Engineering Simplicity

5G Fronthaul Network Using Seamless MPLS Segment Routing—Juniper Validated Design (JVD)

Published
2025-06-05

Table of Contents

5G Fronthaul Network Using 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 provides a Juniper Validated Design (JVD) for a 5G Fronthaul network using Seamless Segment Routing Multiprotocol Label Switching (SR-MPLS) with Juniper Network's next-generation ACX7000 platform. Through careful testing, we found that the ACX7100-32C/48L and ACX7509 are excellent choices for access and aggregation purposes. They offer enhanced performance and a wide range of advanced features, outperforming previous ACX platforms in most scenarios. Both the ACX7100 and ACX7509 are particularly well-suited for the Hub Site Router (HSR) or Lean Edge segments, as they provide the necessary scale, bandwidth, feature velocity, and performance capabilities. While the ACX7100 surpasses the requirements for access nodes, it's also an ideal option for 400G Fronthaul or Metro Access deployments.

Solution Benefits

A 5G Fronthaul network using seamless MPLS segment routing offers several advantages:

- **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 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 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**: SR-MPLS 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 quality of service (QoS) management, which are essential for delivering high-performance 5G services.
- **Network Resiliency**: 5G networks require high availability and resiliency to ensure uninterrupted service. SR-MPLS offers fast rerouting mechanisms and supports protection and restoration schemes, allowing the network to quickly recover from failures and maintain service continuity.
- **Simplified Operations**: SR-MPLS 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: Fronthaul, Midhaul, and Backhaul as shown in Figure 1 on page 2.

The primary focus of the JVD is on providing design options and performance validation of the Fronthaul network considered in the context of the end-to-end converged 5G xHaul network.

Figure 1: 5G xHaul Reference Network

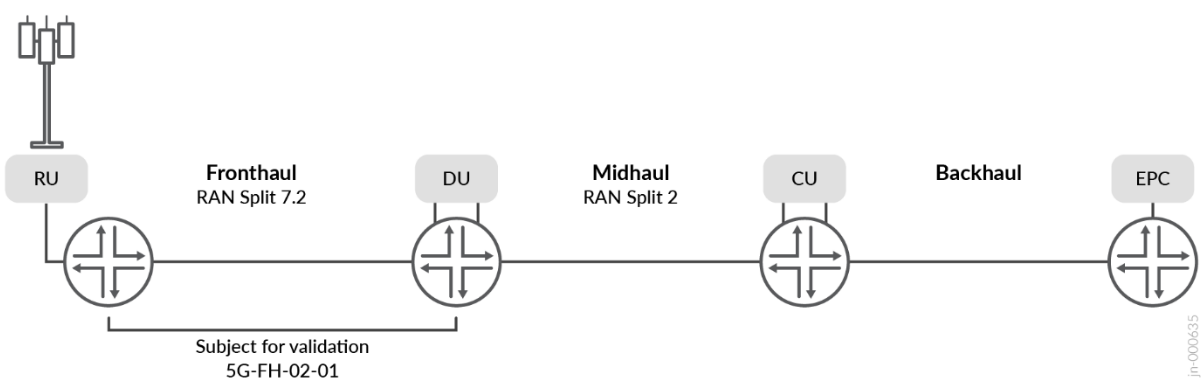

Fronthaul segment enables Layer 2 connectivity between Open Radio Unit (O-RU) (cell site) and Open Distributed Unit (O-DU) (shown as RU and DU in Figure 1 on page 2) in the radio access network (RAN), allowing them to communicate for control, data, and management traffic. It also ensures time and frequency 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. The current solution for 5G Fronthaul transport is based on O-RAN Alliance architecture [ORAN-WG9.XPSAAS.0-v00.01].

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 like 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 3 (as proposed by the O-RAN Alliance), summarizes the deployment scenarios for the RAN according to the ITU-T for simultaneous support of 4G and 5G.

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

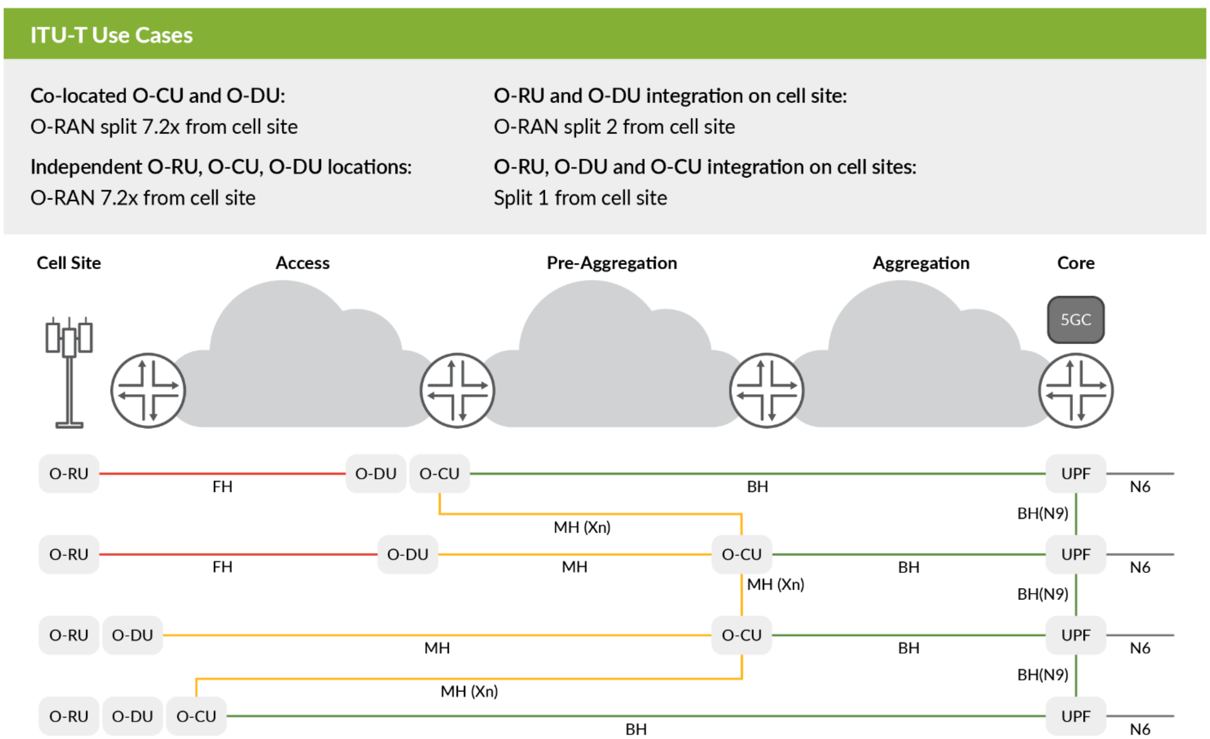

Figure 3 on page 4 shows an end-to-end 5G xHaul network, modeled after common topology [O-RAN.WG9.XPSAAS-v02.00], which defines four segments of transport infrastructure: access, pre-aggregation, aggregation, and transport core.

Figure 3: 5G xHaul Network Topology

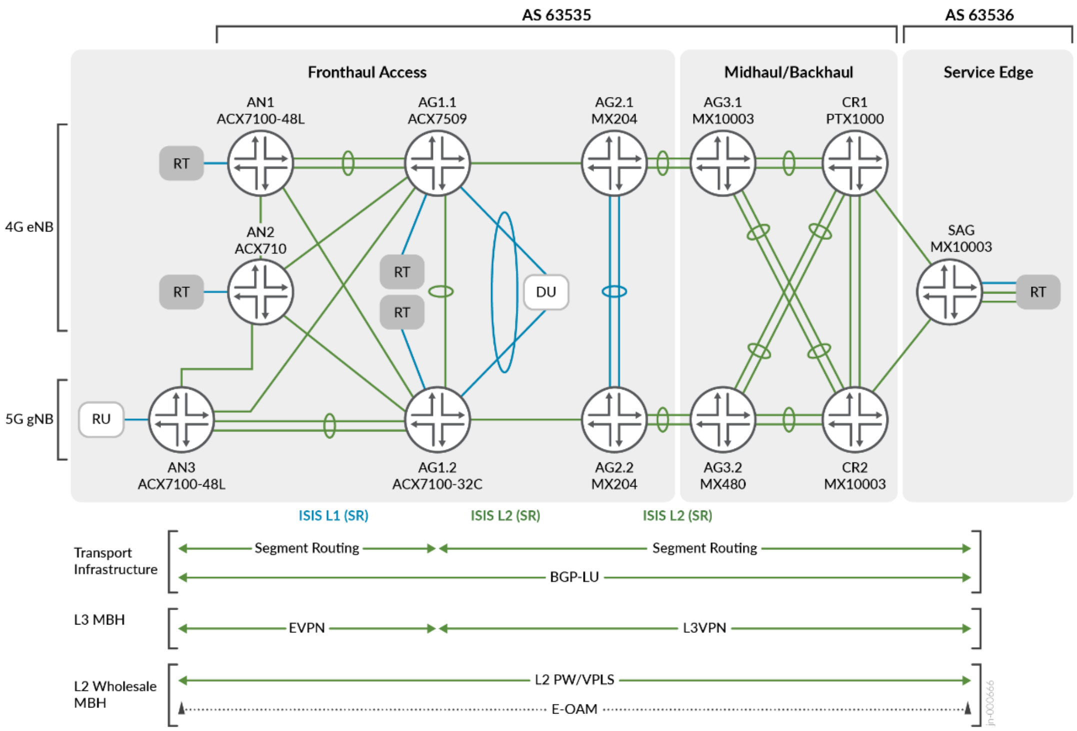

Topology Definitions:

1. **Access**: ACX7100-48L as DUT (AN3), ACX7100-48L (AN1), ACX710 (AN2) cell site routers.

2. **Pre-Aggregation**: ACX7509 (AG1.1) and ACX7100-32C (AG1.2) DUT hub site routers.

3. **Aggregation**: MX204s (AG2.1/AG2/2), MX10003 (AG3.1), MX480 (AG3.2) aggregation routers.

4. **Core Network**: PTX1000 (CR1) and MX10003 (CR2) core routers. MX10003 (SAG) services router.

The network topology in this use case combines Seamless MPLS and ISIS Segment Routing technologies. These technologies enable the support of multiple 5G and 4G services, such as:

- Layer 3 VPN (IPv4 and IPv6)
- Ethernet VPN–virtual private wireless service (EVPN-VPWS)
- EVPN Flexible Cross Connect (EVPN-FXC)
- L2VPN
- L2Circuit
- Border Gateway Protocol-Virtual Private LAN Service (BGP-VPLS)

Solution Architecture and Design

IN THIS SECTION

The reference solution architecture for end-to-end xHaul, shown in Figure 4 on page 6, deploys a spine-leaf access topology (Fronthaul) and core, peering and services gateway roles comprising the xHaul infrastructure.

The network setup includes Segment Routing MPLS 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 MPLS connectivity, BGP-Labeled Unicast (BGP-LU) is enabled at the border nodes.

Figure 4: JVD 5G Fronthaul Services Topology

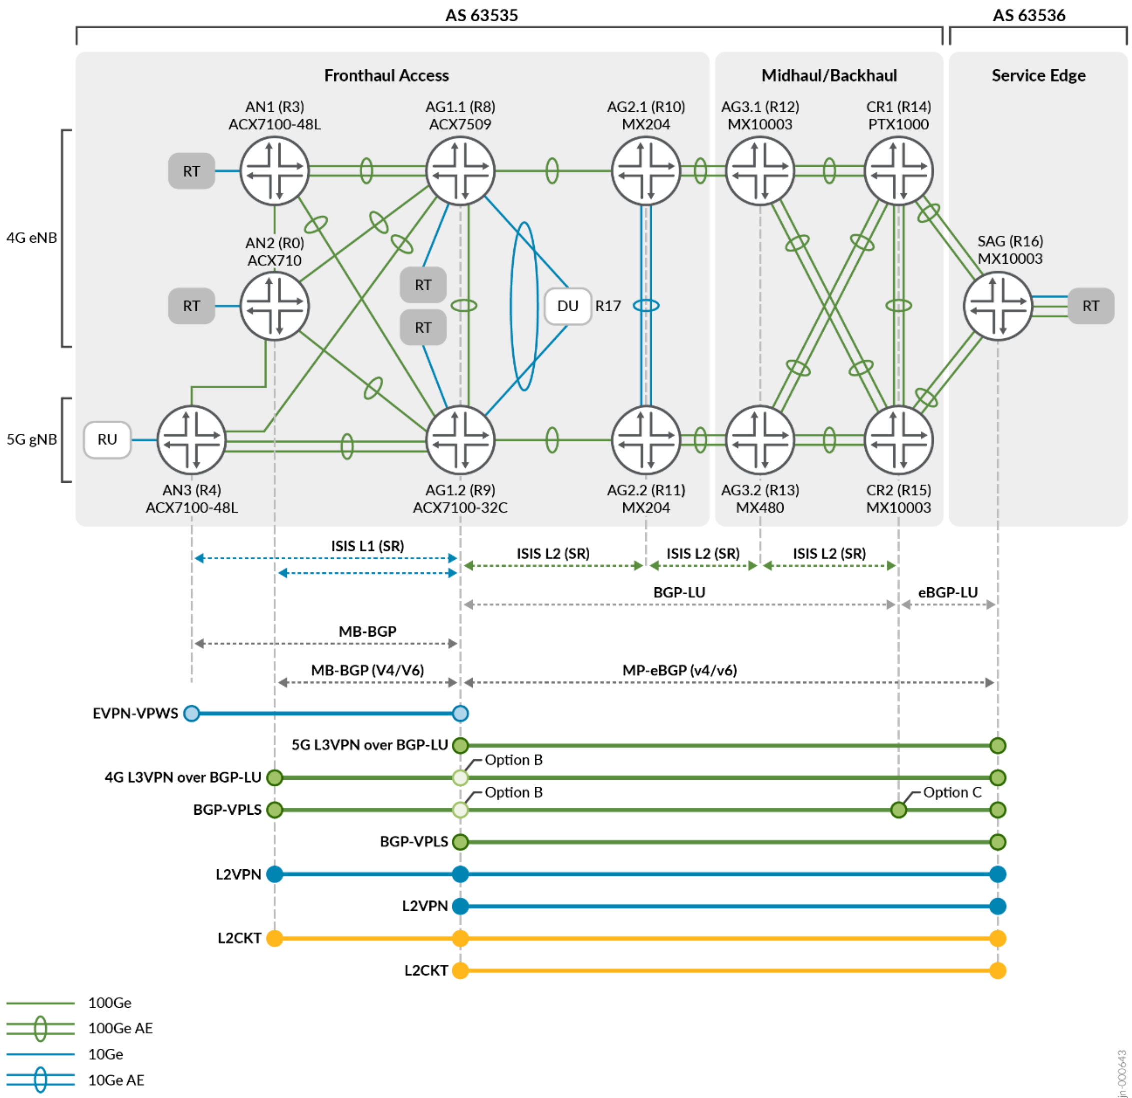

To manage current network scale and allow future growth, two sets of route reflectors are used at CR1 and CR2, primarily serving the westward HSR (AG1) clients. AG1.1/AG1.2 act as redundant route reflectors specifically for the access Fronthaul segment. Inter-AS Option-B solutions are supported through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

To ensure the performance and reliability of the network, the design proposes additional enhancements, including E-OAM performance monitoring and Flow-Aware Transport Pseudowire (FAT-PW) load-balancing. These enhancements help monitor the network performance and balance the traffic load efficiently.

Overlay Services

The overlay services in the network use different combinations of VLAN operations. These operations are applied to each Layer 2 service types, including EVPN, L2Circuit, VPLS, and L2VPN. Additionally, L3VPN services incorporate IPv6 tunnelling to validate 6vPE functionality.

As shown in Figure 4 on page 6, inter-domain VPN Option-B is used at the HSR/AG1 region, and Inter-AS Option-C is used between the SAG and CR border nodes. These nodes are enabled by BGP-LU.

This combination of VPN's is designed in a way to allow following traffic flows in the 5G xHaul network:

- 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

For more detailed information about this architecture, contact your Juniper Networks representative.

Fronthaul Network Design and 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.

As shown in Figure 5 on page 8, the Fronthaul network consists of ACX7000 series routers interconnected by high-capacity links. The ACX7100-48L serves as the CSR, providing connectivity between O-RU and HSR aspects of the RAN. 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.

Figure 5: ACX Platforms Positioning Within Fronthaul Network Topology

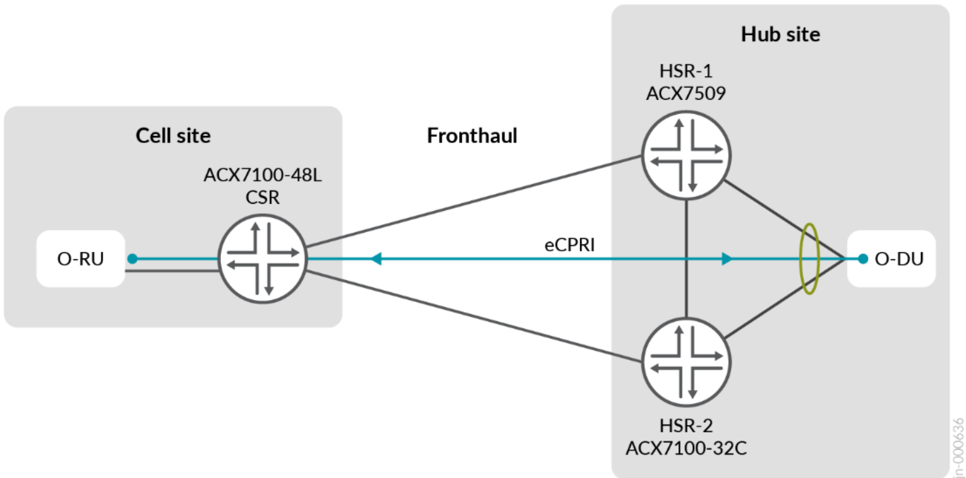

Fronthaul Layer 2 Connectivity Models

There are several possible connectivity models between O-RU and O-DU, but for the purposes of this JVD, we've validated these two:

- EVPN-VPWS Single-Homed supporting dedicated MAC for eCPRI without redundancy
- EVPN-VPWS with A/A LAG DU attachment

Figure 6 on page 9 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. Currently, Ethernet OAM with performance monitoring is supported for the single-homed configuration in this model.

Figure 6: EVPN-VPWS Single-Homed Supporting Dedicated MAC for eCPRI Without Redundancy

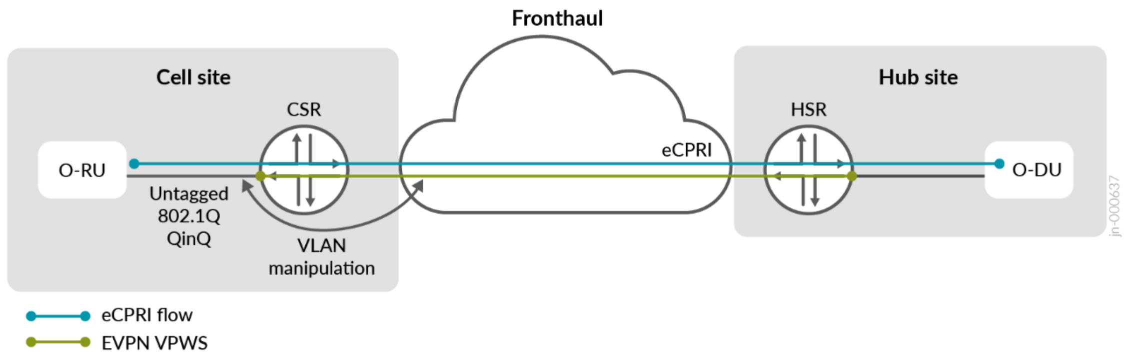

Figure 7 on page 9 illustrates the second connectivity model. In this scenario, the network uses two configurations between the CSR (AN3) and HSR-1/HSR 2:

- EVPN-VPWS with active-active multihoming
- EVPN-VPWS with Flexible Cross Connect (FXC) active-active multihoming

To enable traffic load sharing, an Active/Active Ethernet Segment Identifier (ESI) Link Aggregation Group (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 Aggregate Ethernet (AE) with both links actively carrying traffic.

In this scenario, eCPRI packets might arrive on either of the O-DU links from the HSRs. Similarly, eCPRI packets can be sent across either of the HSR-1 or HSR-2 uplinks to support active-active functionality. This setup ensures flexibility and redundancy in the network for improved performance.

Figure 7: O-RAN Fronthaul Network A/A EVPN-VPWS

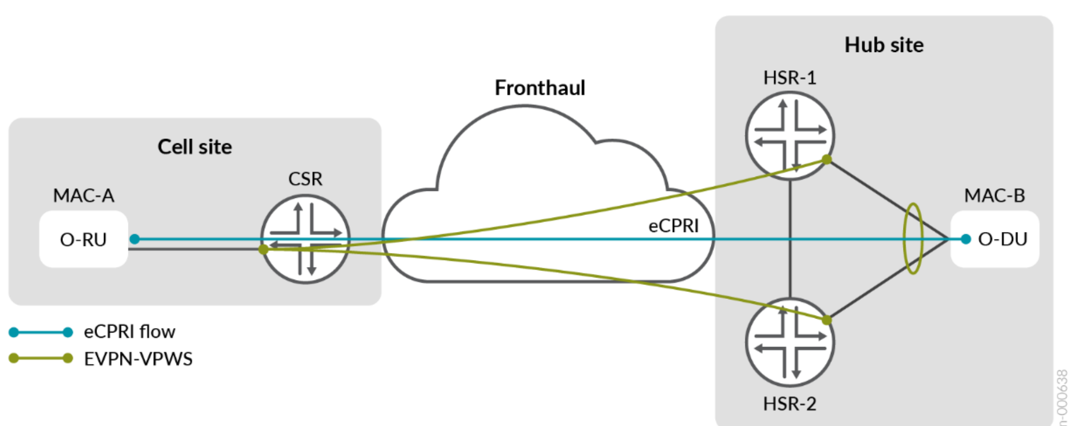

Layer 3 Connectivity Models

We chose L3VPN protocol to facilitate Layer 3 connectivity between O-DU and vCU/vEPC elements of the 5G xHaul. The JVD proposes two different connectivity models between the Hub Site Router (HSR) and the O-DU, with both supporting L3 multihoming between O-DU and pair of HSRs:

- EVPN IRB anycast gateway with L3VPN, refer to Figure 8 on page 10
- BD with IRB and static MAC/ARP with L3VPN, refer to Figure 9 on page 11
- We tested both connectivity models independently and collected data on their coexistence and convergence. We validated that both models are effective in facilitating the necessary connectivity between the HSR and O-DU.

For further details about the testing for these connectivity models, contact your Juniper Networks representative.

Figure 8: EVPN IRB Anycast Gateway with L3VPN

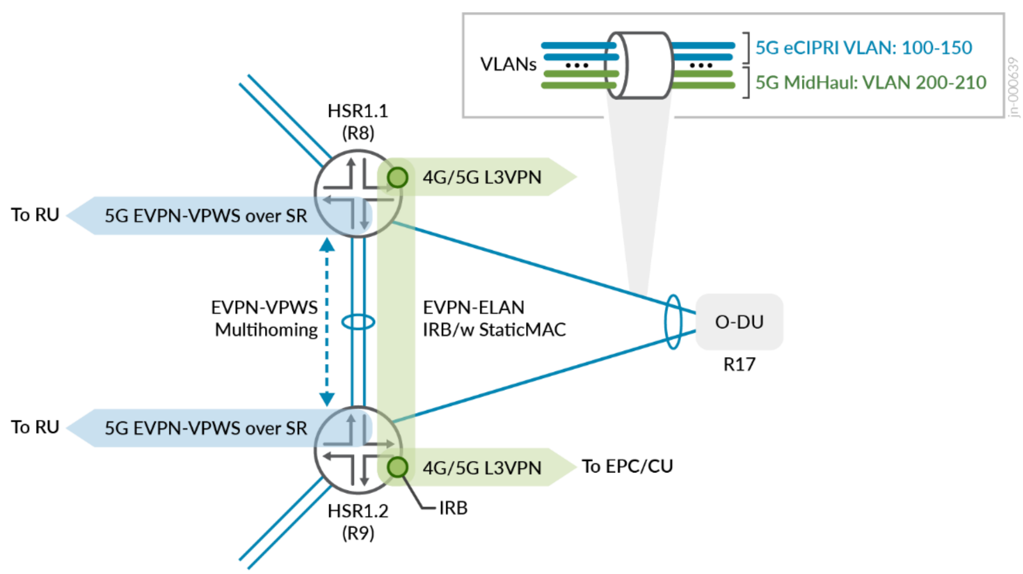

Figure 9: BD with IRB and Static MAC/ARP with L3VPN

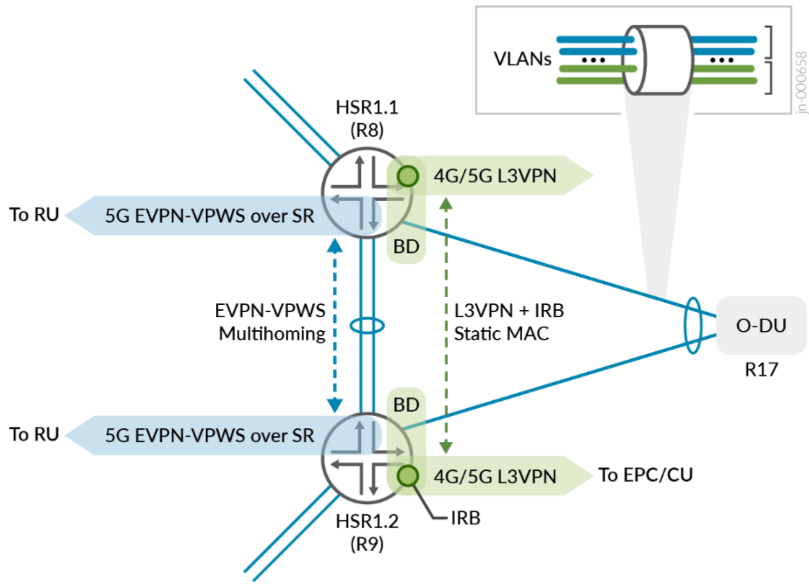

VLAN Operations

The ACX7000 series offers a wide range of VLAN operations. The proposed VLAN operations specifically cater to EVPN eCPRI-emulated traffic flows and support various normalization scenarios:

- Untagged frames
- Tagged (802.1q)
- Dual-tagged (802.1ad)

While the main emphasis of this JVD is on the Fronthaul segment, we also validated most of the VLAN operations for each VPN service type. To see the full test report, contact your Juniper Networks representative.

Figure 10: Ethernet Encapsulation and VLAN Normalization of eCPRI Flows

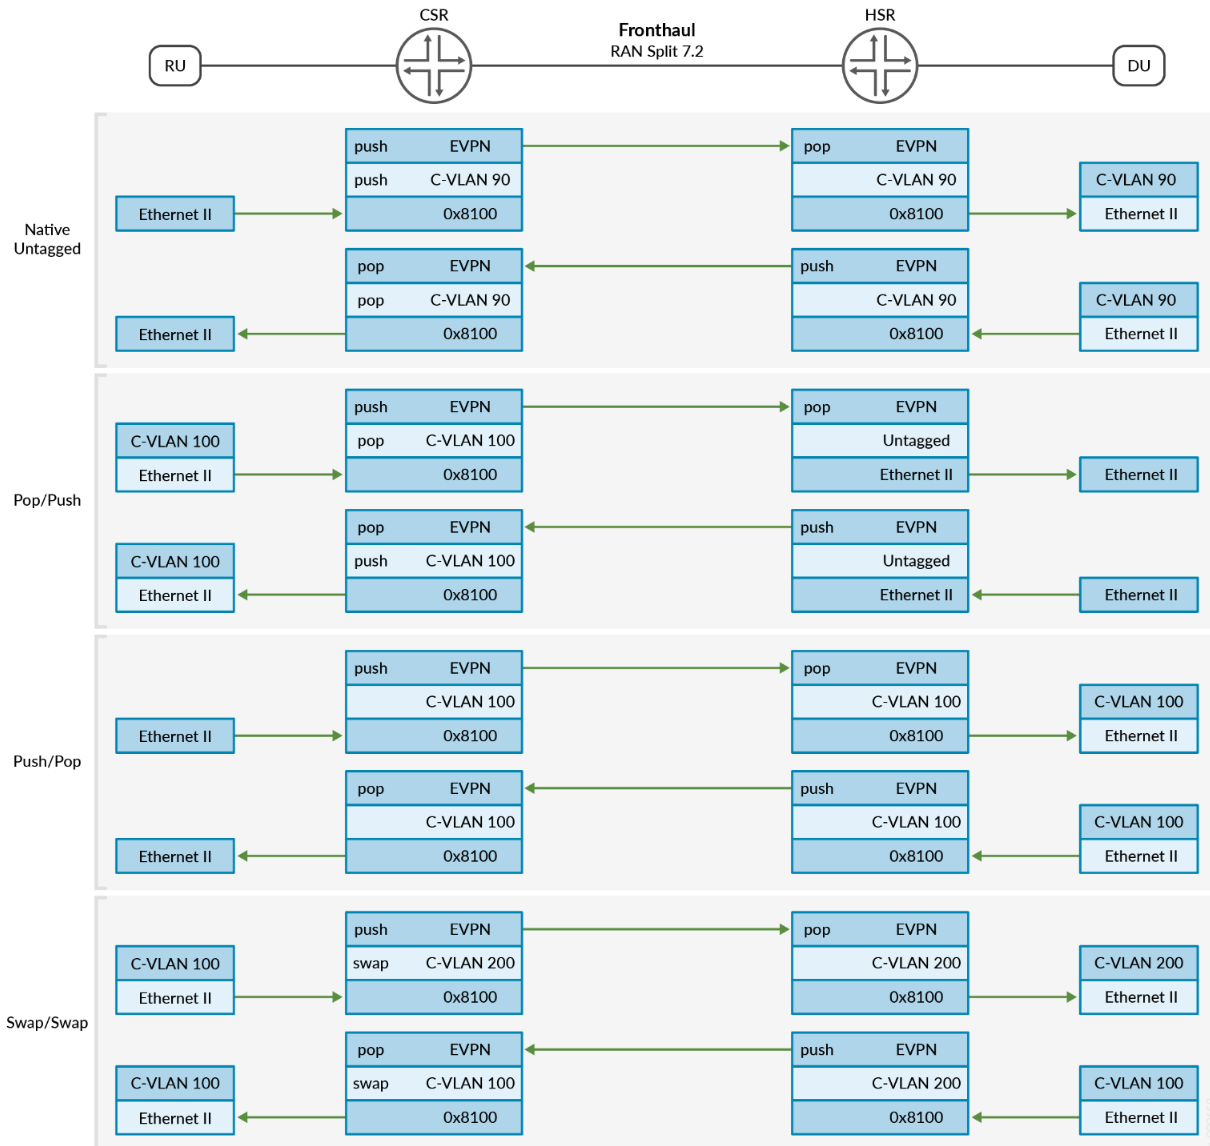

Figure 11: Ethernet Encapsulation and VLAN Normalization of eCPRI Flows (cont'd)

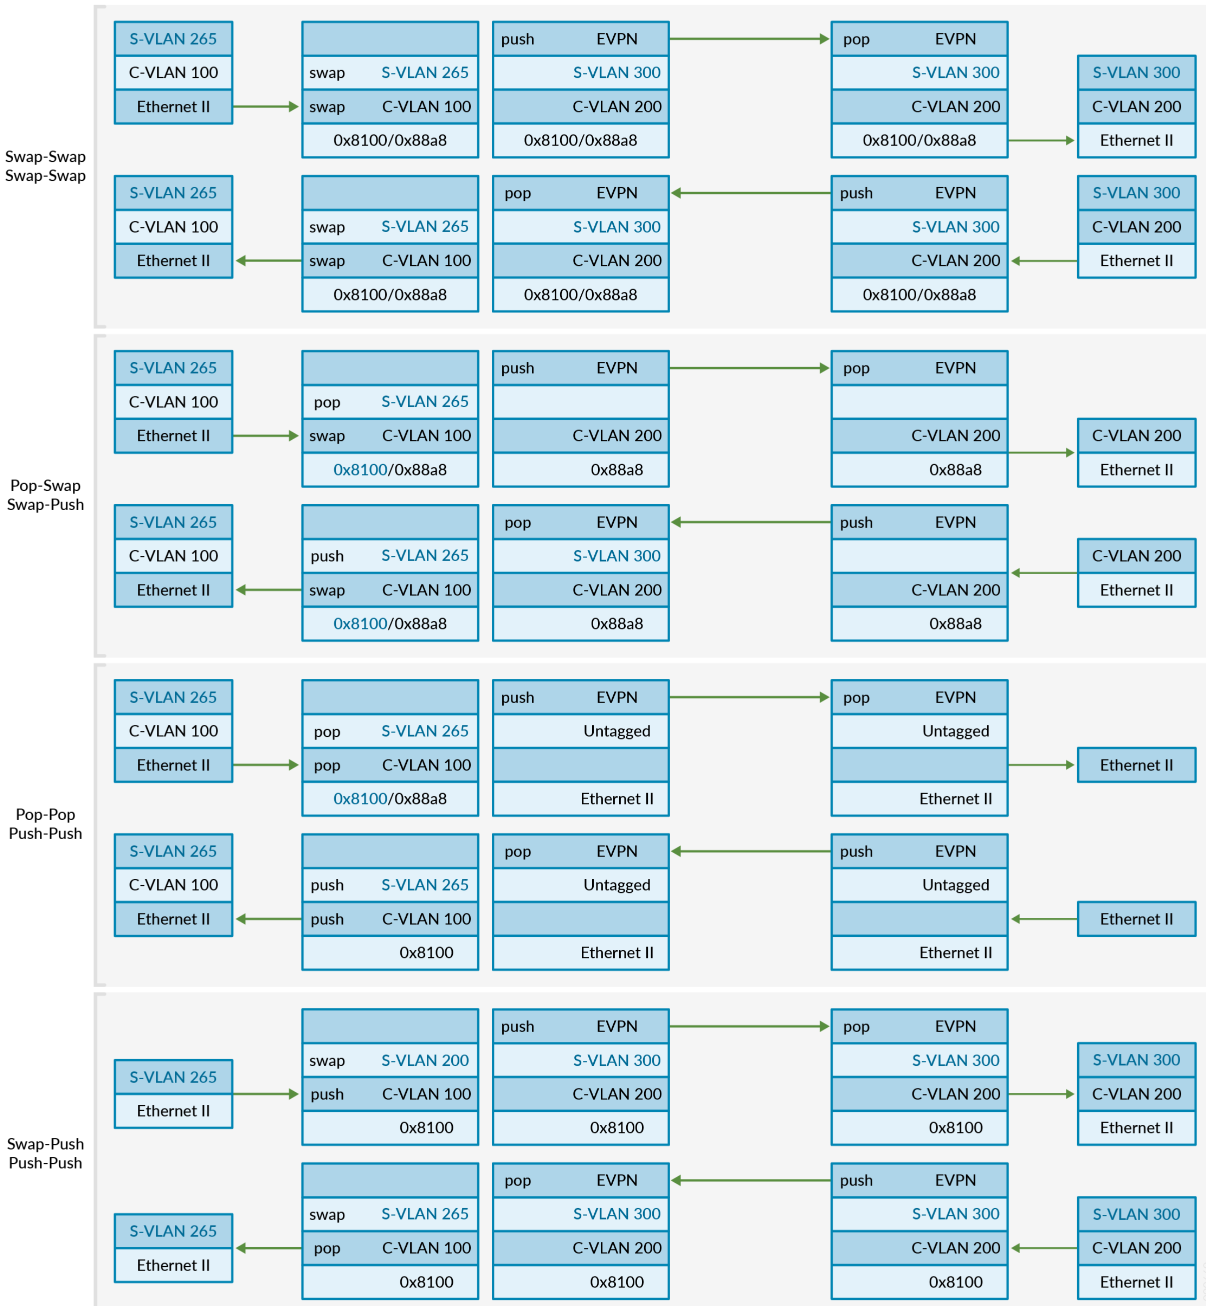

Solution and Validation Key Parameters

IN THIS SECTION

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 Design and Infrastructure

- 5G xHaul MBH reference architecture
- Seamless MPLS across xHaul IGP domains (Inter-AS BGP-LU)
- Fast failover & detection TI-LFA, BFD, OAM, ECMP, etc.
- Segment Routing L-ISIS
- Redundant Route Reflectors
- EVPN-VPWS and FXC A/A Multihoming
- EVPN-ELAN with Multihoming and EVPN IRB
- BGP-VPLS, L2VPN, L2Circuit, L3VPN

- Inter-AS Option B/C
- O-RAN Alliance [ORAN.WG9.XPSAAS-v02.00]

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, kindly 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, please reach out to your Juniper Networks representative

Table 1: Scale Summary

Feature	AN3 (ACX7100-48L) Access / CSR	AG1.1 (ACX7509) Pre-Agg / HSR	AG1.2 (ACX7100-32C) Pre-Agg / HSR	SAG (MX10003) Services Agg
RIB/FIB	351k/123k	303k/121k	295k/121k	381k/116k
VLANS	2115	7300	7300	8414
EVPN-VPWS SH	1000	1000	1000	0
EVPN-VPWS MH	101	101	101	0
EVPN-ELAN (L3VPN)	0	500	500	0
Bridge Domain (L3VPN)	0	500	500	0
L2Circuit	200	1000	1000	2400
L2VPN	200	1000	1000	2410

Table 1: Scale Summary *(Continued)*

Feature	AN3 (ACX7100-48L) Access / CSR	AG1.1 (ACX7509) Pre-Agg / HSR	AG1.2 (ACX7100-32C) Pre-Agg / HSR	SAG (MX10003) Services Agg
L3VPN (OSPF)	100	100	100	1101
L3VPN (BGP)	100	100	100	100
VPLS	200	1000	1000	2400
MAC (VPLS)	10.6k	22k	102k	145k
CFM UP MEP	1200	1000	1000	1000

Anticipated Fronthaul Network Resiliency and Latency

The solution creates a resilient network system that safeguards against link or node problems in the Fronthaul segment. It ensures that if a link fails between CSR and HSR nodes, the traffic flow will be restored within 50 milliseconds or less. Additionally, the expected delay for packets between RU and DU should not exceed 120 microseconds.

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)

Contact your Juniper Networks representative for the full test report and feature list.

Test Bed

Figure 12 on page 18 illustrates the test bed 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, specifically the Remote Tester (RT), for scalability purposes.
- **Midhaul and Backhaul segments:** These are represented by ring topologies and serve aggregation and core roles, but they are not the main focus of this discussion.

Figure 12: JVD 5G Fronthaul Lab Topology

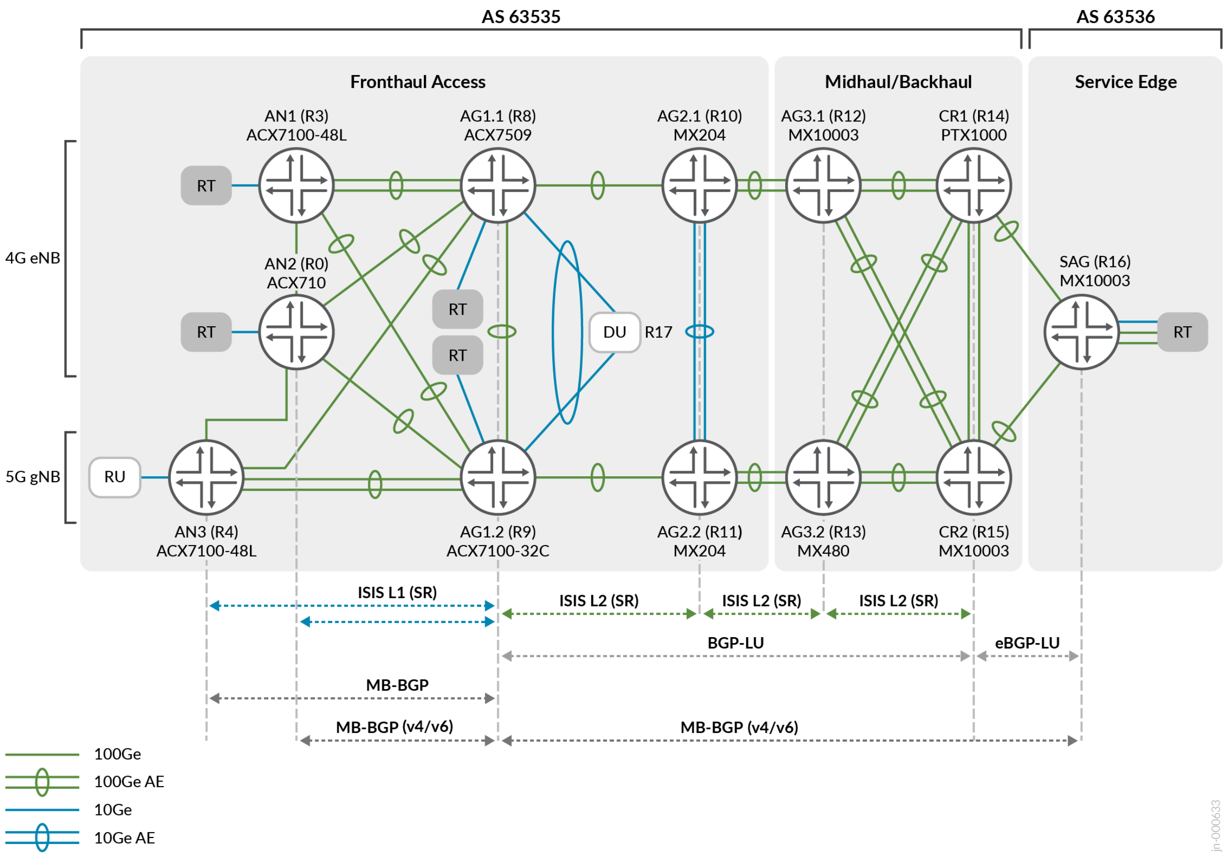

Table 2: Topology Definitions

Layer	Devices Under Test
Access	ACX7100-48L as DUT (AN3), ACX7100-48L (AN1), ACX710 (AN2) cell site routers
Pre-Aggregation	ACX7509 (AG1.1) and ACX7100-32C (AG1.2) DUT 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

In the testing process, we used three main devices as Devices Under Test (DUT) to run all the test cases. The ACX7100-48L router was used as the Access Node Cell Site Router (CSR) for both 5G and 4G. The ACX7100-32C and ACX7509 routers were used as the Pre-Aggregation Hub Site Routers (HSR) for the AG1 role.

To support and expand the testing environment, we introduced additional helper nodes. We added another ACX7100-48L router (AN1) to scale the setup and emulate identical traffic flows between each HSR DUT, ensuring accurate and uniform validation against both ACX7509 and ACX7100-32C HSRs. MX routers were used in different roles, such as Aggregation (AG2/AG3), Core (CR2), and services router (SAG). The PTX platform was used for the Core Router (CR1).

For a complete test report with detailed information about the hardware and software used, contact your Juniper Networks representative.

Solution Validation Goals

Here are the major test goals:

- 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 ACX7100-48L (AN3), ACX7100-32C (AG1.2) and ACX7509 (AG1.1), with adjacent link/node failures for all traffic types.
- Measure solution resilience of layer 2/3 flows from specified DUTs, including:
    - Access Node (AN) to Pre-Aggregation AG1 (O-RU to O-DU)
    - Pre-Aggregation AG1 (including O-DU) to Services Gateway (SAG)
    - Access Node (AN) to SAG (EPC)
- Validate network stability for major 4G/5G traffic flows at scale with each VPN service type over SR-MPLS during normal and stress conditions.
- Validate consistency and resiliency of the ACX7100-48L, ACX7100-32C and ACX7509 systems against negative stress conditions (enabled/disable control and data plane daemons, add/delete configurations, etc.)
- Identify product limitations, anomalies, and open Problem Reports (PRs) exposed during validation stage.
- OAM¾ Connectivity Fault Management (CFM), Y.1731¾over VPLS and EVPN-VPWS services.
- Flow-Aware Transport Pseudowire (FAT-PW) load-balancing.

Solution Validation Non-Goals

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

- Class of Service
- Underlay MPLS/SR transport other than specified in the Test Goals
- Network Slicing
- Flex Algo
- Classful Transport, BGP-CT, SSR
- BGP-PIC
- Management and Automation
- Chassis-based High Availability (as applicable)
- End-to-End Timing and Synchronization Distribution: Synchronous Ethernet, IEEE1588v2
- SLA Monitoring: RFC 2544, Y.1564, TWAMP, Active Assurance
- Telemetry

Results Summary and Analysis

IN THIS SECTION

Within the course of validation, we were able to demonstrate a robust solution for 5G Fronthaul and Midhaul transport infrastructure utilizing Seamless MPLS. The JVD generates reasonable multi-vector scale of L2/L3 connectivity services as compared with MNO and MAN operator expectations for real network deployments, while satisfying stringent SLA requirements. The scale reference in Table 1 on page 15 characterizes primary multidimensional KPI's represented in the validated profile.

ACX7100-48L, ACX7100-32C and ACX7509 as the DUT routers have each successfully passed 186 test cases curated to support the given reference architecture.

With this network design, the architecture can deliver fast restoration within 50ms for most traffic flows transported over ISIS-SR with TI-LFA. Load distribution and optimization knobs were shown to improve service restoration against link/node failures. Link events consistently achieve <50ms convergence, see Table 4 on page 24. Node failures were more disruptive as expected, exacerbated by scale and in some rare cases, production limitations. Contact your Juniper Networks representative for a complete results report and analysis.

ECMP Load Balancing

The design minimizes traffic issues when links or nodes fail by distributing the load across equal cost paths and making use of TI-LFA protection mechanisms. Multiple load-sharing methodologies are incorporated, such as adjusting IGP metrics, BGP multipath, ECMP fast-reroute, and VPN-unequal-cost for L3VPN services. We also activated FAT-PW label on ACX7000 devices where supported, including Layer 2 VPN and Layer 2 Circuits, which enables additional granularity to the hash computation. Hash-key configuration is mandatory on ACX7000 series to perform traffic load balancing. By default, there are no hash-keys enabled. This validation enables all relevant hash-keys for extracting L2, L3 and MPLS fields.

The full test report includes details about the load sharing configurations enabled for this validation. Junos OS Evolved 22.3R1-S1/R2 and later are recommended versions for ACX7000 series and includes additional load sharing enhancements. Contact your Juniper Networks representative if you'd like to receive the complete report.

Table 3: ACX7000 ECMP Summary Results

ACX7000 ECMP Load Balancing Performance [1]							
Service: DUT (Traffic Path)	ECMP Links	Flow	FAT-PW	Link1 LAG (#)	Link2 LAG (#)	Link3 LAG (#)	Link4 LAG (#)
L2VPN: ACX7100-48L (AN3 to SAG)	3	100kfps	Y	ae20(2) 36.6kpps	ae23(1) 33.1kpps	et-0/0/50 33.1kpps	N/A

Table 3: ACX7000 ECMP Summary Results *(Continued)*

ACX7000 ECMP Load Balancing Performance [1]							
Service: DUT (Traffic Path)	ECMP Links	Flow	FAT-PW	Link1 LAG (#)	Link2 LAG (#)	Link3 LAG (#)	Link4 LAG (#)
L2CKT: ACX7100-48L (AN3 to SAG)	2	100kfps	Y	49.2kpps	et-0/0/50 53.2kpps	N/A	N/A
VPLS: ACX7100-48L (AN3 to SAG) *[2]*	3	100kfps	N	ae20(2) 36.0kpps	ae23(1) 33.2kpps	et-0/0/50 33.7kpps	N/A
EVPN: ACX7100-48L (AN3 to DU)	3	100kfps	N	ae20(2) 36.0kpps	ae23(1) 33.4kpps	et-0/0/50 33.0kpps	N/A
L3VPN: ACX7100-48L (AN3 to AG1.2)	3	100kfps	N	ae20(2) 36.9kpps	ae23(1) 33.7kpps	et-0/0/50 33.4kpps	N/A
L2VPN: ACX7509 (AG1.1 to SAG)	2	100kfps	Y	ae4(2) 50kpps	ae5(1) 49.6kpps	N/A	N/A
L2CKT: ACX7509 (AG1.1 to SAG)	2	100kfps	Y	ae4(2) 53.2kpps	ae5(2) 50.2kpps	N/A	N/A

Table 3: ACX7000 ECMP Summary Results *(Continued)*

ACX7000 ECMP Load Balancing Performance [1]							
Service: DUT (Traffic Path)	ECMP Links	Flow	FAT-PW	Link1 LAG (#)	Link2 LAG (#)	Link3 LAG (#)	Link4 LAG (#)
VPLS: ACX7509 (AG1.1 to SAG) *[2]*	2	100kfps	N	ae4(2) 53.3kpps	ae5(2) 50.1kpps	N/A	N/A
EVPN: ACX7509 (AG1.1 to AN1)	4	100kfps	N	ae4(2) 25.0kpps	ae19(2) 25.0kpps	ae22(1) 25.0kpps	ae23(1) 25.0kpps
L3VPN: ACX7509 (Unit Test)incrementing src/dst ip x3	4	50kfps	N	et-1/0/8 13kpps	et-1/0/10 12kpps	et-6/0/8 12kpps	et-6/0/13 13kpps

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

[2] Only Known Unicast is shown. VPLS BUM traffic should not load balance over ECMP. This is an expected behavior.

Convergence Briefing

Network convergences were validated against the following failure events:

- DUT Adjacent link failures
- DUT Indirect link failures
- DUT Node failures
- Daemon failure and configuration add/delete events

Overall convergence results are within expectations for the given network design. In the Fronthaul (CSR to HSR), failure and restoration events were well within 50ms for all VPN traffic types. The latest generation ACX7000 product family demonstrated optimal convergence, improving upon previous generation Broadcom-based platforms. For all segment routing links, TI-LFA with loose mode allows transitioning to link-protection should node-protection become unavailable.

L2-L3 interworking scenarios with longer convergence can be attributed to control plane triggered global repair events, which may be improved with additional mechanisms such as BGP Prefix Independent Convergence (BGP-PIC). These models include L3VPN to EVPN-ELAN IRB Anycast Gateway, and L3VPN to Bridge Domain with IRB Static Virtual-MAC. EVPN convergence may be further improved with dynamic-list-next-hop (DLNH) and evpn-egress-link-protection (EPL).

Table 4 on page 24 and Table 5 on page 25 show the convergence results in the given scenarios:

- TI-LFA without ECMP (default for the ACX7000 Fronthaul segment).
- TI-LFA with ECMP (default for MX/PTX paths in core).

ACX7000 does not simultaneously support FRR and ECMP. During certain path failure events, TI-LFA is prioritized over ECMP to avoid triggering global repair.

Table 4: Convergence Times for 5G Fronthaul Failure Events Per Flow Type

Flow Type	5G EVPN-VPWS (msec)		5G EVPN-FXC (msec)		5G L3VPN-BD IRB (msec)		5G L3VPN-EVPN IRB (msec)		L3VPN (msec)
Events	CSR->DU	DU->CSR	CSR->DU	DU->CSR	DU->SAG	SAG->DU	DU->SAG	SAG->DU	CSR<>HSR
AN3 to AG1.2 disable	7	8	0	6	-	-	-	-	6
AN3 to AG1.2 enable	1	0	0	0	-	-	-	-	0
AN3 to AG1.1 disable	5	1	0	0	-	-	-	-	4
AN3 to AG1.1 enable	10	0	0	0	-	-	-	-	1
AG1.2 to AG1.1 disable	5	0	0	0	0	0	5	0	1
AG1.2 to AG1.1 enable	5	0	0	1	0	0	0	0	1

Table 4: Convergence Times for 5G Fronthaul Failure Events Per Flow Type *(Continued)*

Flow Type	5G EVPN-VPWS (msec)		5G EVPN-FXC (msec)		5G L3VPN-BD IRB (msec)		5G L3VPN-EVPN IRB (msec)		L3VPN (msec)
Events	CSR->DU	DU->CSR	CSR->DU	DU->CSR	DU->SAG	SAG->DU	DU->SAG	SAG->DU	CSR<>HSR
AG1.2 to AG2.2 disable	-	-	-	-	0	0	0	0	-
AG1.2 to AG2.2 enable	-	-	-	-	1	0	0	10	-
AG2.2 to AG3.2 disable	-	-	-	-	0	0	0	0	-
AG2.2 to AG3.2 enable	-	-	-	-	0	0	1	0	-
CR2 to SAG disable	-	-	-	-	0	0	0	1	-
CR2 to SAG enable	-	-	-	-	1	0	0	0	-

Table 5: Convergence Times for MBH and Midhaul Failure Events Per Flow Type

Flow Type	5G EVPN-VPWS (msec)		5G EVPN-FXC (msec)		5G L3VPN-BD IRB (msec)		5G L3VPN-EVPN IRB (msec)	
Events	CSR<>SAG	HSR<>SAG	CSR<>SAG	HSR<>SAG	CSR<>SAG	HSR<>SAG	CSR<>SAG	HSR<>SAG
AN3 to AG1.2 disable	2	-	5	-	1	-	0	-
AN3 to AG1.2 enable	1	-	5	-	0	-	0	-

Table 5: Convergence Times for MBH and Midhaul Failure Events Per Flow Type *(Continued)*

Flow Type	5G EVPN-VPWS (msec)		5G EVPN-FXC (msec)		5G L3VPN-BD IRB (msec)		5G L3VPN-EVPN IRB (msec)	
Events	CSR<>SAG	HSR<>SAG	CSR<>SAG	HSR<>SAG	CSR<>SAG	HSR<>SAG	CSR<> SAG	HSR<>SAG
AN3 to AG1.1 disable	5	-	6	-	5	-	0	-
AN3 to AG1.1 enable	1	-	5	-	5	0	0	-
AG1.2 to AG1.1 disable	2	2	0	1	2	0	0	4
AG1.2 to AG1.1 enable	2	0	0	5	0	0	0	0
AG1.2 to AG2.2 disable	0	144	4	5	0	4	4	4
AG1.2 to AG2.2 enable	0	1	35	2	0	5	0	0
AG2.2 to AG3.2 disable	0	65	0	80	0	65	0	78
AG2.2 to AG3.2 enable	2	0	55	0	0	2	1	0
CR2 to SAG disable	0	48	0	60	0	50	0	44

Table 5: Convergence Times for MBH and Midhaul Failure Events Per Flow Type *(Continued)*

Flow Type	5G EVPN-VPWS (msec)		5G EVPN-FXC (msec)		5G L3VPN-BD IRB (msec)		5G L3VPN-EVPN IRB (msec)	
Events	CSR<>SAG	HSR<>SAG	CSR<>SAG	HSR<>SAG	CSR<>SAG	HSR<>SAG	CSR<>SAG	HSR<>SAG
CR2 to SAG enable	0	0	0	0	0	0	0	0

Recommendations

The ACX7100-32C/48L and ACX7509 are excellent choices for access and aggregation purposes. They offer enhanced performance and a wide range of advanced features, delivering significant scale and performance over previous generation ACX platforms. Both the ACX7100 and ACX7509 are particularly well-suited for the HSR or Lean Edge metro segments, as they provide the necessary scale, bandwidth, and performance characteristics. The ACX7100 feature-rich capabilities surpass most access nodes requirements but are ideal for supporting 400G Fronthaul or Metro Access deployments.

EVPN-VPWS is primarily tested at scale, but Flexible Cross Connect (FXC) may further improve HSR resiliency by aggregating services and reducing control plane scale. However, no issues were seen on ACX7000 series handling these services at scale with E-OAM for single-homed EVPN.

Segment routing is a recommended underlay architecture to support end-to-end Seamless MPLS stitching with BGP-LU across multiple IGP and inter-AS domains; this can be further enhanced with seamless segment routing and BGP-CT once supported, starting with Junos OS 23.1R1-EVO. By utilizing TI-LFA and ECMP mechanisms, we can achieve fast failover and resilience in the reference architecture. The ACX7100 and ACX7509 platforms can be confidentially deployed to support the featured protocols and services.

While the reference design of this JVD is the 5G xHaul infrastructure, the technologies and practical solutions covered can be leveraged as building blocks from which additional designs may evolve to support multidimensional network architectures.