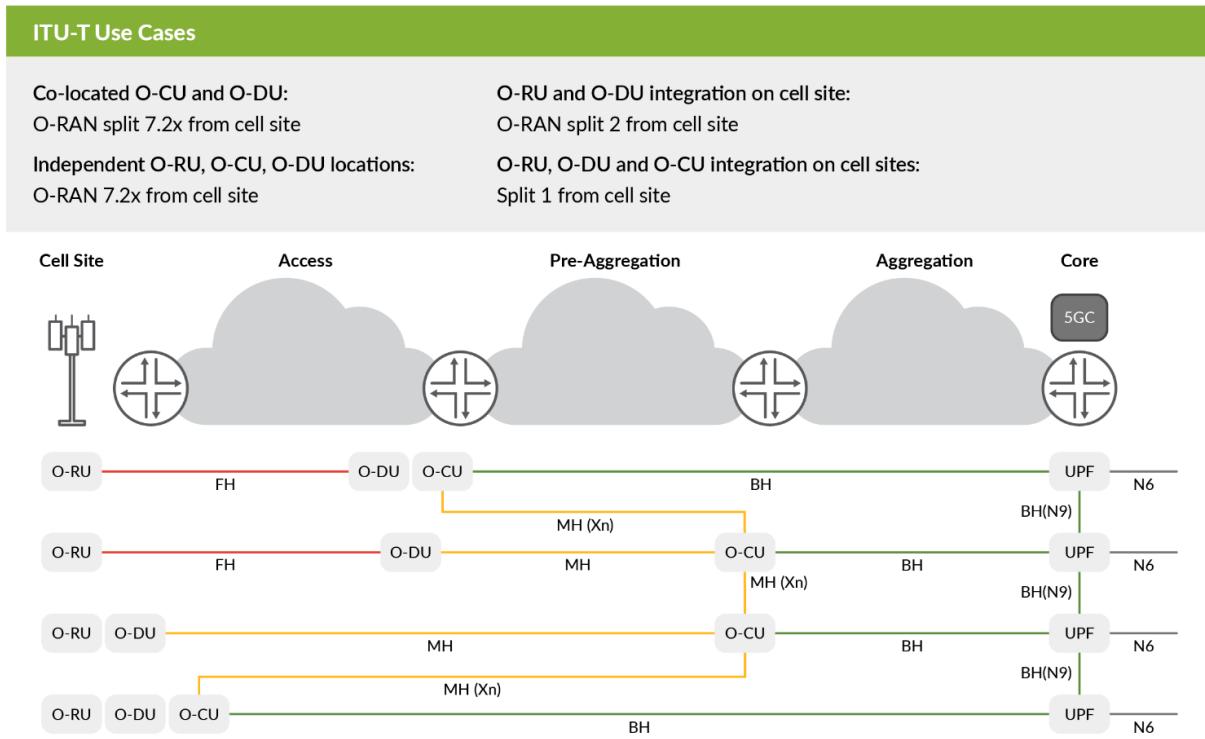


Fronthaul segment enables Layer 2 connectivity between the Open Radio Unit (O-RU) and Open Distributed Unit (O-DU) (shown as RU and DU in the [Figure 1 on page 3](#)) in the RAN. This allows them to communicate for control, data, and management traffic to ensure time and frequency synchronization between 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.

The advancement of the RAN involves different architectures for 4G, including distributed, centralized, and virtual setups, which need to coexist with the 5G disaggregated O-RAN. These diverse ecosystems provide flexibility for the placement of components such as O-DU and O-CU. This JVD does not cover all possible scenarios but closely aligns with O-RAN split 7.2x, where the O-RU connects to the CSR and the O-DU is located within the HSR infrastructure. If needed, additional insertion points can be implemented to support disaggregation between the Midhaul and Backhaul segments by extending appropriate services.

[Figure 2 on page 4](#) summarizes the deployment scenarios for the RAN according to the ITU-T for simultaneous support of 4G and 5G as proposed by the O-RAN Alliance.

Figure 2: RAN Deployment Scenarios for Simultaneous Support of 4G and 5G



The ACX7024 Universal Cloud Metro Router is ideally built for the CSR role, supporting 24 ports of 1/10/25 GbE and 4x100 GbE shared across 360 Gbps system throughput. The use case under consideration is the insertion of the ACX7024 as part of the 5G xHaul solution ([5G Fronthaul Network Using Seamless MPLS Segment Routing](#) and [5G Fronthaul Class of Service](#) JVDs) and validation of its scale, performance, and functional capabilities when used as a CSR in the 4G/5G Fronthaul network.

Solution Design and Architecture

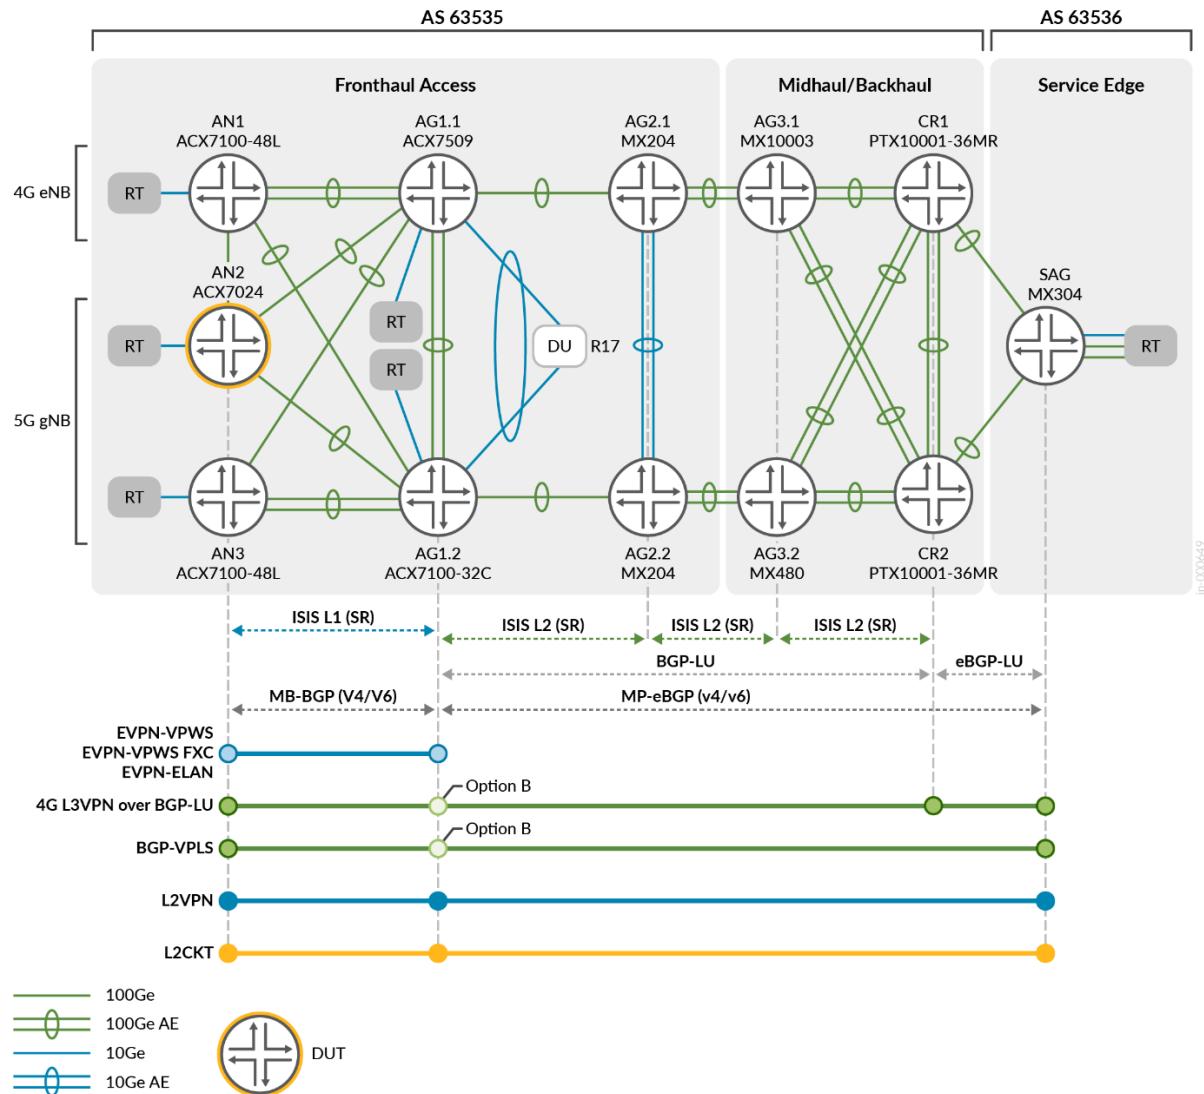
IN THIS SECTION

- [Overlay Services | 7](#)
- [Connectivity Models | 8](#)
- [Layer 3 Connectivity Models | 9](#)
- [5G QoS Identifier \(5QI\) Model | 10](#)
- [QoS Profiles | 12](#)

- Class of Service Building Blocks | [13](#)
- Service Carve Out | [15](#)
- VLAN Operations | [15](#)

[Figure 3 on page 6](#) shows an end-to-end 5G xHaul network, modeled after common topology [O-RAN.WG9.XPSAAS-v02.00], which defines four segments of transport infrastructure: access, pre-aggregation, aggregation, and transport core. Foundational technologies incorporate modern and legacy VPN services over segment routing.

Figure 3: 5G Fronthaul Services Topology



The Fronthaul network deployment scenarios were carefully designed to support both the traditional 4G MBH 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 network underlay features SR-MPLS across multiple ISIS domains and inter-AS. Access nodes are placed into an ISIS L1 domain with adjacencies to L1/L2 HSR nodes where L2 domain extends from aggregation to core segments. Seamless MPLS is achieved by enabling BGP Labeled Unicast (BGP-LU) at border nodes.

Table 1 on page 7 summarizes the choice of protocols.

Table 1: Transport Layer

	Fronthaul	Midhaul/Backhaul	SAG
IGP	ISIS L1/L2	ISIS L2	-
Intra-Domain MPLS Tunnel	SR-ISIS	SR-ISIS	-
Protection	TI-LFA	TI-LFA	-
Inter-Domain Transport	-	BGP-LU (Option B)	BGP-LU (Option C)

To handle the increased network scale, 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 in the network use different combinations of VLAN operations. These operations are applied to various Layer 2 service types such as EVPN-ELAN, EVPN-VPWS, EVPN-FXC, L2Circuit, VPLS, and L2VPN. Starting with Junos OS Evolved Release 22.3R1, Flow Aware Transport Pseudowire Label (FAT-PW) is supported 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 IPv6 PE functionality.

The following combination of VPNs is designed in a way to allow following traffic flows in the 5G xHaul network:

- Layer 2 eCPRI (emulated) between O-RU to O-DU traffic flows%5G Fronthaul
- 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 3 IP packet flows between 5G O-DU and CU/EPC%5G Midhaul and Backhaul
- Layer 2 Midhaul flows emulating additional attachment segments%4G Midhaul and Backhaul

Connectivity Models

There are two connectivity models between O-RU and O-DU. These models leverage the following EVPN-VPWS, EVPN-FXC or EVPN-ELAN services:

1. EVPN-VPWS single-homed supporting dedicated MAC for eCPRI without redundancy
2. EVPN-FXC VLAN-AWARE single-homed supporting dedicated MAC for eCPRI without redundancy
3. EVPN-VPWS with A/A ESI LAG DU attachment
4. EVPN-FXC VLAN-AWARE with A/A LAG DU attachment
5. EVPN-ELAN with A/A ESI LAG DU attachment

[Figure 4 on page 8](#) illustrates the first connectivity model. In this scenario, the network utilizes EVPN-VPWS single-homing connectivity. This setup supports dedicated MAC for eCPRI without redundancy. Additionally, it uses Ethernet OAM with performance monitoring. However, it is important to note that Ethernet OAM with performance monitoring is only supported for the single-homed configuration in this model.

Figure 4: O-RAN Fronthaul Single-Homed EVPN-VPWS/FXC

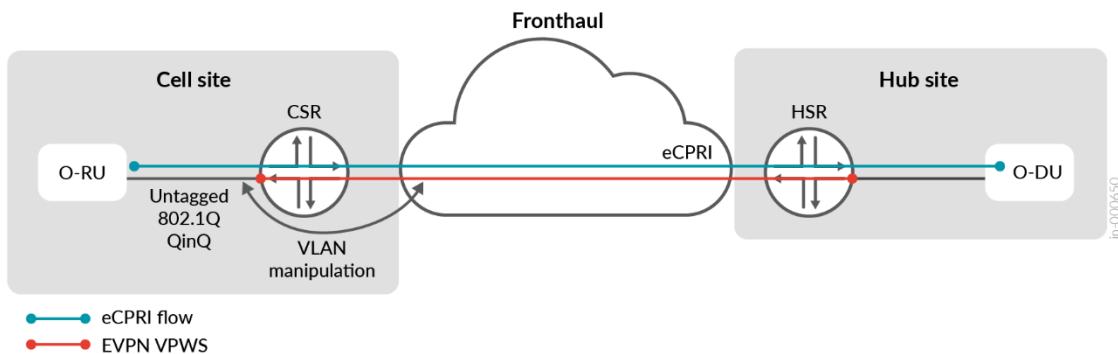


Figure 5: O-RAN Fronthaul A/A EVPN-VPWS/FXC/ELAN

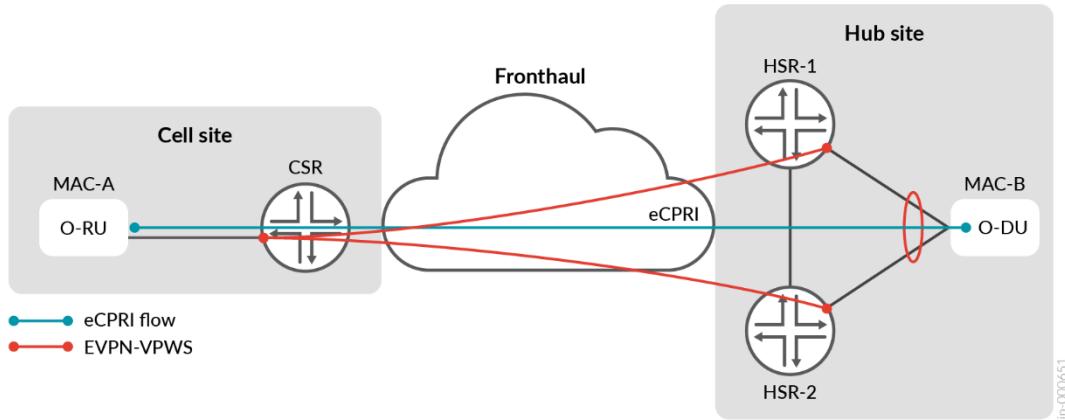


Figure 5 on page 9 illustrates the second connectivity model. This model uses either EVPN-VPWS or EVPN-ELAN with active/active multihoming. Additionally, it uses EVPN-VPWS with FXC active/active multihoming from CSR (AN4) to HSR (AG1.1/AG1.2). The HSR devices are connected to the O-DU through an active/active Ethernet Segment Identifier (ESI) Link Aggregation Group (LAG), enabling the sharing of traffic load. The links are bundled into an active/active EVPN ESI 10Ge LAG between AG1.1 and AG1.2, as well as to the O-DU, which consists of a two-member Aggregate Ethernet (AE) with both links actively functioning. In this configuration, eCPRI packets might arrive on either O-DU link from the HSRs, while eCPRI packets are transmitted across either HSR uplink for active/active operations.

To enable traffic load sharing, an active/active ESI LAG is established between the HSRs and the O-DU. This allows for balanced distribution of traffic. The links are bundled into an active/active EVPN ESI 10G Ethernet LAG between HSR-1 and HSR-2, as well as to the O-DU. The O-DU includes a two-member AE with both links actively carrying traffic.

Layer 3 Connectivity Models

We chose L3VPN protocol to facilitate Layer 3 connectivity between O-DU and vCU/vEPC elements of the 5G xHaul. Two unique connectivity are proposed, with both supporting Layer 3 multihoming between O-DU and pair of HSRs:

- EVPN IRB anycast gateway with L3VPN
- BD with IRB and static MAC/ARP with L3VPN

The two corresponding models are referred to as *EVPN IRB with L3VPN* and *BD IRB with L3VPN*, respectively, see Figure 6 on page 10 and Figure 7 on page 10. For more details about configurations for these connectivity models, contact your Juniper Networks representative.

Figure 6: EVPN IRB Anycast Gateway with L3VPN

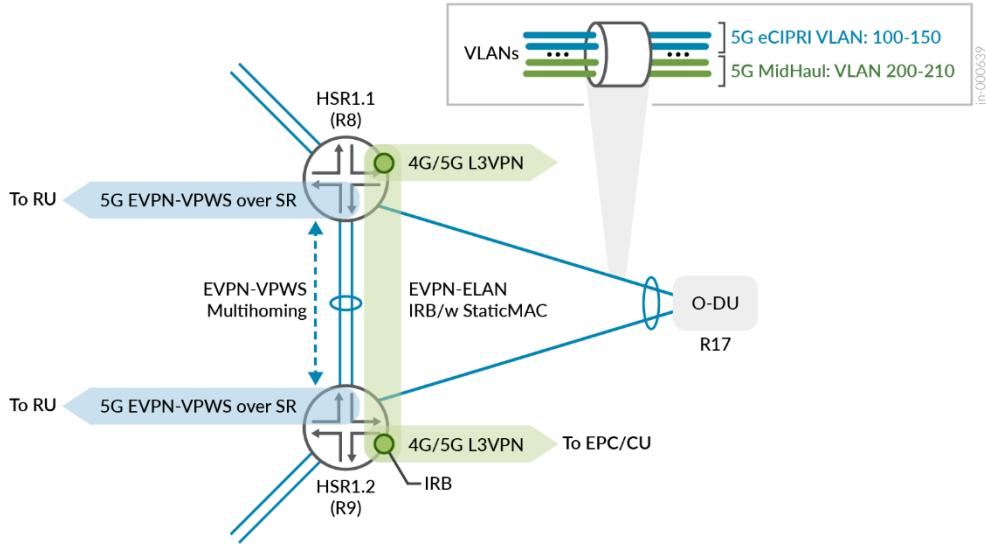
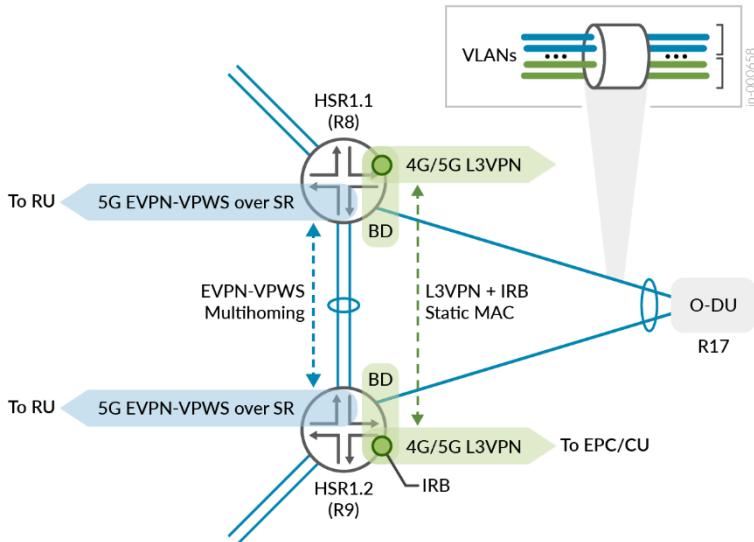


Figure 7: BD with IRB and Static MAC/ARP with L3VPN



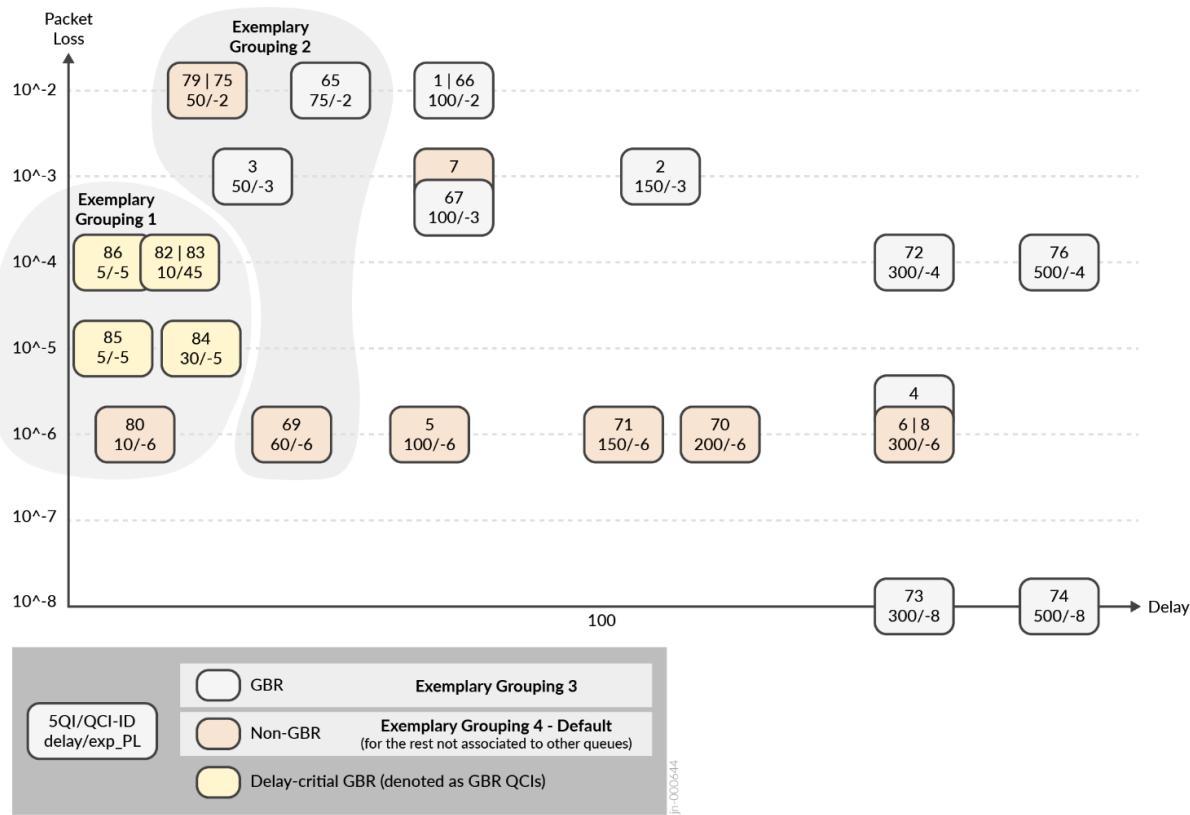
5G QoS Identifier (5QI) Model

When transitioning from the 4G LTE Quality of Service Class Identifier (QCI) model to the flow-based 5G QoS Identifier (5QI) model, most traffic definitions overlap. However, 5G introduces new categories

for delay-critical Guaranteed Bit Rate (GBR) flows. In the 5G Fronthaul segment, eCPRI-based flows handle user and control traffic between the O-RU and O-DU. These flows require high bandwidth and extremely low delay. Therefore, all devices in the access topology must prioritize this traffic type with the highest priority.

The O-RAN specification [O-RAN.WG9.XPSAAS-v02.00] proposes a model to group common QCI and 5QI flow characteristics into four exemplary groups based on their delay budget. This grouping aims to provide a framework for defining the QoS for different types of traffic in the 5G network. Refer to [Figure 8 on page 11](#).

Figure 8: O-RAN 5QI/QCI Exemplary Grouping



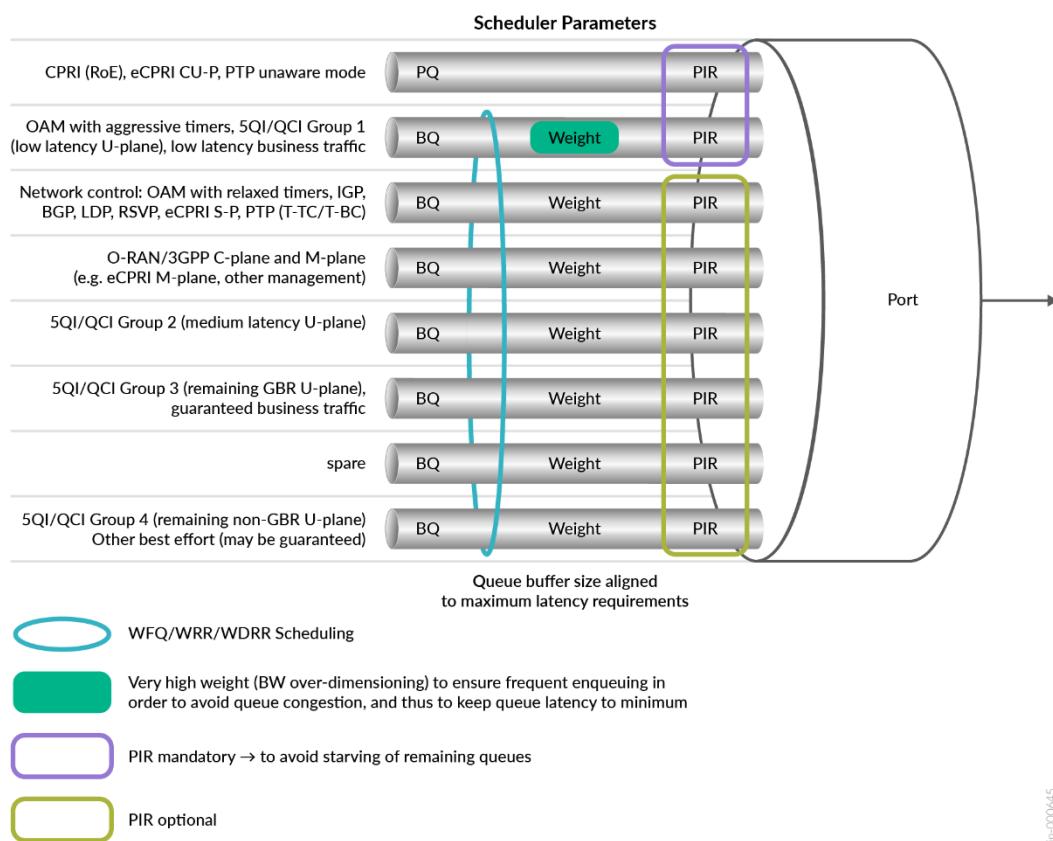
QoS schemas can differ among mobile operators, and this JVD does not endorse a specific design as the recommended one. The objective is to establish predictable behaviors for critical and non-critical traffic flows across various services delivered by the xHaul network. The transport architecture needs to demonstrate its capability to accommodate existing and emerging mobile applications while maintaining the integrity of the delay budget and ensuring traffic priorities.

For more details on the specific latency and delay budgets considered for this JVD, contact your Juniper Networks representative.

QoS Profiles

O-RAN/3GPP proposes two common QoS profiles to meet the requirements of the transport network. In Profile A, illustrated in [Figure 9 on page 12](#), a single priority queue is defined to handle ultra-low latency flows such as Precision Time Protocol (PTP) and eCPRI. This queue is given priority over all other queues. Lower priority queues are then serviced using weighted fair queuing (WFQ) round-robin scheduling. The ACX7000 series are best suited for Profile A.

Figure 9: Single Priority Queue (Profile A)



jh-000645

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.

As of Junos OS Evolved Release 22.3R2, 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

Class of Service (CoS) governs how traffic is forwarded, stored, or dropped in conjunction with mechanisms to manage and avoid congestion. CoS is comprised with the following basic building blocks:

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

CoS models differ between operators based on unique traffic profiles and characteristics. [Table 2 on page 13](#) defines a pseudo-customer model that we used for this JVD. For more details on this CoS model, contact your Juniper Networks representative.

Table 2: Validated Scheduling Profiles

Forwarding Classes	Scheduling Parameters				Classification & Rewrite				Traffic Profile			QCI/5C I Mapping
	Queue	Priority	Transmit rate	Buffer size	802.1p	DSCP	MPLS EXP	Packet Loss Priority	Resource Type	Traffic Type		
Business	5	Low	20%	20%	4	CS4, AF4x	4	Low	GBR	Guaranteed U-Plane Business Conversational Real Time Gaming/Video		QCI1-4, 6
Network Control	4	Low	5%	2%	7	CS7	7	Low	GBR	Protocol, Timing		QCI65-67
Real Time	2	Strict High	40%	30%	5	CS5	5	Low	Delay-Critical GBR	eCPRI		CPRI QCI82-90
			Shaped			EF						

Table 2: Validated Scheduling Profiles (*Continued*)

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

We validated two styles of ingress classification:

- 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 are 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 a maximum of eight queues per interface. All platforms represented support eight queues in total. For this JVD, we used six queues and associated forwarding classes to accommodate the traffic scheme requirements. For Profile-A, we used only one strict-high queue, which is shaped (PIR) to prevent starving low priority queues. We configured all other queues as low priority and serviced as weighted fair queuing (WFQ) 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. The ACX series supports rewriting only the outer tag, which is the default. In most cases, it is preferred to preserve and transmit the inner (C-TAG) 802.1p bits transparently.

Service Carve Out

As a best practice, ultra-low latency services (eCPRI) are assigned the highest priority. MBH applications might have varying treatments. [Table 3 on page 15](#) lists the priority mappings that we have used for this JVD, grouped by service type.

Table 3: 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

The ACX7000 series supports a comprehensive set of VLAN manipulation operations compared to previous generation ACX platforms. This JVD doesn't include all possible permutations, but does validate 80 VLAN combinations across L2Circuit, L2VPN, EVPN-VPWS, and EVPN-ELAN 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 4 on page 16](#) summarizes the explicit VLAN normalization operations that we validated for each Layer 2 VPN type. For the comprehensive test report for each operation, contact your Juniper Networks representative.

Table 4: 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

Table 4: Validated VLAN Operations (*Continued*)

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

Solution and Validation Key Parameters

IN THIS SECTION

- [Supported Platforms | 18](#)
- [Service Profiles | 18](#)
- [Scale and Performance | 19](#)
- [Key Feature List | 21](#)
- [Test Bed | 22](#)
- [Solution Validation Goals | 25](#)
- [Class of Service Validation Points | 26](#)
- [Solution Validation Non-Goals | 28](#)
- [Failure Scenarios | 28](#)

This section outlines solution key parameters and validation objectives for this JVD.

Supported Platforms

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.

Service Profiles

[Table 5 on page 18](#) and [Table 6 on page 19](#) show the list of Fronthaul and Midhaul service profiles respectively and associated network services which were used during validation. Note that Fronthaul profiles were in the focus of the validation, while Midhaul profiles and associated traffic flows were used for sake of completeness. The Fronthaul profiles were the focus of the validation, Midhaul profiles and their associated traffic patterns were used to ensure validation completeness.

Table 5: Fronthaul Service Profiles

Use Case	Service Overlay Mapping	End Points
4G L3VPN MBH	End-to-End L3VPN between CSR (AN4) to SAG	AN4/SAG IPv4; IPv6
5G Fronthaul	Fronthaul EVPN-VPWS + FXC single-homing from AN4 to HSRs (AG1) with E-OAM Performance Monitoring	AN4/AG1 Untagged, Single/Dual Tag
5G Fronthaul	Fronthaul EVPN-VPWS + FXC with Active/Active Multihoming from AN4 to HSRs (AG1)	AN4/AG1 Untagged, Single/Dual Tag
5G Fronthaul	Fronthaul EVPN-ELAN with Active/Active Multihoming from AN4 to HSRs (AG1)	AN4/AG1 Untagged, Single/Dual Tag
L2VPN MBH	End-to-End L2VPN between CSR (AN4) to SAG with FAT-PW	AN4/SAG Untagged, Single/Dual Tag
L2Circuit MBH	End-to-End L2Circuit between CSR (AN4) to SAG with FAT-PW	AN4/SAG Untagged, Single/Dual Tag

Table 5: Fronthaul Service Profiles (*Continued*)

Use Case	Service Overlay Mapping	End Points
BGP-VPLS MBH	End-to-End VPLS between CSR (AN4) to SAG with E-OAM Performance Monitoring	AN4/SAG
		Untagged, Single/Dual Tag

Table 6: Midhaul Service Profiles

Use Case	Service Overlay Mapping	End Points
5G Midhaul	EVPN IRB anycast gateway with L3VPN multi-homing	AG1/SAG IPv4; IPv6
5G Fronthaul	Bridge Domain IRB anycast static MAC/IP with L3VPN	AN4/AG1 Untagged, Single/Dual Tag
L2VPN Midhaul	Midhaul L2VPN HSR (AG) attachments (AG1) to SAG with FAT-PW	AG1/SAG
L2Circuit Midhaul	Midhaul L2Circuit attachments between HSR (AG) to SAG with FAT-PW	AG/SAG

Scale and Performance

This section contains key performance indexes (KPIs) used in solution validation targets. Validated KPIs are multi-dimensional and reflect our observations in customer networks or reasonably represent solution capabilities. These numbers do not indicate the maximum scale and performance of individual tested devices. For uni-dimensional data on individual SKUs, contact your Juniper Networks representatives.

The Juniper JVD team continuously strives to enhance solution capabilities. Consequently, solution KPIs may change without prior notice. Always refer to the latest JVD test report for up-to-date solution KPIs. For the latest comprehensive test report, contact your Juniper Networks representative.

The scale reference in [Table 8 on page 20](#) provides an overview of KPIs represented in the validated profile.

To validate CoS functionality, we tested the classification, scheduling, shaping, and rewriting behaviors of the ACX7024 across services utilizing the 5G xHaul infrastructure. As part of the testing, we measured the latency for critical Fronthaul traffic types.

Based on the network design, the architecture can deliver fast restoration within 50ms for most traffic flows transported over ISIS-SR with Topology Independent Loop-Free Alternate (TI-LFA) protection mechanisms. Load distribution and optimization features were shown to improve service restoration in the event of link or node failures. Link events consistently achieved convergence in less than 50ms. The ACX7024 with Junos OS Evolved Release 22.3R2 can deliver the solutions outlined here across intra- and inter-domain architectures and is ideally situated for the CSR access role.

Table 7: KPI Scale Summary

Feature	AN4 (ACX7024)–Access / CSR	AG1.1 (ACX7509)– Pre-Agg / HSR	AG1.2 (ACX7100-32C)– Pre-Agg / HSR	SAG (MX10003)– Services Agg
RIB/FIB	200k/100k	400k/375k	400k/375k	640k/430k
IFLs	1498	11145	11010	16288
EVPN-VPWS SH	200	700	700	0
EVPN-VPWS MH A/A	100	200	200	0
EVPN-FXC SH	50	50	0	0
EVPN-FXC MH	50	50	50	0
EVPN-ELAN	50	50	50	0
L2Circuit	100	1000	1000	2500
L2VPN	50	1000	1000	2450
L3VPN	100	100	100	100
VPLS	100	1000	1000	2500

Table 7: KPI Scale Summary (Continued)

Feature	AN4 (ACX7024)–Access / CSR	AG1.1 (ACX7509)–Pre-Agg / HSR	AG1.2 (ACX7100-32C)–Pre-Agg / HSR	SAG (MX10003)–Services Agg
L3VPN BD (Midhaul)	0	500	500	500
MAC (VPLS)	10k	29k	111k	176k
CFM UP MEP (1s)	300	100	100	100
Total VPN Services	800	4650	4600	8050

Key Feature List

- EVPN-VPWS
- EVPN-ELAN
- EVPN-FXC
- L3VPN
- BGP-VPLS
- L2Circuit
- L2VPN
- Segment Routing ISIS
- TI-LFA (link/node)
- ISIS
- BGP
- BGP-LU
- BFD

- Community-based Routing Policy
- Route Reflection
- IPv4
- IPv6
- LACP
- AE
- CFM
- LFM
- VLAN (802.1q)

For the full test report and feature list, contact your Juniper Networks representative.

Test Bed

[Figure 10 on page 23](#) illustrates the test bed that we used. The network consists of four layers: access, pre-aggregation, aggregation, and transport core.

- **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. This JVD does not focus on these segments.

Figure 10: 5G Fronthaul Lab Topology

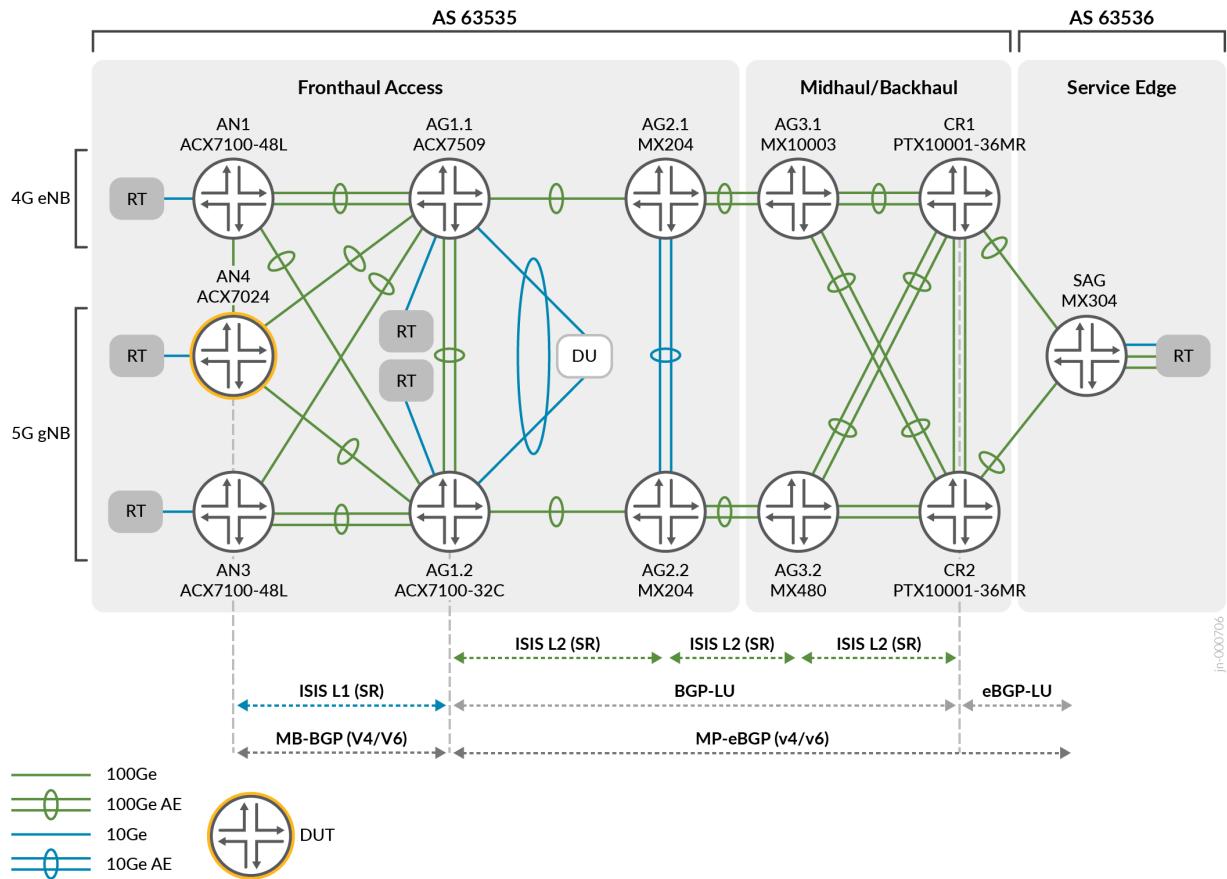
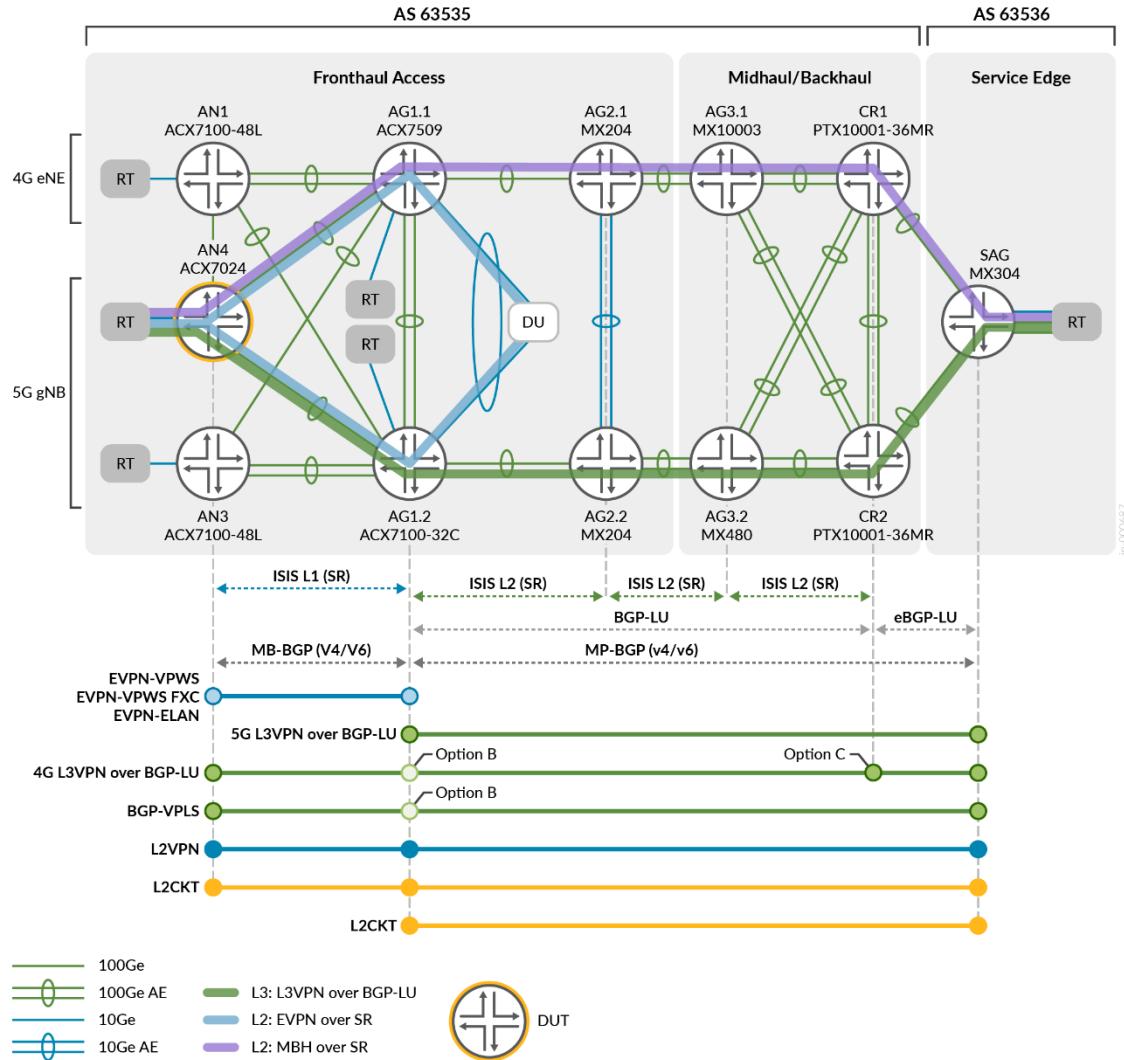


Table 8 on page 23 lists the topology definitions.

Table 8: Topology Definitions

Layer	Devices Under Test
Access	ACX7100-48L (AN3), ACX7100-48L (AN1), ACX710 (AN2) CSRs
Pre-Aggregation	ACX7509 (AG1.1) and ACX7100-32C (AG1.2) HSRs
Aggregation	MX204s (AG2.1/AG2.2), MX10003 (AG3.1), MX480 (AG3.2) aggregation routers
Core Network	PTX1000 (CR1) and MX10003 (CR2) core routers. MX10003 (SAG) services router

Figure 11: End-to-End 4G/5G Traffic Flows and Network Architecture



The flows are generated in the same way from AN4 (ACX7024) towards both AG1.1 and AG1.2, as well as AN4 to SAG. Load sharing is applied whenever possible. The network paths are chosen based on IGP metrics. The packet sizes for most VPN services range from 128 to 1000 bytes.

For additional details on validation scenarios and full archive of the test bed configuration used for this JVD, contact your Juniper Networks representative.

Solution Validation Goals

The main goal was to validate the reference design of a unified 5G xHaul network with a specific focus on the Fronthaul segment. To achieve this, we used Seamless MPLS over ISIS Segment Routing (ISIS-SR), enabling the support of multiple 4G/5G services including:

- VLAN-aware services including L3VPN (IPv4 and IPv6 virtual private networks)
- Active-Active Multihoming for EVPN-ELAN
- EVPN-VPWS and EVPN Flexible Cross Connect (FXC) VLAN-aware services
- Single-homed services such as EVPN-VPWS, EVPN-FXC, BGP-Virtual Private LAN Service (VPLS), Layer 2 Virtual Private Network (L2VPN), and L2Circuit

Here are the major test goals for this JVD:

- Validate VPN services, including L3VPN, EVPN-VPWS, EVPN-FXC, EVPN-ELAN, BGP-VPLS, L2Circuit, and L2VPN over SR-MPLS transport architecture.
- Validate TI-LFA redundancy mechanisms over Segment Routing with Seamless MPLS/BGP-LU.
- Validate network resiliency, traffic restoration, and measured convergence time for ACX7024 (AN4) with adjacent link failures for all traffic types.
- Measure solution resilience of Layer 2 and Layer 3 flows from Access Node (AN) to Pre-Aggregation AG1 (O-RU to O-DU).
- Validate input/output VLAN operations for the normalization of all VPN services.
- Validate the basic mechanisms of CoS:
 - Classification of traffic based on DSCP, 802.1p and EXP with Packet Loss Priority (PLP) high and low.
 - Preservation of QoS codepoints end-to-end for inner and outer tags.
 - Support for ingress classification using fixed and behavior aggregate styles.
 - Creation of at least six forwarding classes and six queues (all featured platforms support eight queues).
 - Support for a two-priority queue scheduling system, consisting of a strict-high priority and a low priority. The system should allocate a certain percentage of time and buffers to each priority queue (Traffic Rate).
 - Strict-high priority queues pre-empt low priority queues.
 - Strict-high priority queue shaping prevents starving low priority queues.

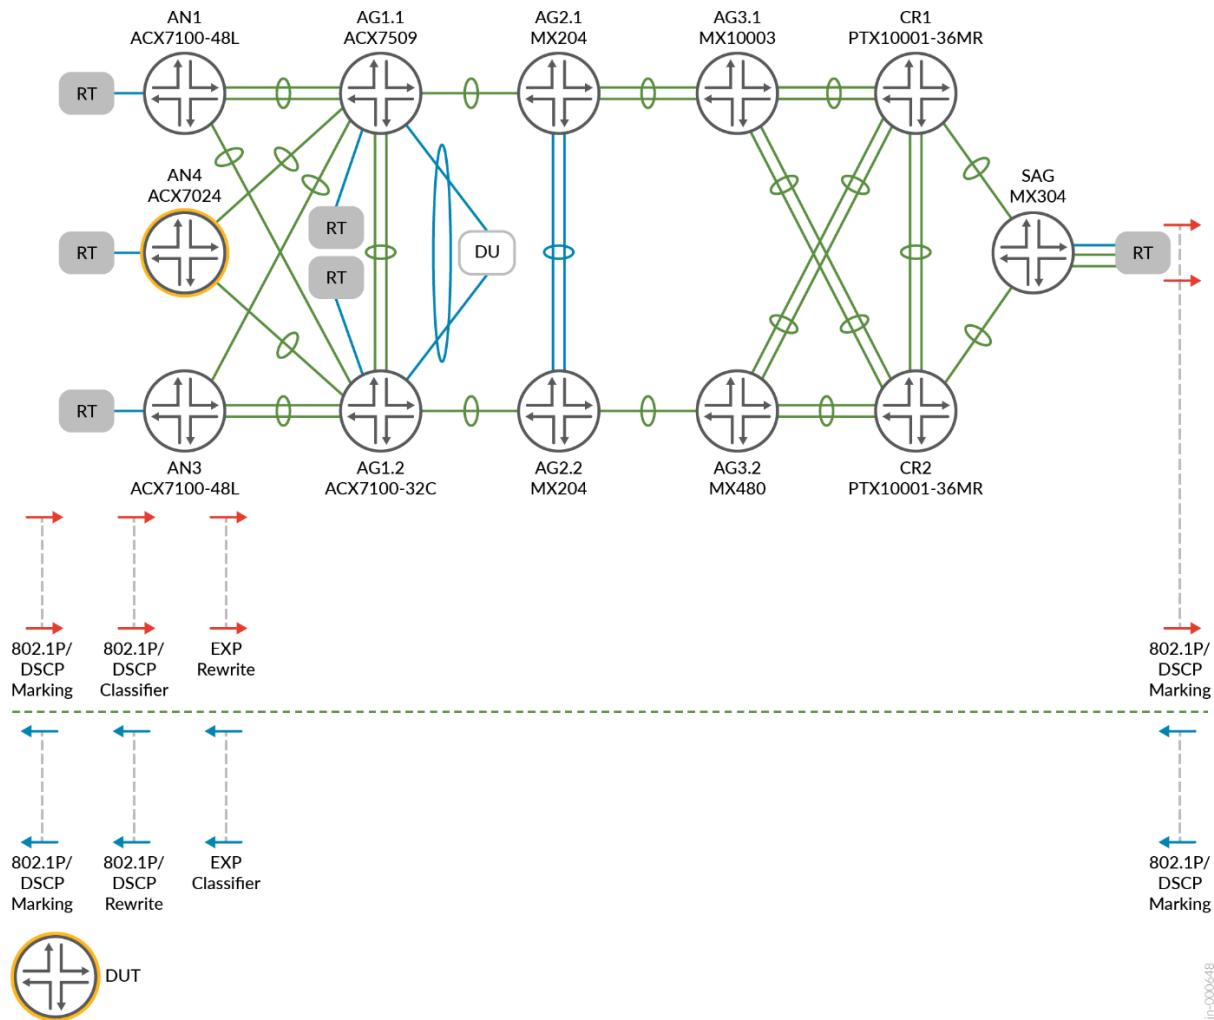
- The port shaper inherits the scheduler characteristics.
- Rewrite operations, based on queue assignment, support 802.1p, DSCP and EXP.
- Rewrite for single-tagged and dual-tagged (outer only) frames.
- Validate latency budgets for non-congested scenarios where <100% line rate is offered while strict-high queue is in-profile:
- O-RU-to-O-DU latency averages $\leq 10\mu$ per device ($\leq 6\mu$ s single DUT).
- RU-to-SAG latency is $\leq 10\text{ms}$ (expected $\leq 150\mu$).
- Validate congestion scenarios:
 - Preservation of highest priority (eCPRI) Fronthaul traffic.
 - Traffic priorities are maintained across shared links.
 - Traffic priorities are maintained within and between VPN services that share common links.
- Validate consistency and resiliency of the ACX7024 against negative stress conditions (enabled/disable control and data plane daemons, add/delete configurations, and so on.)
- Identify product limitations, anomalies, and open Problem Reports (PRs) exposed during validation stages.
- Attempt to resolve and verify opened PRs during validation.

Class of Service Validation Points

We tested CoS operations and performance requirements to maintain the reliability of important 5G Fronthaul traffic between RU and DU. In [Figure 12 on page 27](#), the DUTs are:

- ACX7024 as the CSR to facilitate traffic flows
- PTX10001-36MR for the core and peering role
- MX304 as the services edge platform

Figure 12: Class of Service Validation Points



In Figure 12 on page 27, the traffic flows from IXIA (RT) are directed through the ACX7024 (AN4) towards the O-DU or SAG (Services Aggregation Gateway). These flows are classified based on Layer 2 (802.1p) or Layer 3 (DSCP) codepoints at specific positions called classifiers. The codepoints are then mapped to EXP values across the SR-MPLS topology.

To ensure the expected behavior, queue statistics are monitored to confirm that the classification and scheduling process yields the desired outcomes. Additionally, rewrite operations are performed at designated positions to modify certain packet fields. Packet captures are taken to verify that DSCP, 802.1p, or EXP bits are correctly rewritten or preserved.

In the opposite direction, flows sent through the SAG are marked and validated once they exit the AN4, ensuring that the marking process operates as intended.

For the full test report including complete details on the hardware and software, contact your Juniper Networks representative.

Solution Validation Non-Goals

Non-goals represent protocols and technologies outside the scope of the current validation.

- Underlay MPLS/SR transport other than specified in the Solution Validation Goals section
- Latency validation under congestion scenarios
- Temporal transmit rate or buffer (elastic buffer is used)
- BGP PIC-Edge at border routers
- Multifield classification to forwarding class mapping
- Custom drop profiles (WRED) (defaults are used)
- Hierarchical CoS and Traffic Control Profiles, IFD/IFL policers
- 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

- DUT Adjacent link failures
- DUT Indirect link failures
- DUT Node failures
- Link congestion
- Queue congestion
- Process restart

Results Summary and Analysis

IN THIS SECTION

- [ACX7024 Functions and Performance | 29](#)
- [ECMP Load Balancing | 29](#)
- [Network Convergence | 31](#)
- [Class of Service Validation | 31](#)
- [Congestion Scenarios | 34](#)
- [Latency Budgets | 35](#)

ACX7024 Functions and Performance

During the validation process, we successfully demonstrated a robust solution for 5G xHaul transport infrastructure using Seamless MPLS with Segment Routing. The JVD achieved a reasonable scale of L2/L3 connectivity services, meeting the expectations of Mobile Network Operators (MNOs) and Metropolitan Area Network (MAN) operators for real network deployments. The solution also met stringent Service Level Agreement (SLA) requirements.

ECMP Load Balancing

The design reduces traffic impact during link/node failure events by enabling load sharing ECMP operations across all devices. Several ECMP mechanisms were configured (as supported) including adjusting IGP metrics, BGP multipath, ECMP fast-reroute, and VPN-unequal-cost for L3VPN services. In addition, FAT-PW label is enabled on the ACX7000 series for L2VPN and L2Circuits. EVPN FAT-PW is supported starting in Junos OS Evolved Release 23.1R1.

For a copy of the full test report, including details on hash-keys enabled for this validation and traffic load sharing limitations, contact your Juniper Networks representative.

Table 9: ACX7024 ECMP Summary Results

ACX7000 ECMP Load Balancing Performance [*]							
Service: DUT (Traffic Path)	ECMP Links	Flow	FAT- PW	Link1	Service: DUT (Traffic Path)	ECMP Links	Flow
L2VPN: ACX7024 (AN4 to SAG)	4	100kfps	Y	ae22	ae25(1)	et-0/0/5	et-0/0/6
				20.2kp ps	30.3kpps	30.3kpps	20.2kpps
L2CKT: ACX7024 (AN4 to SAG)	3	100kfps	Y	ae22	et-0/0/5	et-0/0/6	N/A
				35kpps	33kpps	32kpps	
VPLS: ACX7024 (AN4 to SAG) [1]	4	100kfps	N	ae22(1)	ae25(1)	et-0/0/2	et-0/0/3
				24.2kp ps	26.7kpps	24.2kpps	26.7kpps
EVPN: ACX7024 (AN4 to DU)	4	100kfps	N	ae22(1)	ae25(1)	et-0/0/2	et-0/0/3
				24.7kp ps	26.2kpps	24.7kpps	26.2kpps
L3VPN: ACX7024 (AN4 to AG1.1)	4	100kfps	N	ae22(1)	ae25(1)	et-0/0/2	et-0/0/3
				25.7kp ps	25.9kpps	25.7kpps	25.9kpps

[*] For complete ECMP results with all outputs, contact your account representative.

[1] Only Known Unicast is shown. VPLS BUM traffic should not load balance over ECMP routed links. Expected behavior.

In terms of ECMP performance, the ACX7024 performed similarly to the previously tested [5G Fronthaul Network Using Seamless MPLS Segment Routing](#) JVD. However, there was a slight imbalance in the distribution of L2VPN traffic due to the hash computation on the ACX7024. Similar results were observed when using three ECMP links, with the ACX7024 exhibiting a distribution of 33kpps/38kpps/30kpps, while the ACX7100-48L achieved nearly perfect balance. For a detailed report on the test results, including information on ACX7024 ECMP Load-Balancing, contact your Juniper Networks representative.

Network Convergence

Overall convergence results are within expectations for the given network design. In Fronthaul (CSR to HSR), ACX7024 failure and restoration events were well within 50ms recovery where expected. ACX7024 performance was comparable to ACX7100-48L in the CSR role with the ACX7100-48L reasonably achieving slightly better convergence results. All ACX7000 series demonstrate improved convergence compared to previous generation ACX5448/ACX710 where CSR-to-DU reported up to four seconds of traffic loss during EVPN-VPWS failure events.

[Table 10 on page 31](#) summarizes convergence performance validations across all represented VPN services, which includes single-homing or active-active multi-homing. Traffic is sent as known-unicast. Higher convergence can be expected for BUM traffic in MAC-learned services. These results are also recorded in the full test report.

Table 10: Convergence Times for 5G Fronthaul Failure Events Per Flow Type

Flow Type	EVPN-VPWS (msec)		EVPN-FXC (msec)		EVPN- ELAN (msec)		VPLS (msec)	L2VPN (msec)	L2CKT (msec)	L3VPN (msec)
Single/ Multihoming	SH	A/A MH	SH		A/A MH	A/A MH	SH	SH	SH	SH
AN4 to AG1.1 disable	18	10		10	11	0	18	21	15	
AN4 to AG1.1 enable	2	2		2	0	0	5	2	8	
AN4 to AG1.2 disable	31	20		0	20	15	0	5	22	
AN3 to AG1.2 enable	0	4		2	2	4	2	4	8	

Class of Service Validation

Across the end-to-end topology, classification and rewrite was performed on 802.1p, DSCP, and EXP as outlined in [Figure 13 on page 32](#). [Table 11 on page 33](#) 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.

Figure 13: Class of Service Functional Diagram

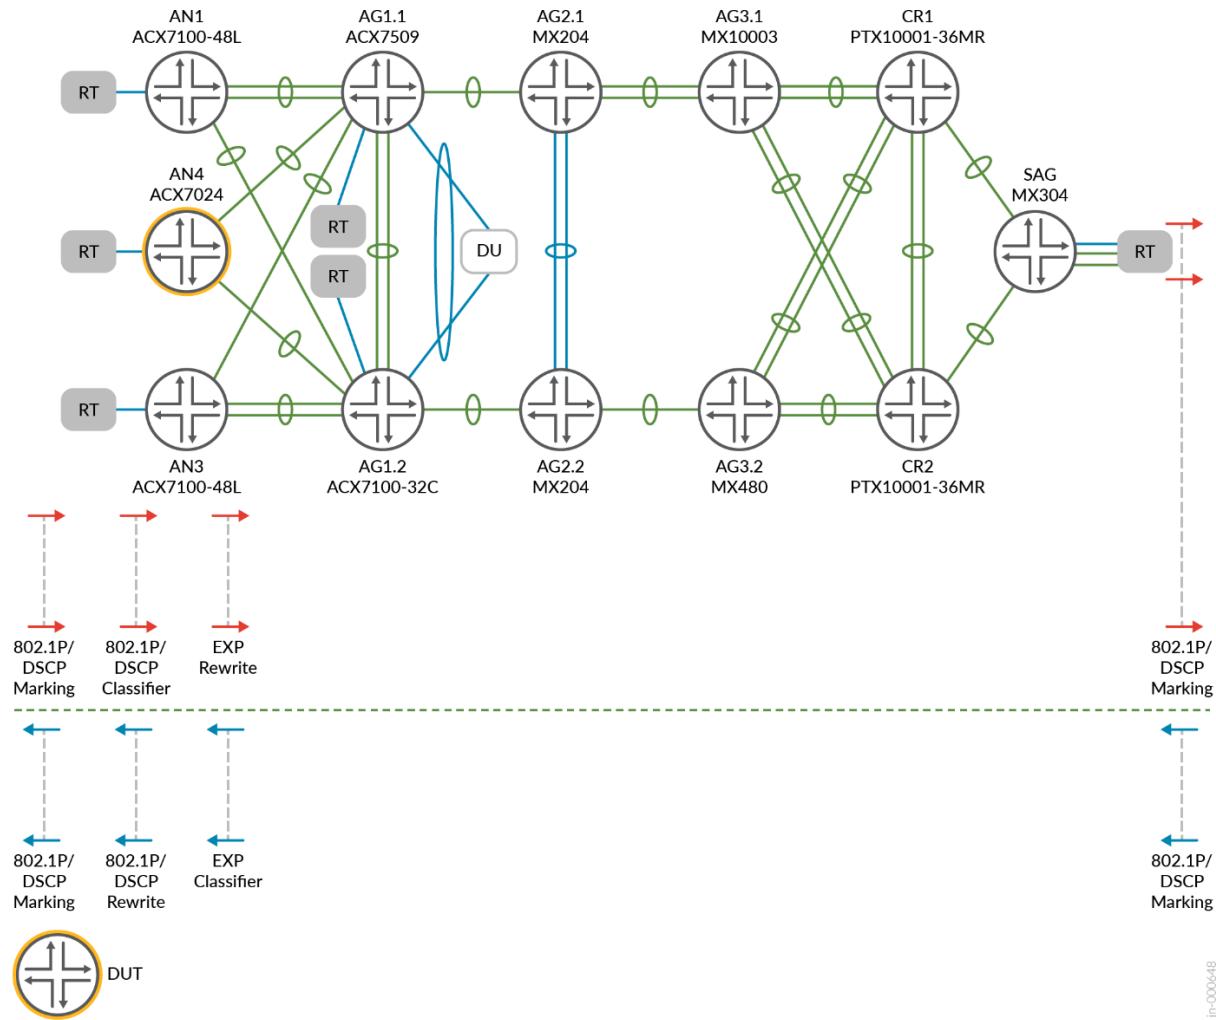


Table 11 on page 33 summarizes VLAN operation scenarios that we executed and the corresponding results. VLAN Tags are represented as Untagged (UT), Single-Tagged (ST), and Dual-Tagged (DT). All listed input/output VLAN mapping operations were validated across L2Circuit, L2VPN, EVPN-VPWS, and EVPN-ELAN services.

For the full test report, which includes an analysis explaining the results for each function, contact your Juniper Networks representative.

Table 11: CoS Summarized Results

Traffic Scenario	VLAN	Ingress Classification Mapped to FC			Scheduler Honored	Rates	Codepoints Rewritten			Bits Preserved
Fixed Classifier	TAG	802.1p	DSC P	EXP	SH	LOW	802.1p	DSC P	EXP	E2E
EVPN-VPWS	UT	--	--	✓	✓	--	✓	--	✓	✓
EVPN-ELAN	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	DSC P	EXP	SH	LOW	802.1p	DSC P	EXP	E2E
L3VPN	UT	--	✓	✓	--	✓	--	✓	✓	✓
L2VPN	UT	✓	--	✓	--	✓	✓	--	✓	✓

Table 11: CoS Summarized Results (Continued)

Traffic Scenario	VLAN	Ingress Classification Mapped to FC			Scheduler Honored	Rates	Codepoints Rewritten			Bits Preserved
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	✓	--	✓	--	✓	✓	--	✓	✓

Congestion Scenarios

The validation included various congestion scenarios outlined in the Solution Validation Goals section. Congestion constitutes one or more conditions where traffic exceeds the configured scheduler transmit-rate, shaped-rate, or port speed and results in expected traffic loss. The major objective is to ensure critical priority traffic is uninterrupted even during periods of congestion.

During key congestion events, we observed the following:

- Strict-high queue was serviced ahead of low priority queues (up to queue shaped rate). Even during periods of high congestion, when low priority queues are dropping packets, critical flows were guaranteed without any packet loss.
- Low priority queues are guaranteed up to configured transmit-rate (CIR) (strict-high queue is shaped).
- Low priority queues are serviced as WFQ when operating in excess regions and bandwidth is available.
- Low priority remainder queue was granted a transmit rate consistent with leftover bandwidth.
- Scheduler percentages correctly inherit the configured port-shaper as port speed.
- Queue shaping rate is deducted from total bandwidth with transmit-rates applied to the remaining bandwidth.
- Priority hierarchies are honored across and within VPN services that share common links.

For the full test report with details on all test cases, contact your Juniper Networks representative.

Latency Budgets

5G xHaul infrastructure defines strict latency budgets and particularly in the Fronthaul segment where supporting ultra-low latency flows are required. Total budget factors elements such as fiber length, connected devices, and transport design. O-RAN mandates a maximum of 100 μ s Fronthaul one-way latency from O-RU to O-DU, with each device \sim 10 μ s. But operations are demanding device latency closer to \sim 5-6 μ s. This is a massive paradigm shift from the requirements of earlier 4G architectures.

First, we looked at how ACX7024 performs, taking latency measurements as a standalone platform. Then, we validated how the complete Fronthaul and MBH infrastructure performs with ACX7024 as the CSR.

Topology 1, shown in [Figure 14 on page 36](#), was used to validate the performance of the ACX7024 device as the CSR. It offers the most accurate representation of the ACX7024's individual performance without considering additional hops in the network path. The traffic is generated by Ixia, excluding self-latency. We simulated critical traffic flows that represent eCPRI using burst or continuous streams, with packet sizes of 64b, 512b, and 1500b.

Figure 14: Topology 1 for Single DUT

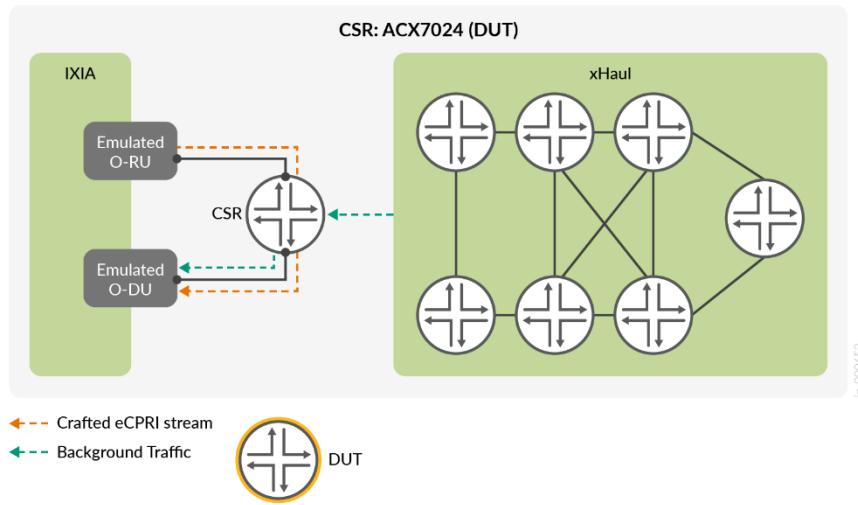


Table 12 on page 36 displays the latency measurements of the ACX7024. In this scenario, a single-DUT utilizes a bridge-domain, with traffic mapped to the strict-high queue. However, it is worth noting that there is a minimal difference observed whether the queue is set to strict-high or low when there is no congestion.

Table 12: Single DUT Latency Measurements

DUT	Queue Priority	Min (μs) Latency	Ave (μs) Latency	Max (μs) Latency	Frame Size	Traffic Pattern	Port
ACX7024	SH	5.44μs	5.46μs	5.94μs	64b	Continuous	10G
ACX7024	SH	5.33μs	5.37μs	5.84μs	512b	Continuous	10G
ACX7024	SH	4.62μs	4.65μs	6.12μs	1500b	Continuous	10G
ACX7024	SH	5.44μs	5.47μs	6.03μs	64b	Burst	10G
ACX7024	SH	5.34μs	5.37μs	5.68μs	512b	Burst	10G
ACX7024	SH	4.63μs	4.66μs	5.87μs	1500b	Burst	10G

For complete outputs of all latency measurements, contact your Juniper Networks representative.

Topology 2, shown in [Figure 15 on page 37](#), was used to measure performance across the xHaul, including both CSR and HSR devices in the Fronthaul segment. The ACX7024 is the CSR DUT and the ACX7509 is the HSR.

The Fronthaul segment consists of three hops:

1. CSR ACX7024
2. HSR ACX7509
3. O-DU QFX5110-48S/O-DU

EVPN single-homed services are between ACX7024 and ACX7509 with QFX being Layer 2 passthrough.

The Midhaul to Backhaul segment (L2Circuit and L3VPN) consists of six hops:

1. CSR ACX7024 (start)
2. HSRs ACX7100-32C and ACX7509
3. AG2 MX204s
4. AG3 MX480/MX10003
5. Core PTX10001-36MRs
6. SAG with MX304 (end)

Figure 15: Topology 2 for Fronthaul

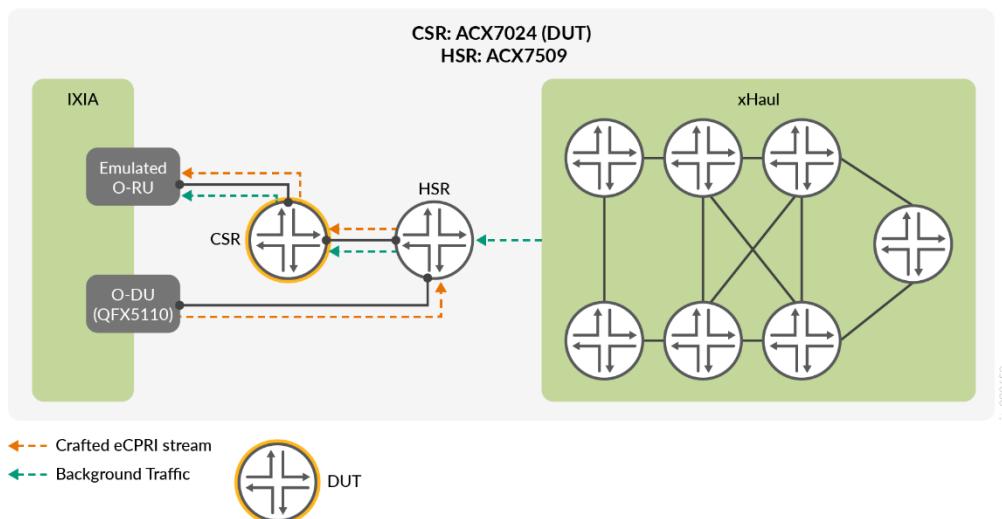


Table 13 on page 38 compares the Fronthaul and Midhaul to Backhaul performance across different service types terminating on the ACX7024 CSR. Total latency factors number of hops, for example EVPN-VPWS with three hops measured 14.63μs, amounting to 4.9μs per hop in the Fronthaul segment.

For the full test report, with complete results detailing the minimum/average/maximum latency across all featured feature types (EVPN, L2VPN, L2Circuit, and L3VPN), contact your Juniper Networks representative.

Table 13: Latency Measurements (No Congestion)

Service Type	Queue Priority	Min (μs)	Ave (μs)	Frame Size	Traffic Pattern	Port	Segment	Hop #
EVPN-VPWS	SH	10.53μs	14.63μs	64b	Continuous	10G	FH	3
EVPN-VPWS	SH	11.46μs	16.26μs	512b	Continuous	10G	FH	3
EVPN-VPWS	SH	11.82μs	18.47μs	1500b	Continuous	10G	FH	3
EVPN-VPWS	SH	10.51μs	13.56μs	64b	Burst	10G	FH	3
EVPN-VPWS	SH	11.45μs	15.19μs	512b	Burst	10G	FH	3
EVPN-VPWS	SH	11.63μs	17.38μs	1500b	Burst	10G	FH	3
EVPN-VPWS	SH	--	11.5μs	512b	Continuous	100G	FH	3
EVPN-VPWS	SH	--	15μs	512b	Continuous	1G shaper	FH	3
L2Circuit	LOW	51.19μs	72.8μs	64b	Continuous	10G	MBH	6
L2Circuit	LOW	50.23μs	77.05μs	512b	Continuous	10G	MBH	6

Table 13: Latency Measurements (No Congestion) (Continued)

Service Type	Queue Priority	Min (μs)	Ave (μs)	Frame Size	Traffic Pattern	Port	Segment	Hop #
		Latency	Latency					
L2Circuit	LOW	49.21μs	82.35μs	1500b	Continuous	10G	MBH	6
L2Circuit	LOW	50.96μs	70.42μs	64b	Burst	10G	MBH	6
L2Circuit	LOW	50.09μs	74.94μs	512b	Burst	10G	MBH	6
L2Circuit	LOW	49.07μs	80.81μs	1500b	Burst	10G	MBH	6
L2Circuit	LOW	--	65.8μs	512b	Continuous	100G	MBH	6
L2Circuit	LOW	--	122.9μs	512b	Continuous	1G shaper	MBH	6
L3VPN	LOW	38.31μs	100.72μs	64b	Continuous	10G	MBH	6
L3VPN	LOW	40.53μs	106.30μs	512b	Continuous	10G	MBH	6
L3VPN	LOW	40.42μs	132.37μs	1500b	Continuous	10G	MBH	6
L3VPN	LOW	--	99.9μs	512b	Continuous	100G	MBH	6
L3VPN	LOW	--	146.5μs	512b	Continuous	1G shaper	MBH	6

The priority queue delivers strict latency performance compared to low priority queues.

Recommendations

The ACX7024 is a reliable choice as a CSR, offering an enhanced feature-set and improved performance compared to previous ACX platforms in most situations. It is specifically designed for the CSR role, catering to the scale, bandwidth, and performance requirements associated with this function.

As compared to the ACX7100-48L, the ACX7024 demonstrates identical feature-set and comparable performance for the CSR role and scale. Compared to previous generation ACX5448/ACX710, we observed overall convergence improvements across all services in the featured design.

Segment Routing is a suggested underlay architecture for enabling seamless MPLS stitching with BGP-LU across multiple IGP and inter-AS domains. This setup can be further improved by incorporating Seamless-SR and BGP-CT once they are fully supported. By utilizing TI-LFA and ECMP mechanisms, we can achieve quick failover and resilience. However, currently ACX7000 series does not support simultaneous ECMP and Fast Reroute (FRR) functionalities.

The ACX7024 supports deterministic and effective QoS, performing within expectations.

Layer 3 (DSCP), MPLS (EXP), and Layer 2 (802.1p) traffic, whether single-tagged or dual-tagged, were accurately classified based on their respective codepoints. The priority hierarchies were properly maintained within the guaranteed and excess regions, as defined by the transmit rate. It is important to note that only the low priority queue supports weighted fair queuing (WFQ), which is necessary for weighted distribution in the excess region. The strict-high priority queue, as is typical, does not have an excess region and needs to be shaped to avoid starving the low priority queues. The CoS model implemented in this scenario ensured that the low priority queues always received the committed information rates (CIR) as configured, using transmit-rate percentages.

When a port is shaped, the CoS scheduling parameters are appropriately adjusted to match the new port speed. Codepoint preservation was successfully maintained during all tested VLAN manipulation sequences, as expected. In terms of latency performance, the ACX7024 demonstrated comparable results to the ACX7100-48L. The average single-DUT latency was slightly better, with around 4-5 μ s for the ACX7024 compared to 5-8 μ s for the ACX7100. However, the minimum achieved latency of ACX7100 was 3.5 μ s and minimum latency observed for ACX7024 was 4.6 μ s.

As compared to the previously validated ACX5448/ACX710 reference network design, we observed differences in some CoS behaviors, which are worth understanding when planning migration to ACX7000 series.

- The ACX7000 series 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. ACX5448/ACX710 perform a default 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.
- The ACX7000 series doesn't support routing-instance classification or rewrite functionalities.
- The ACX7000 series doesn't support the ability to lock the buffer using temporal or exact configuration.

For further details, contact your Juniper Networks representative for the full test report.

Although this JVD primarily focuses on the convergence of 5G xHaul infrastructure, the technologies and practical solutions discussed can serve as building blocks for developing various network architectures. These concepts can be leveraged to support multidimensional network designs and enable further advancements in network infrastructure.