Engineering Simplicity

5G CSR xHaul Seamless Segment Routing —Juniper Validated Design (JVD)

Published
2025-06-05

Table of Contents

5G CSR xHaul Seamless 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 focusses on 5G xHaul network reference design with seamless segment routing using ACX710 and ACX5448 Universal Metro Routers in the role of cell site Access Node (AN); ACX5448 as Pre-Aggregation (AG1); MX204 Universal Router as Aggregation (AG2); MX10003 and MX480 Universal Routers as Aggregation (AG3); PTX1000 Router and MX10003 Universal Router as Core Routers (CR) and MX10003 Universal Router as a Services Aggregation Gateway (SAG) router.

The seamless segment routing solution is designed to provide resilient fast failover and low latency to meet 5G network requirements and deliver cohesive end-to-end solutions across domains boundaries. This solution leverages the xHaul architecture and implementation guidance described in the Open radio access network (O-RAN) Alliance technical specification [ORAN-WG9.XPSAAS.0-v00.01] for transport network domain. The advent of 5G technologies requires transport networks to support resilient multidimensional and border agnostic architectures. This JVD incorporates use cases and requirements in Metro mobile Backhaul and 5G initiatives.

The foundational technologies of the 5G xHaul JVD incorporate current and legacy VPN services over segment routing. Interdomain decoupling of transport and service layers is accomplished by leveraging seamless segment routing with BGP Labeled Unicast (BGP-LU) and Prefix-SID. Network pre-slicing concepts using Flexible Algorithm and BGP Classful Transport (BGP-CT) create new layers of network abstraction and provide capabilities to support end-to-end service level agreements (SLA) requirements. A key focus of this solution is the ability to color-map individual overlay services, enabling granular support of 4G and 5G applications over xHaul infrastructure.

Summary of the platforms for 5G xHaul seamless segment routing solution is as follows:

Solution	AN	AG1	AG2	AG3	CR	SAG
5G xHaul Seamless Segment Routing	ACX710 ACX5448	ACX5448	MX204	MX480/ MX10003	PTX1000 and MX10003	MX10003

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.

Figure 1: 5G xHaul Reference Network

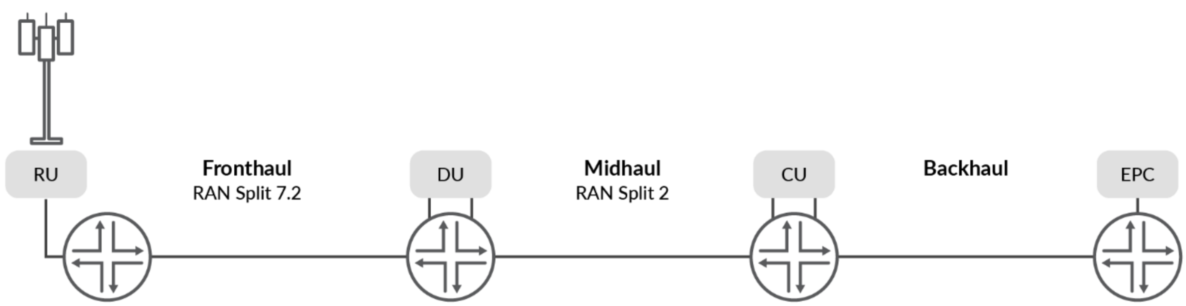

The validated design addresses the convergence of 5G xHaul by leveraging seamless segment routing technologies.

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].

Midhaul segment of the 5G network infrastructure defines the disaggregated RAN, responsible for supporting gNodeB-Centralized unit (gNB-CU) to gNB-DU communication. De-aggregated gNB consists of O-CU Control Plane (O-CU-CP) and CU User Plane (O-CU-UP), connected to the corresponding O-

DU over F1-C (control) and F1-U (data) interfaces. Interconnecting CU-CP and CU-UP by E1 interface specified by 3rd Generation Partnership Project (3GPP) TS 38.401.

Backhaul network defines the infrastructure connecting the 5G mobile core to CU. Control and user planes are assigned into unique VPNs for separation between user data and 3GPP control plane.

Midhaul and Backhaul architectures do not require the same latency budgets as Fronthaul and are represented as ring, mesh, hub-and-spoke, or spine-and-leaf topologies.

The validated design establishes these domains and device roles by following recommendations of O-RAN Alliance [ORAN-WG9.XPSAAS.0-v00.01]. The key attributes of the transport network architectures are considered with establishment of strict latency budgets, traffic prioritization, packet loss tolerance, fault management, and bandwidth requirements for meeting the demands across Fronthaul, Midhaul, and Backhaul segments. For the 5G xHaul JVD, the primary focus is on the mobile backhaul (MBH) transport architecture and services at scale. The infrastructure is expected to support diverse residential, business, and wholesale services concurrently with end-to-end 4G/5G mobile operator applications. The JVD proposes color-aware service mapping that enables granular end-to-end traffic steering across disparate domains.

The evolution of the RAN encompasses distributed, centralized, and virtual 4G architectures that must coexist with the 5G disaggregated O-RAN. These diverse ecosystems allow flexible insertion points for O-DU/O-CU components. The dissection of RAN functional splits proposed by 3GPP and O-RAN for supporting disaggregated deployment models are outside the scope of this document but an important attribute in the proposed architecture.

Figure 2 on page 4 as referenced by O-RAN Alliance [O-RAN.WG9.XPSAAS-v02.00], summarizes RAN deployment scenarios according to the ITU-T GSTP-TN5G for concurrent 4G and 5G support.

Figure 2: ITU-T Use Cases

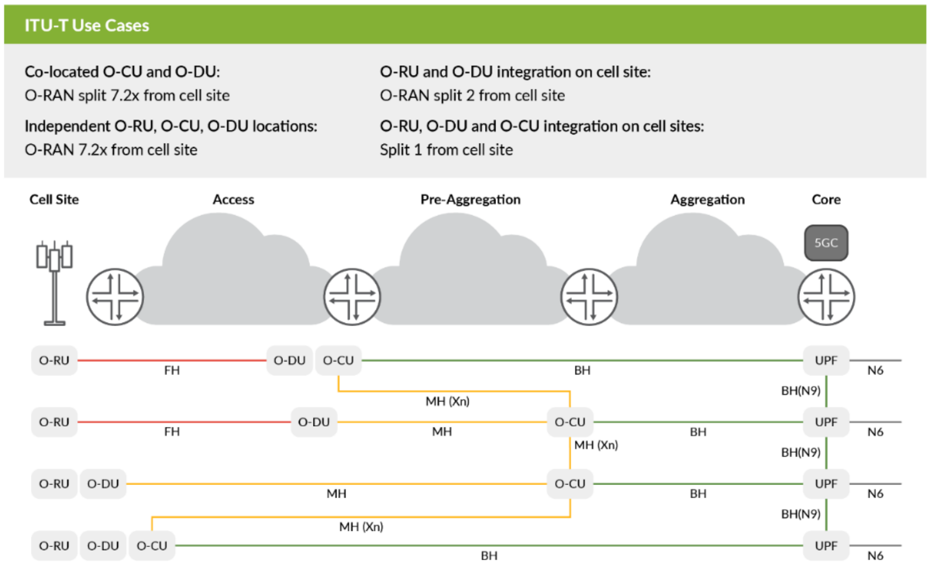

This JVD does not attempt to cover all possibilities and instead deploys an integrated Backhaul model. Additional insertion points for split architectures can be engineered to support disaggregation between Midhaul and Backhaul segments by extending appropriate services.

The ability to support complex concurrent network architectures further prioritizes delivery of seamless end-to-end solutions. A foundational emphasis of the 5G xHaul JVD is in demonstrating network pre-slicing solutions where overlay services might be mapped to underlay transport across a seamless converged mobile network.

Figure 3 on page 5 shows a seamless end-to-end MBH 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 Infrastructure

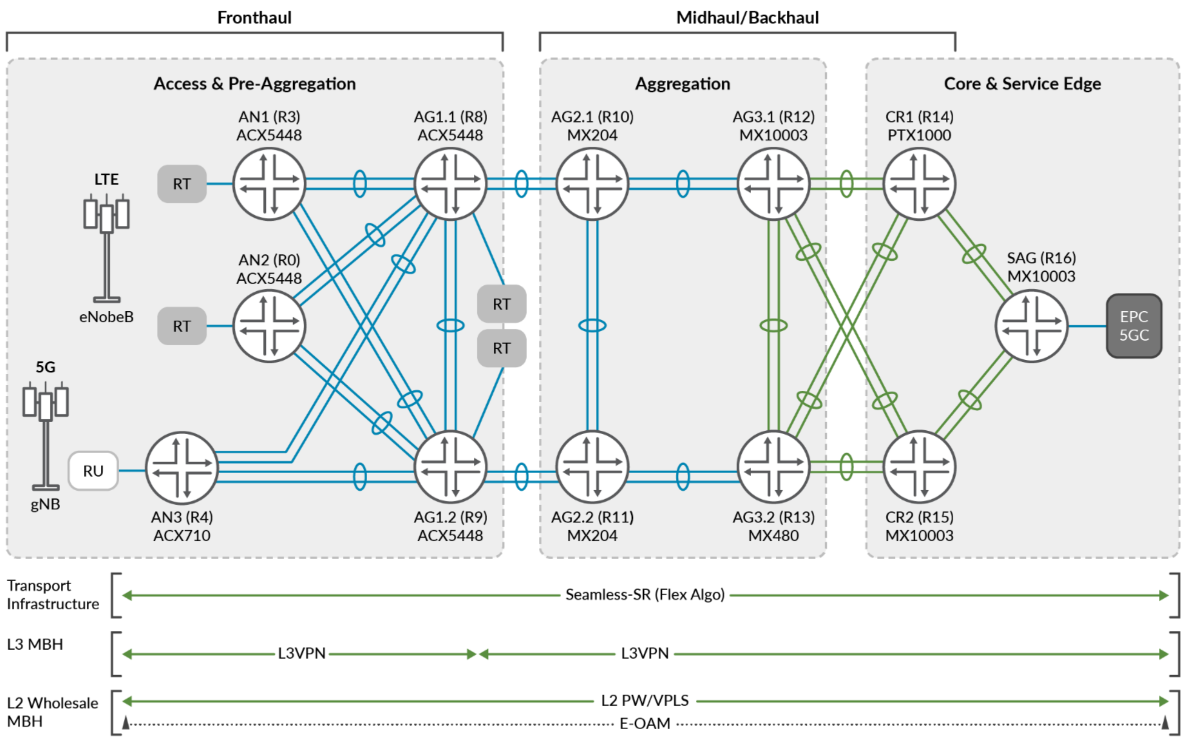

Topology Definitions:

- **Access**—Includes ACX710 & ACX5448 Access Nodes (AN1, AN2, AN3) cell site routers.
- **Pre-Aggregation**—Includes ACX5448 hub site routers AG1.1 and AG1.2.
- **Aggregation**—Includes aggregation routers AG2.1/AG2.2 (MX204) and AG3.1/AG3.2 (MX10003/MX480).
- **Core Network**—Includes core routers CR1/CR2 (PTX1000/MX10003) and service router SAG (MX10003).

Spine-leaf access topology establishes 5G Fronthaul segment with redundant hub site routers (HSR) – AG1.1 and AG1.2 in Figure 3 on page 5 and collapses 4G Pre-Aggregation and 5G HSR roles. Pre-Aggregation AG1 nodes include additional access insertion points (RT) for scaling the environment. Midhaul and Backhaul segments, represented by ring topologies, are combined to include Aggregation and Core roles.

Underlay consists of seamless segment routing L-ISIS, Prefix-SID, and BGP-LU stitching at border nodes, as shown in Figure 4 on page 6:

- Access nodes and Pre-Aggregation segments reside within Level 1 ISIS domain.
- Pre-Aggregation segment includes L1/L2 ISIS routers.

- Pre-Aggregation, Aggregation, and Core segments reside within a Level 2 ISIS domain.
- TI-LFA node redundancy is supported within each domain instantiation.
- Flex-Algo underlay slicing further partitions the SR-MPLS xHaul network into colored paths with two Flexible Algorithm Definitions (FAD), including red (128) and blue (129) using IGP and TE metrics.

Figure 4: 5G xHaul Converged Topology with Pre-Slicing Architecture

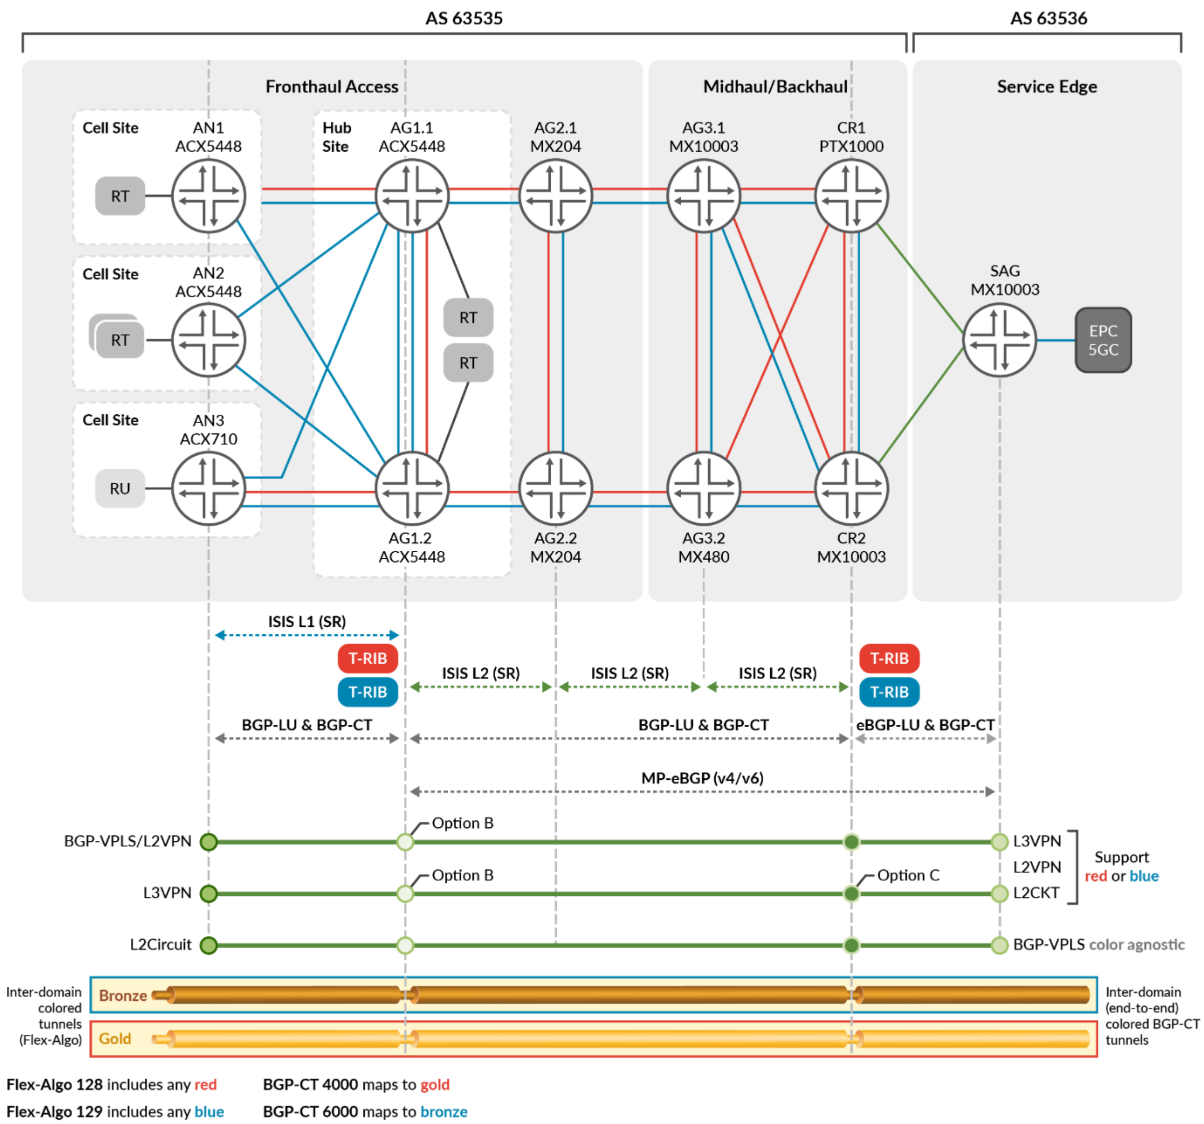

While BGP-LU and Inter-AS solutions can meet seamless E2E requirements, these do not satisfy service aware traffic engineering SLAs. BGP-CT provides network slicing and dynamic traffic steering based on transport classes, which are determined by route target attributes.

Color-aware overlay services are mapped end-to-end across domains with BGP-CT. Gold and bronze transport classes are defined so that gold is mapped to paths which include only red (128) algo and bronze is mapped to paths which include only blue (129) algo. In the event gold or bronze paths are unavailable, the network is designed to fallback to inet.3 (color agnostic) path selection.

The validated design includes the following overlay services:

- **Layer 3 VPN** includes termination points on all Access Nodes, AG1.1/AG1.2, and SAG. L3VPN supports BGP-CT per-prefix color-mapping. Routes received with protocol direct or OSPF are mapped to transport class gold and traverse the flex-algo slice denoted as red (128), while BGP-learned routes matching specific attributes are mapped to bronze transport class and blue algo (129).

- **Layer 2 VPN** includes termination points on all Access Nodes, AG1.1/AG1.2, and SAG. MP-BGP L2VPN supports BGP-CT color-mapping. All customer routes exported by the L2VPN instance are mapped to either gold or bronze transport classes and traverse only the corresponding flex-algo slice.

- **Layer 2 Circuit** includes termination points on AN3, AG1.1/AG1.2, and SAG. L2Circuit supports BGP-CT as a community mapping assignment. L2Circuits are assigned gold or bronze traffic class and traverse only corresponding algo.

- **BGP-VPLS** includes termination points on AN1, AN3, AG1.1/AG1.2, and SAG. VPLS is deployed to traverse the topology as color-agnostic services, demonstrating the coexistence with color-aware service mapping. BGP-VPLS supports BGP Classful Transport with service mapping from Junos OS Release 21.4R1 and later releases.

Solution and Validation Key Parameters

IN THIS SECTION

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

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

Table 1: Solution Design and Infrastructure

Attribute	Value
Transport	• Seamless SR: ISIS-SR, flex-algo, transport classes • BGP-CT/BGP-LU
Resilience	• ISIS L1 and L2 TI-LFA • Two Flex-Algos Definitions (128 red and 129 blue) • Route Reflectors, add-path, multipath, ECMP, BFD, OAM/CFM/LFM, LACP
Services	• L2VPN, L3VPN, BGP-VPLS, L2 Circuit • Color-based service mapping for all services except VPLS • Inter-AS Options B/C
BGP-CT	• Bronze (6000) and Gold (4000) • Bronze maps to blue Algo • Gold maps to red Algo

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 Networks 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 2: Scale and Performance

Service	AN1—ACX5448	AN2—ACX5448	AN3—ACX710	AG1—ACX5448	AG2—ACX5448	SAG—MX10003
L3VPN	10	100	100	320	120	320
L2VPN	10	10	100	10	1000	1130
L2CKT	0	0	100	700	1000	1800
VPLS	0	10	100	15	1000	1125
VLANS	30	121	410	1045	3120	4376
RIB	22K	278K	290K	1.1M	193K	1.2M
FIB	11K	102K	102K	492K	95K	504K

Key Feature List

Validation includes but not limited to the following key features spin on DUT nodes:

- L3VPN
- VPLS
- L2Circuit

- L2VPN
- Segment Routing ISIS
- TI-LFA (link/node)
- Flex-Algo (IGP/TE metrics)
- ISIS, BGP
- BGP-LU
- BGP-CT
- Per-Service/Prefix color mapping
- Flex-Algo with Transport Class
- Bidirectional Forwarding Detection (BFD)
- Community-based Routing Policy
- Route Reflection
- IPv4
- IPv6
- Link Aggregation Control Protocol (LACP),
- Aggregated Ethernet (AE)
- Connectivity fault management (CFM),
- Link fault management (LFM),
- VLAN (802.1q).

Contact your Juniper Networks representative for a complete feature list.

Solution Validation Goals

Here are the major test goals:

- Validation of network stability and pre-slicing functionalities for Seamless SR 5G xHaul design architecture:
    - End-to-End Seamless SR across ISIS L1 and L2 domains and Inter-AS BGP-LU stitching.

    - Flex-Algo with red (128) and blue (129) FADs IGP & TE metrics.
    - BGP-CT creation of transport RIBs for underlay and service mapping functions using bronze and gold extended color communities.
    - Color-mapping results in deterministic path selection and traffic steering.
- Validation of overlay service mapping (service slicing):
    - BGP-CT service/prefix mapping per supported VPN type (L3VPN, L2VPN, and Layer 2 circuit (L2CKT))
    - Validate BGP-CT Resolution Scheme fallback operations.
    - BGP-VPLS performs color agnostic path selection.
- Validation of network stability for major 4G/5G traffic flows at scale with each VPN service type over Seamless SR during normal and stress conditions.
- Validation of network resiliency, the ability for seamless traffic restoration, and measured convergence time for ACX710 AN and ACX5448 AN/AG, along with adjacent links and node failures for all services/traffic flows.
    - Solution resilience of layer 2/3 flows from AN to AG1, AG1 to SAG and AN to SAG (EPC).
- Validation of consistency and resiliency of the ACX710/ACX5448 systems against negative stress conditions (enable or disable control and data plane daemons, add or delete configurations, and so on).
- Identify product limitations and anomalies exposed during validation stages.

Solution Validation Non-Goals

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

- Class of Service (CoS).
- Underlay MPLS/SR transport other than specified in the Solution Goals section.
- Multi-Instance ISIS.
- OSPF as L3VPN routes only.
- EVPN over colored transport BGP-CT.
- Management and Automation.

- BGP PIC-Edge at border routers.
- End-to-End Timing and Synchronization Distribution:
  - Synchronous Ethernet
  - IEEE1588v2
- SLA Monitoring: RFC 2544, Y.1564, TWAMP, Active Assurance
- Telemetry
- Loopback filters on DUT

The regulatory trends are as follows:

- 3GPP Forum
- O-RAN Alliance

Results Summary and Analysis

The JVD team validated the reference design of the converged 5G xHaul network and MBH architectures over seamless segment routing with color-mapping pre-slicing capabilities supported by ACX710 and ACX5448 Universal Metro Routers. ACX710 and ACX5448 DUT routers with Junos OS Release 21.2R3 successfully passed the validation test cases with scale and performance parameters observed within anticipated limits.

The validation testing focused on delivering a Seamless MBH solution across BGP-LU stitched domains with ACX710 and ACX5448 in roles supporting 4G/5G Access Node and Pre-Aggregation packet-switched transport CSR-to-SAG, CSR-to-HSR, and HSR-to-SAG.

Contact your Juniper Networks representative for test result reports.

The JVD validation demonstrated a robust solution for 5G xHaul transport infrastructure utilizing Seamless SR and color-aware network pre-slicing concepts. The architecture provides for resilient and agile deployment, spanning multiple IGP domains and autonomous systems with end-to-end flexibility, path management, and traffic steering.

Key results are as follows:

- JVD topology generates reasonable multi-vector scale of L2/L3 connectivity services as compared with mobile network operator (MNO) and metro area network (MAN) operator expectations for real network deployments, while satisfying SLA requirements.

- The scale reference characterizes primary multidimensional KPI's represented in the validated profile.
- With given 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 while node failures were more disruptive than expected, exacerbated by scale and in some rare cases, production limitations.

ACX710 and ACX5448 Universal Metro Routers with Junos OS release 21.2R3 can deliver the solutions outlined herein across intra and interdomain architectures. Both platforms are ideally situated for access roles and ACX5448 Universal Metro Router can be further leveraged in an aggregation HSR role.

Load sharing operations are enabled on all devices for ECMP opportunities across Fronthaul, Midhaul, and Backhaul segments. Several ECMP mechanisms were configured including adjusting IGP metrics, BGP multipath and add-path, and ECMP fast-reroute. Colored underlay slices are available for all supported services. For example, bronze services are load shared across only blue paths and gold services are load shared across only red paths. Contact your Juniper Networks representative for the colored traffic distribution information. Color agnostic services are shown to always select and disperse over best paths. The distribution of traffic flows resulted in significantly reduced impact for most failure scenarios.

Contact your Juniper Networks representative for hash computation and ECMP validation information.

The goal of this JVD is to identify anomalous behaviors which might impact solution design.

The solution resiliency was validated against multiple stress conditions and failure scenarios which include but not limited to:

- Adjacent link failures
- Indirect link failures
- Node failures

Table 3 on page 14 summarizes some key results for Layer 2/Layer 3 services traffic flows restoration under failure conditions. The color-coded Events column represents the expected network slice that must be impacted by the described failure condition. For all test scenarios, we run nearly 200 traffic flows—with each bundling congruent VPN services and encompassing all VPN types—between each Access Node (CSR) to AG1s (HSR), ANs to SAG and AG1s to SAG. In greater majority of cases, the failure recovery measured <50ms but always the traffic item with the highest impact is recorded.

For example, in Table 3 on page 14, the event AN2(R0) link disable (R0-R4), a red link is disabled between R0-R4, and the expected impact is to services which are mapped to red or color-agnostic.

Table 3: Convergence Times for Access Failure Events Per Flow Type

	L3VPN (ms)			L2VPN (ms)		L2CKT (ms)		VPLS (ms)
Events	FG-128	FG-129	IPv6 (no FG)	FG-128	FG-129	FG-128	FG-129	No FG
AN2(R0) link disable (R0-R4)	263	0	0	2	0	0	0	0
AN2(R0) link enable (R0-R4)	186	0	0	2	0	0	0	1
AN3(R4) link disable (R4-R9)	0	14	0	0	3	0	290	0
AN3(R4) link enable (R4-R9)	0	0	0	0	5	0	5	2
AN3(R4) link disable (R4-R9)	8	0	0	8	0	0	0	2
AN3(R4) link enable (R4-R9)	143	0	0	10	2	0	5	4

Table 4 on page 15 lists the results that were initially measured with higher failover. It was determined that the convergence delay was ECMP routing table update following the failure and hash algorithm calculating the new paths. By enabling equal-cost-multipath fast reroute protection (ecmp-fastreroute) on the SAG (R16), we were able to refresh ECMP information without waiting for route table update.

Table 4: Convergence times for Pre-Aggregation and Aggregation Failure Events Per Flow Type

	L3VPN (ms)			L2VPN (ms)		L2CKT (ms)		VPLS (ms)
Events	FG-128	FG-129	IPv6 (no FG)	FG-128	FG-129	FG-128	FG-129	No FG
AG1.2 (R9) link disable (R9-R8)	10 (1093)	0	0	11 (1094)	0	1	0	1
AG1.2 (R9) link enable (R9-R8)	0	0	0	5	0	5	0	1
AG1.2 (R9) link disable (R9-R8))	0	0	0	0	3	0	5	0
AG1.2 (R9) link disable (R9-R8)	0	0	0	4	5	0	2	1
AG1.2 (R9) link disable (R9-R11)	34	0	7	40	0	43	0	8 (6388)
AG1.2 (R9) link disable (R9-R11)	0	0	0	2	2	3	10	2
AG1.2 (R9) link disable (R9-R11)	0	13	1	0	18	0	20	2

Table 4: Convergence times for Pre-Aggregation and Aggregation Failure Events Per Flow Type *(Continued)*

	L3VPN (ms)			L2VPN (ms)		L2CKT (ms)		VPLS (ms)
AG1.2 (R9) link disable (R9-R11)	0	0	0	5	2	2	1	0
AG2.2(R 11) link disable (R11-R13)	63	3	33	66	0	66	0	45
AG2.2(R 11) link enable (R11-R13)	0	0	0	2	0	1	0	0
AG2.2(R 11) link disable (R11-R13)	0	59	29	0	56	0	54	45
AG2.2(R 11) link enable (R11-R13)	0	6	0	0	11	0	12	0

Table 5: Convergence Times for Core and SAG Failure Events Per Flow Type

	L3VPN (ms)			L2VPN (ms)		L2CKT (ms)		VPLS (ms)
Events	FG-128	FG-129	IPv6 (no FG)	FG-128	FG-129	FG-128	FG-129	No FG
AG3.2 (R13) link disable (R13-R15))	0 (38)	0	14	0 (45)	0	0 (36)	0	22
AG3.2 (R13) link enable (R13-R15)	1	0	0	5	0	4	0	0
AG3.2 (R13) link disable (R13-R15)	0	43	19	0	50	0	44	28
AG3.2 (R13) link enable (R13-R15)	0	0	0	0	5	0	4	0
CR2 (R15) link disable (R15-R16)	57	69	67	60	67	63	64	61 (2006)

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

	L3VPN (ms)			L2VPN (ms)		L2CKT (ms)		VPLS (ms)
Events	FG-128	FG-129	IPv6 (no FG)	FG-128	FG-129	FG-128	FG-129	No FG
CR2 (R15) link enable (R15-R16)	0	0	0	0	0	0	0	0

There is further delay in scenarios where global repair is triggered rather than managed by Broadcom ASIC with HW FRR. Additionally, BFD FRR is not supported on ACX5448 or ACX710 Universal Metro Router. Adjusting revertive timers can help control service restoration but overall, these results were mostly within expectation.

Recommendations

ACX710 Universal Metro Router can be recommended as CSR in small to midsize 4G/5G MBH applications with access segments consisting of 50-70 access nodes per IGP domain (multiple IGP domains can comprise the access and aggregation segments). ACX5448 Universal Metro Router can be used in equivalent and expanded roles, including pre-aggregation and aggregation for midsized engagements.

Seamless segment routing is a recommended underlay architecture to support end-to-end stitching with BGP-LU and BGP-CT transport classes across multiple IGP and inter-AS domains. Using TI-LFA and ECMP mechanisms, the solution supports rapid failover and resilience. Network pre-slicing is accomplished with color-mapping overlay services (L3VPN, L2VPN, and L2Circuit) to Flex-Algo underlay tunnel slices. In parallel, color agnostic services coexist, taking advantage of the same infrastructure. Both ACX710 and ACX5448 Universal Metro Routers support these protocols and services to enable end-to-end converged 5G MBH.

This JVD uses Junos OS Release 21.2R3 as a minimum version to support the validated design.

Contact your Juniper Networks representative for the full test report with details about convergence times for all failure events.

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