EANTC

Multi-Vendor MPLS SDN Interoperability Test Report

2023

MPLS SD&AI NET WORLD23

Table of Contents

Editor's Note

Going back to Paris in person for the MPLS & SDN World Congress in Spring 2023 is a wonderful thought! Although it is legitimate to read this test report at any place and time, it is best digested on-site, while watching 78 participating devices working as part of the live showcase.

This year, EANTC proudly presents one of the richest, most advanced multi-vendor interoperability showcases we have ever staged. With 17 participating vendors and more than 1,230 successful test combinations achieved, this year's event has certainly overcome any pandemic-related limitations in participation and technical coverage that we still felt at the previous conference edition last year.

The real miracle, in fact, is how the vital ecosystem between vendors and network operators in the packet transport area has remained alive and innovative over such a long time. From the early 2000s when we started MPLS interoperability testing at EANTC, we have always seen 10-20 leading manufacturers participating with distinct yet standards-based, interoperable products. At the same time, network operators—whether carriers, enterprises, or government organizations—have always helped to keep the market diverse. MPLS and its integrated successors continued to flourish while other telecom technologies rose and fell.

MPLS and now SDN/Segment Routing have brought a constant flow of innovations large and small, carrying the technology through the times while most often offering viable upgrade paths (to avoid the word "seamless"). Fights over evolution directions have been carried out and resolved, for example between Segment Routing with MPLS, with VXLAN, or over IPv6. Meanwhile, SRv6 has been fully integrated in the main technology path and 10 vendors have collaboratively shown advances in SRv6 at our event this year.

All this, of course, comes at a cost: Complexity. A technology solution so all-encompassing that it caters for large and small network needs alike, always adapting itself to the latest feature requirements, and being backwards compatible in many ways will inevitably grow into a complex system of protocols and alternatives. For this reason, our test plans and reports have kept growing. This is an effect we are happy to live with. Now let's introduce some of the most important test results:

In the Segment Routing area, vendors put a major focus on SRv6 innovations. Most notably, we tested multi-vendor interoperability of micro segment IDs (µSID) successfully with 11 implementations. This is an industry-first achievement and an important result of negotiations in the Internet Engineering Task Force (IETF), where multiple solutions competed and a single compromise was chosen that works for all vendors. In general, 10 vendors participated in SRv6 testing—more than ever.

That said, SR-MPLS is alive as well and multi-vendor tests of network slices was conducted, confirming that end-to-end slicing in the transport domain is possible for 5G workloads. We also performed performance tests of failover scenarios with TI-LFA, demonstrating that the failover times even in multi-vendor deployments can be reduced to below 35 milliseconds.

EVPN service tests were bustling as always. All basic tests have been completed years ago, and EVPNs are a stable, standardized technology. This year, we focused on IPv6 support for both underlay and overlay. Additionally, vendors tested tunnel stitching with VXLAN.

New Ethernet generations usually find their way into our event quickly; this year, we tested 400 Gigabit Ethernet long-range (ZR) optics with two vendors, proving that multi-vendor 400G connections carrying Segment Routing services are already viable.

We also continued to tackle the complex area of multi-vendor management plane interoperability. SDN controller interaction (PCEP) lacked support for PCC-initiated paths or SRv6; this area has not shown major improvements in 2023. In NETCONF testing, we observed multi-vendor support for standardized OpenConfig models for L2 VPN and L3 VPN services (there were no new tests for IETF Yang models, as support for these models is decreasing). For the first time, one single controller was able to interoperate with routers from four vendors based on standardized Yang models. Vendors added telemetry testing as well.

Finally, in the area of clock synchronization, we benefitted again from an expert group of vendors, most of which have worked together in our interop events for many years. The tests focused on new ITU Class C und D clock requirements. Participants achieved 100 ns (nanoseconds) precise synchronization regularly, and to just 5 ns in some optimal multi-vendor test cases. We tested Enhanced SyncE with six vendors for the first time, reducing the time to frequency synchronization lock and extending holdover times. Vendors built a chain of seven boundary clocks, proving that precise synchronization is possible in such longer chains. Finally, we successfully evaluated three topologies for Open Fronthaul synchronization in mobile Open RAN scenarios.

Altogether, the two weeks of hot-stage testing were as busy as ever. The most rewarding sight is to witness engineers from all vendors working together with a single goal: Making advanced SDN/MPLS-based multi-vendor, standards-based telecom transport networks a reality. This test report provides much more detail to document the state of the art of participating vendor implementations. We hope the report will provide new insights and inspirations!

Our Mission

For over 30 years in the industry, accomplishing various networking projects of different types in all networking fields, EANTC has gained a unique position with special and deep experiences in networking fields and testing. These experiences and knowledge are driving us year after year to bring the leading vendors in the world to test their newest innovations, technologies, and devices, inspecting all the possible interop challenges and issues, giving them the opportunity to discover any anomalies in their Hardware/Software when they connect to each other devices.

In one sentence, our mission is to push the industry and innovation further, finding the boundaries, and stretching them further more, helping the service providers to evolve, renew, and expand their multi-vendors networks with less troubleshooting and difficulties in field.

Why These Areas

The testing areas were selected through thorough analysis and discussions with our partners, aiming to encompass all aspects of service provider networks. Moreover, we consistently prioritize the latest and most relevant topics each year.

Working Process

Preparations:

The preparations for the MPLS/SDN interoperability 2023 began in the fall of the previous year. We initiated discussions about test areas, test ideas, and general test plans with all participating vendors during a kickoff call, followed by three rounds of technical calls. In these calls, we thoroughly discussed technical details, new testing ideas, and the latest standards, ultimately shaping our test plans to be state-of-the-art.

Hot-Staging:

By beginning of March, everything was set, and ready. Newest hardware with latest versions of software, from all over the world were already packed in EANTC lab in Berlin, waiting for the starting signal. Two weeks of non-stopping testing, deep and extensive discussions, racing time to solve some emerged issues, resulted in great testing results for all our vendors.

EANTC engineers observed, and verified all the tests in details, following the test procedures and pre-defined test steps, they judged all tests results independently.

Interoperability Test Results

As usual, this test reports documents only positive results (passed test combinations) individually with vendor and device names. Failed test combinations are not mentioned in the diagrams; they are referenced anonymously to describe the state of the industry. Our experience shows that participating vendors quickly proceed to solve interoperability issues after our test so there is no point in punishing them for their willingness to learn by testing. Confidentiality is vital to encourage manufacturers to participate with their latest - beta - solutions and enables a safe environment in which to test and learn.

Terminology

We use the term tested when reporting on multi-vendor interoperability tests. The term demonstrated refers to scenarios where a service or protocol was evaluated with equipment from a single vendor only.

Test Equipment

With the help of participating test equipment vendors, we generated and measured traffic, emulated and analyzed control and management protocols, and performed clock synchronization analysis.

We thank Calnex, Keysight, and Spirent for their test equipment and support throughout the testing.

As in previous events, a number of Calnex instruments were used in the Time Synchronization test cases. Paragon-neo was used to generate and measure PTP and 1PPS signals with sub-ns accuracy and 250ps resolution, enabling characterization of devices to Class D performance and networks to Level 6C at line rates from 1GbE to 400GbE (configurable from 100MbE to 400GbE). For network tests, the new Calnex Sentry had its debut in a plugfest, allowing four (configurable up to six) 1PPS signals to be measured simultaneously, enabling sync across a network to be monitored, or multiple tests to be run simultaneously, saving time and allowing more to be achieved in the same time. For network tests running at 10G, we also used the inline impairment capabilities of Paragon-X to emulate conditions for PTP in a real-world network environment.

Participants and Devices

Participants	Devices
Arista	7050SX2 7050SX3 7050X3 7280R 7280R2 7280R3
Arrcus / UfiSpace	UfiSpace S9600-72XC UfiSpace S9710-76D
Calnex	Paragon-neo Paragon-neo PAM4 Paragon-X Sentry SNE
Ciena	5166
Cisco	8201-24H8FH 8201-32FH ASR 9901 ASR 9902 ASR 9903 Crosswork NCS 540-24Q8L2DD NCS 540-28Z4C NCS 540X-12Z16G NCS 540X-16Z4G8Q2C NCS 57B1-5DSE NCS 57C1-48Q6 IOS XRd
Ericsson	6273 6673 6675
HPE	Aruba CX8325-48Y8C Aruba CX8360-48Y6C Aruba CX9300-32D Aruba CX10000-48Y6C
Huawei	NCE-IP NetEngine 8000 F8 NetEngine 8000 M4
Intel	E810-XXVDA4T E810-QCDA2T
Juniper	ACX7024-AC ACX7100-32C MX204 Paragon Pathfinder PTX10001-36MR QFX5110 QFX5120 QFX5130
Keysight	IxNetwork
Microchip	TimeProvider 4100
Nokia	7750 SR-1 Network Services Platform (NSP)
RARE/freeRtr	BF2556X-1T
Ribbon	NPT-2100A
Spirent	STC

Table 1: Participating Vendors and Devices

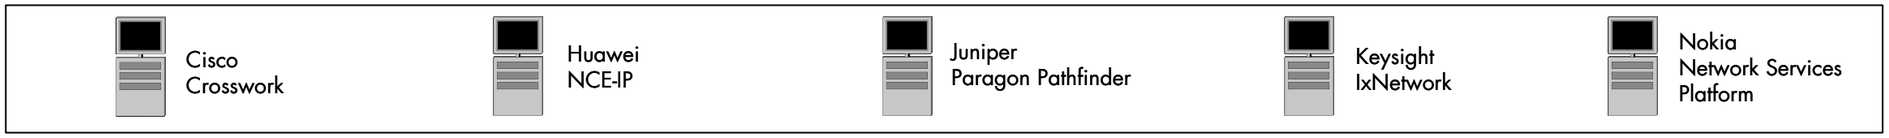
Cisco
Crosswork
Huawei
NCE-IP
Juniper
Paragon Pathfinder
Keysight
IxNetwork
Nokia
Network Services
Platform

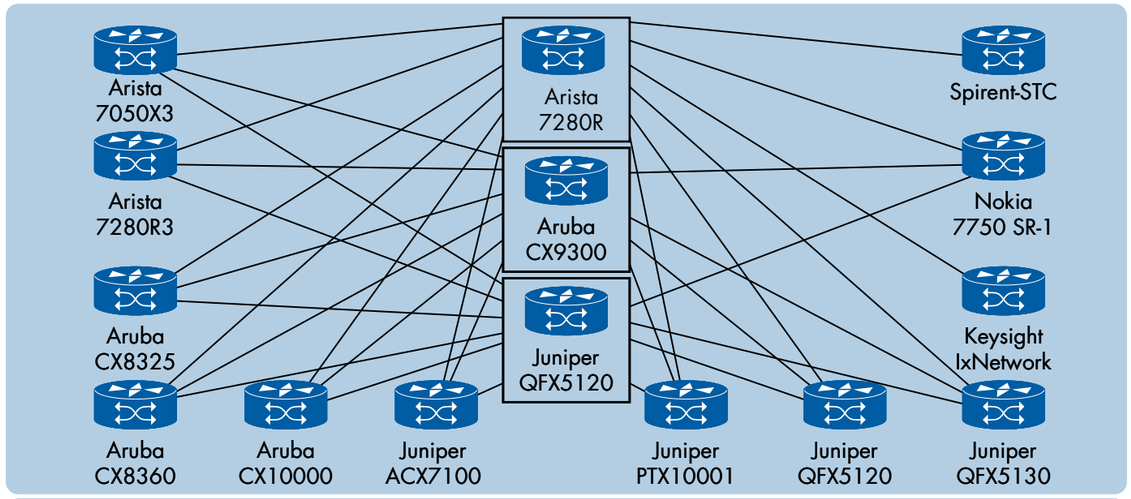
Arista
7050X3
Arista
7280R3
Aruba
CX8325
Aruba
CX8360
Aruba
CX10000
Juniper
ACX7100
Arista
7280R
Aruba
CX9300
Juniper
QFX5120
Juniper
PTX10001
Juniper
QFX5120
Spirent-STC
Nokia
7750 SR-1
Keysight
IxNetwork
Juniper
QFX5130

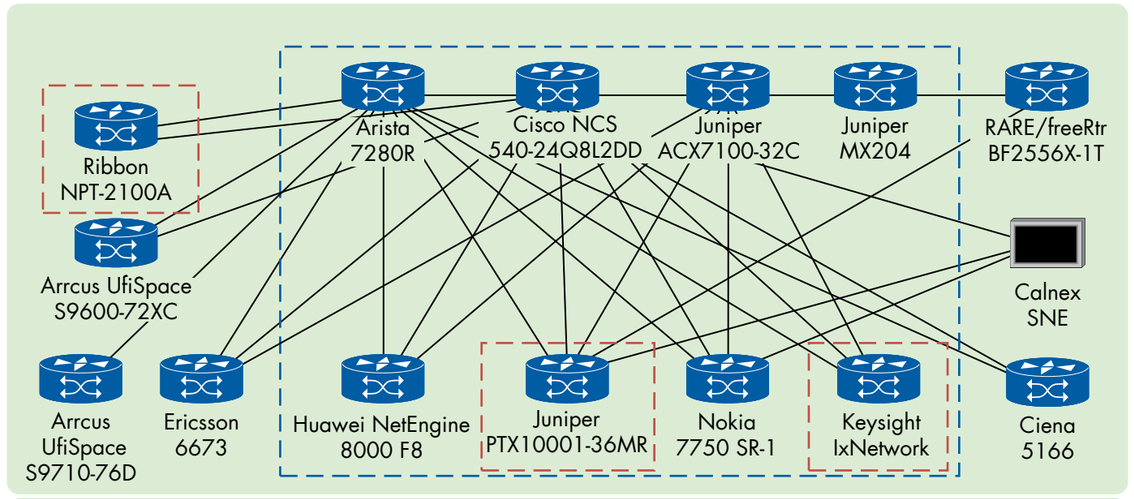
Ribbon
NPT-2100A
Arrcus UfiSpace
S9600-72XC
Arrcus
UfiSpace
S9710-76D
Ericsson
6673
Arista
7280R
Huawei NetEngine
8000 F8
Cisco NCS
540-24Q8L2DD
Juniper
PTX10001-36MR
Juniper
ACX7100-32C
Nokia
7750 SR-1
Juniper
MX204
Keysight
IxNetwork
RARE/freeRtr
BF2556X-1T
Calnex
SNE
Ciena
5166

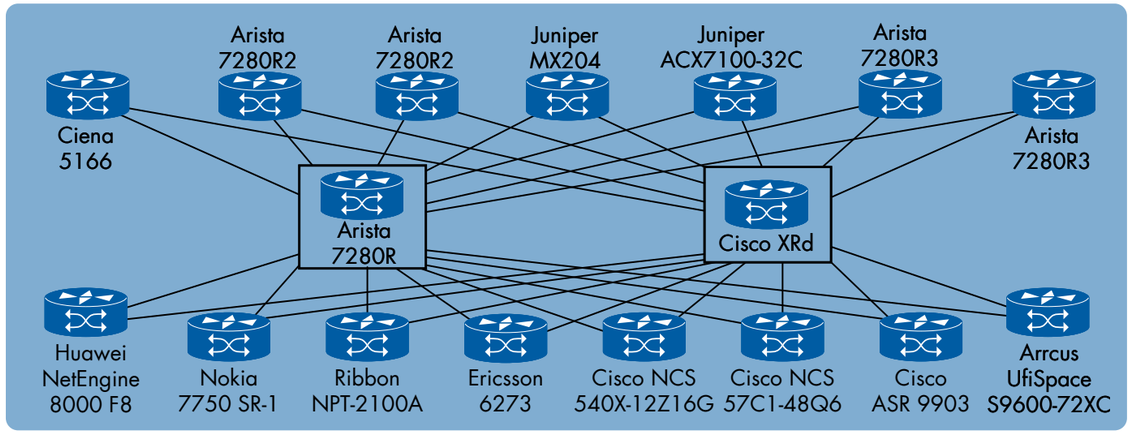
Ciena
5166
Arista
7280R2
Arista
7280R2
Arista
7280R
Juniper
MX204
Juniper
ACX7100-32C
Cisco XRd
Arista
7280R3
Arista
7280R3
Huawei
NetEngine
8000 F8
Nokia
7750 SR-1
Ribbon
NPT-2100A
Ericsson
6273
Cisco NCS
540X-12Z16G
Cisco NCS
57C1-48Q6
Cisco
ASR 9903
Arrcus
UfiSpace
S9600-72XC

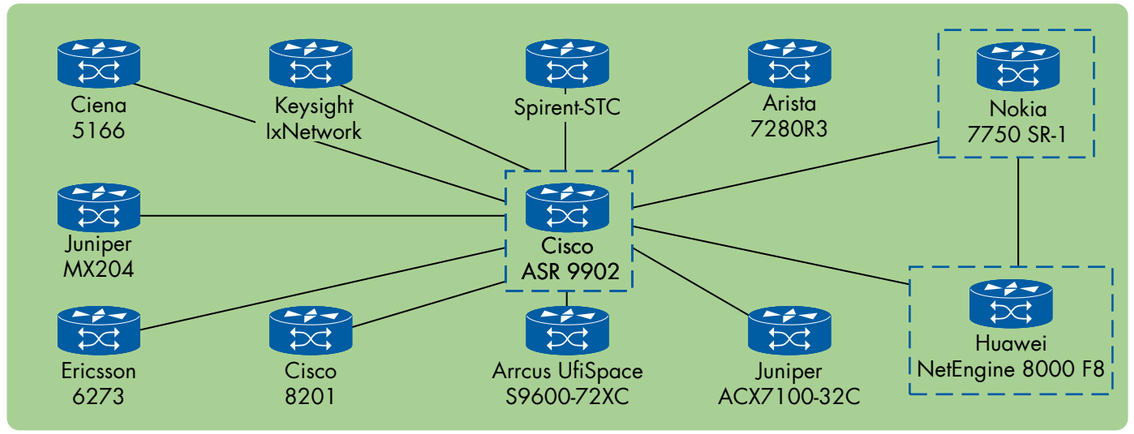
Ciena
5166
Keysight
IxNetwork
Spirent-STC
Arista
7280R3
Nokia
7750 SR-1
Juniper
MX204
Cisco
ASR 9902
Huawei
NetEngine 8000 F8
Ericsson
6273
Cisco
8201
Arrcus UfiSpace
S9600-72XC
Juniper
ACX7100-32C

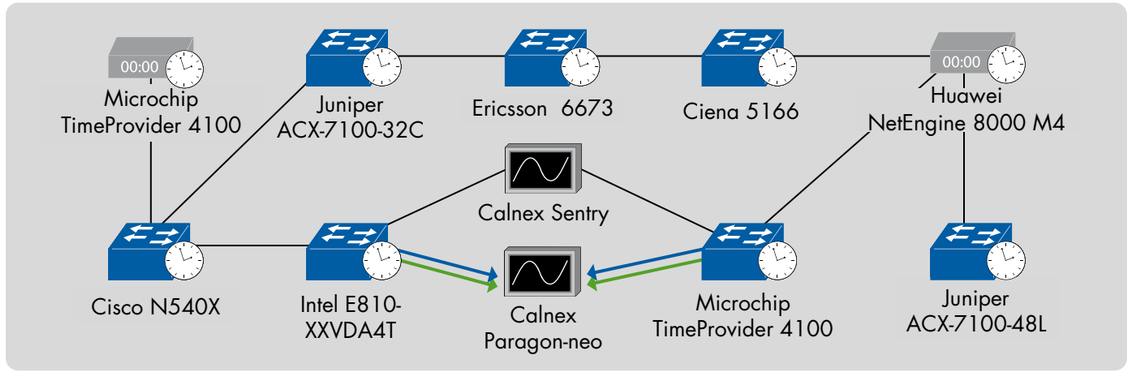
00:00
Microchip
TimeProvider 4100
Juniper
ACX-7100-32C
Ericsson 6673
Ciena 5166
00:00
Huawei
NetEngine 8000 M4
Calnex Sentry
Cisco N540X
Intel E810-
XXVDA4T
Calnex
Paragon-neo
Microchip
TimeProvider 4100
Juniper
ACX-7100-48L

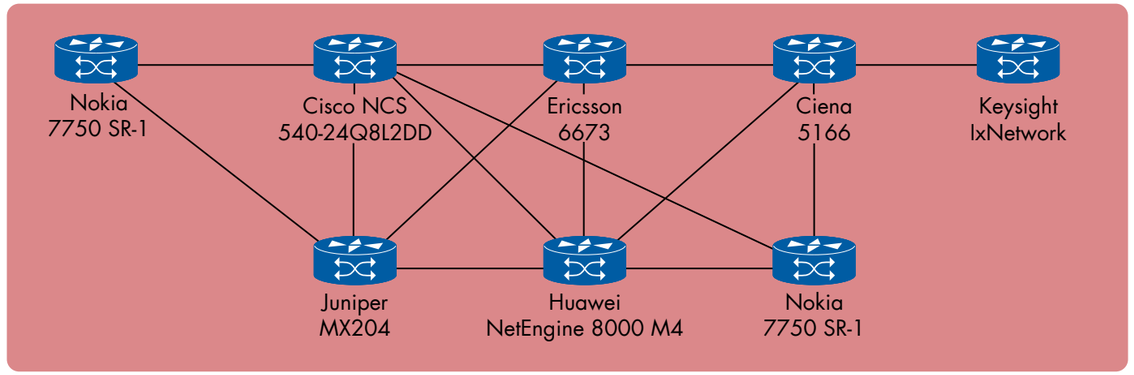
Nokia
7750 SR-1
Cisco NCS
540-24Q8L2DD
Ericsson
6673
Ciena
5166
Keysight
IxNetwork
Juniper
MX204
Huawei
NetEngine 8000 M4
Nokia
7750 SR-1

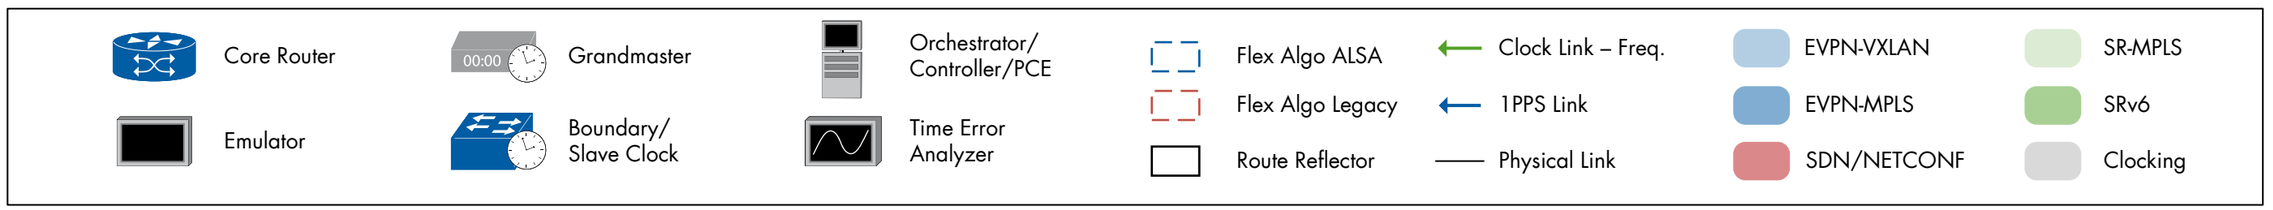
Core Router
Grandmaster
Orchestrator/
Controller/PCE
Flex Algo ALSA
Clock Link – Freq.
EVPN-VXLAN
SR-MPLS
Emulator
Boundary/
Slave Clock
Time Error
Analyzer
Flex Algo Legacy
1PPS Link
EVPN-MPLS
SRv6
Route Reflector
Physical Link
SDN/NETCONF
Clocking

EVPN

EVPN (Ethernet Virtual Private Network) was initially conceived as a BGP-based Layer 2 VPN technology that provides a scalable and efficient way of extending Layer 2 domains over a WAN (Wide Area Network). Over time EVPN became the de-facto VPN standard, not only for Layer 2 VPNs but also for Layer 3, multicast, and other advanced VPN services.

EVPN has become increasingly popular in data centers as it provides a mechanism for distributing MAC (Media Access Control) addresses across the network, which is essential for efficient and flexible VM (Virtual Machine) creation and mobility. EVPN also enables network administrators to create tenant-specific virtual networks that can span multiple data centers, making it an ideal solution for multi-tenant environments.

EVPN supports advanced features such as Network slicing, fast convergence, load balancing, and multi-path forwarding. These features are critical for providing high availability and efficient use of network resources in data centers and 5G networks. Network slicing as a key feature of 5G networks, enables network operators to create multiple virtual networks, each with its own characteristics and service levels, on a single physical infrastructure.

EVPN E-LAN Service

E-LAN is a versatile and easily adjustable networking service that leverages BGP (Border Gateway Protocol) to establish secure and uninterrupted communication channels across remote sites. A multipoint-to-multipoint Ethernet VPN seamlessly interconnects customer sites while presenting each location as a single Ethernet segment to all other sites. EVPN disrupts conventional forwarding-plane MAC address learning and uses BGP extensions for control-plane MAC learning and transmitting. Additionally, this service allows for All-Active Multi-Homing, enabling traffic load-balancing among multi-homed PEs.

We performed the test once with all the vendors' devices participating simultaneously in a mixture of single-homed and multi-homed devices. The second run consisted of multi-homed multi-vendor devices.

In this test, we verified the establishment of ISIS and BGP sessions and the EVPN signaling. In the next step, we observed the DF and non-DF PEs in single-active multi-homed devices and flow-based traffic load balancing in all-active multi-homed devices. Then, we verified zero packet loss during Any-to-Any bidirectional unicast traffic generation, and lastly, we measured the link failover and link recovery out of service time.

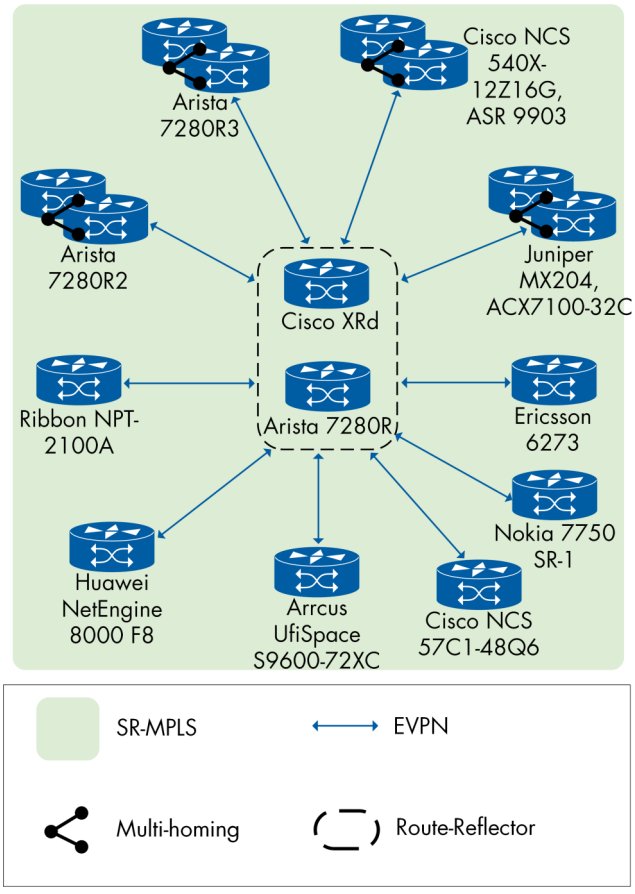

Figure 1: E-LAN Service

Run	Sita A Peerings	Sita B Peerings
1	Cisco NCS 540X-12Z16G with Arista 7280R2	Huawei NetEngine 8000 F8 with Ericsson 6273
2	Cisco NCS 540X-12Z16G with Arista 7280R2	Huawei NetEngine 8000 F8 with Nokia 7750 SR-1
3	Arista 7280R2 with Nokia 7750 SR-1	Huawei NetEngine 8000 F8 with Ribbon NPT-2100A
4	Arista 7280R2 with Juniper MX204	Huawei NetEngine 8000 F8 with Ribbon NPT-2100A
5	Arista 7280R2 with Juniper ACX7100-32C	Huawei NetEngine 8000 F8 with Ribbon NPT-2100A

Table 2: E-LAN Service Multi-vendor Multi-homing

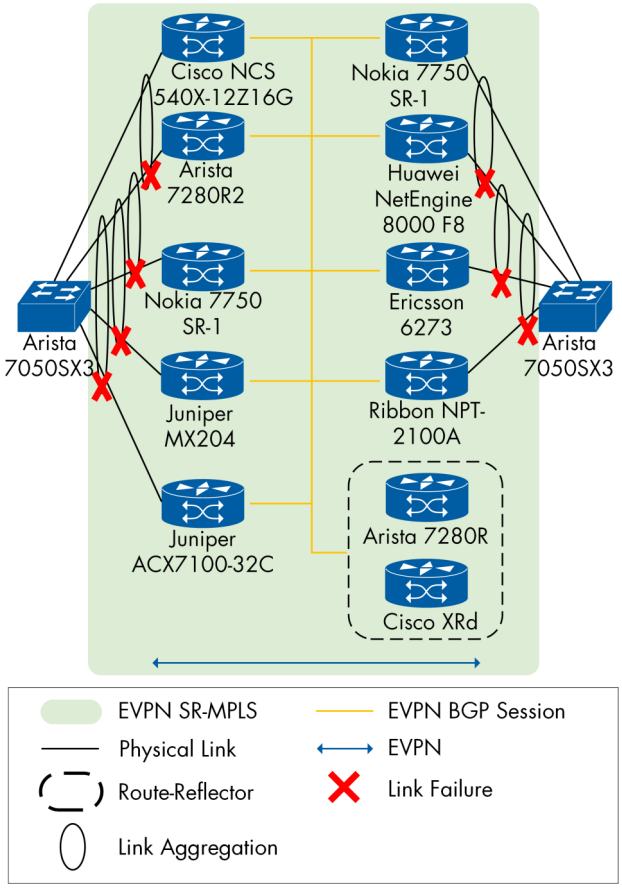

Figure 2: E-LAN Service
Multi-vendor Multi-homing

Arista 7280R2, Arista 7280R3, Cisco NCS 540X-12Z16G with ASR 9903, and Juniper MX204 with ACX7100-32C participated as multi-homed PEs.

Cisco NCS 57C1-48Q, Arrcus UfiSpace S9600-72XC, Huawei NetEngine 8000 F8, Nokia 7750 SR-1, Ericsson 6273, and Ribbon NPT-2100A participated as single-homed PEs.

In both scenarios, Arista 7050SX3 was the CE device, Spirent-STC the Traffic Generator.

VLAN-based and VLAN-aware Bundle Symmetric IRB

EVPN is an ideal solution that succeeds the VPWS and VPLS services. Once EVPN enters the Data Center network, the inter-subnet routing becomes mandatory with the tenant numbers increasing. Therefore, RFC 9135 and 9136 offer the solution for the inter-subnet routing function, which is integrated routing and bridge function. RFC 9136 introduces the EVPN RT-5 to solve the interaction between MAC-VRF and IP-VRF because IP-VRF sometimes summarizes the route, and RT-5 would help the remote end to understand the subnet behind it.

Virtual Extensible LAN (VXLAN) is a technology used to solve the scalability issue of VLAN. In the DC network, tons of tenants need their network isolated from each other. However, VLAN has a hard limitation of 4096, which is far below the needs in DC, whose tenants can be more than 10K most of the time. VXLAN is proposed under such background to solve the VLAN limitation and give the tenant the flexibility to extend and control their traffic.

We performed three runs of the test case in both multi-homed and single-homed scenarios in the VXLAN area and four runs in single-homed, multi-homed, and multi-homed multi-vendor with Route Type-5 and VPNv4 routes in the SR-MPLS area.

In VXLAN area, the first run was the multi-homed scenario. And the second and third runs were the single-homed scenarios. We verified the BGP session status, route table, and VXLAN encapsulation. We sent bidirectional full-mesh inter- and intra-subnet unicast traffic to confirm no packet loss, and bridging and routing worked as expected. The difference between multi-homed and single-homed scenarios is that we verified the load-balancing function of the all-active multi-homed device, shut down one of the active links, and restored the link afterward.

After all the tests, we confirmed that the all-active function works as expected, and the switchover time is no more than three seconds for all all-active DUTs. There was also no packet loss in all VLAN-based scenarios.

During the tests with SR-MPLS underlay, we verified the IGP, BGP, EVPN sessions, and route table. Then, we sent bi-directional full-mesh traffic toward the PE devices to confirm no packet loss. During the traffic generation, we verified traffic load balancing on the multi-homed all-active DUTs. We demonstrated the Link failure and recovery on multi-homed PEs.

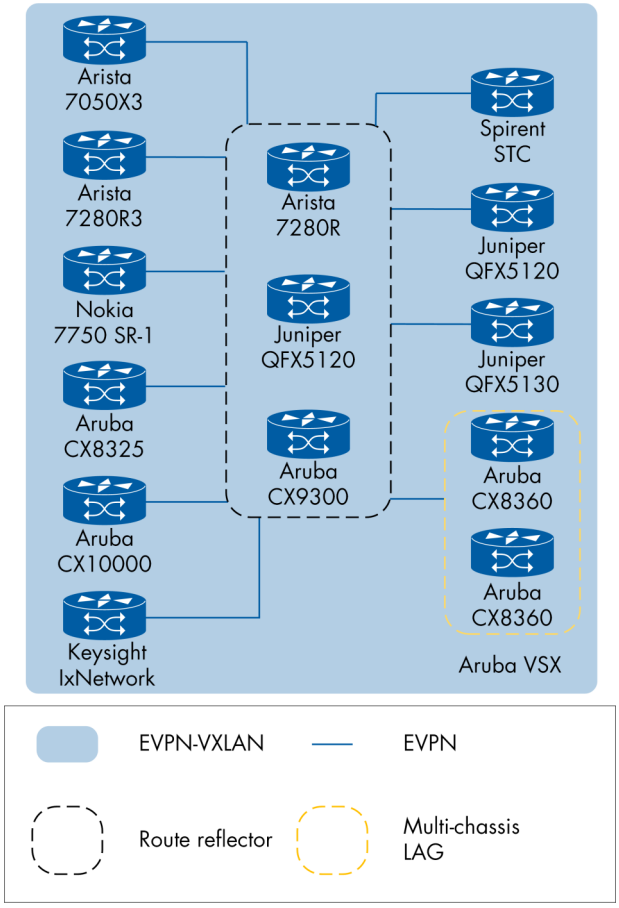

Figure 3: VLAN-based Symmetric IRB—single-homed

The following Vendors participated successfully in this test case :

Single-homed PEs: Arista 7050X3, Arista 7280R3, Aruba CX8325, Aruba CX8360, Aruba CX10000, Juniper QFX5120, Juniper QFX5130, Keysight IxNetwork, Nokia 7750 SR-1, Spirent-STC

CE: Arista 7050SX, Juniper QFX5110

Traffic generator: Keysight IxNetwork, Spirent-STC

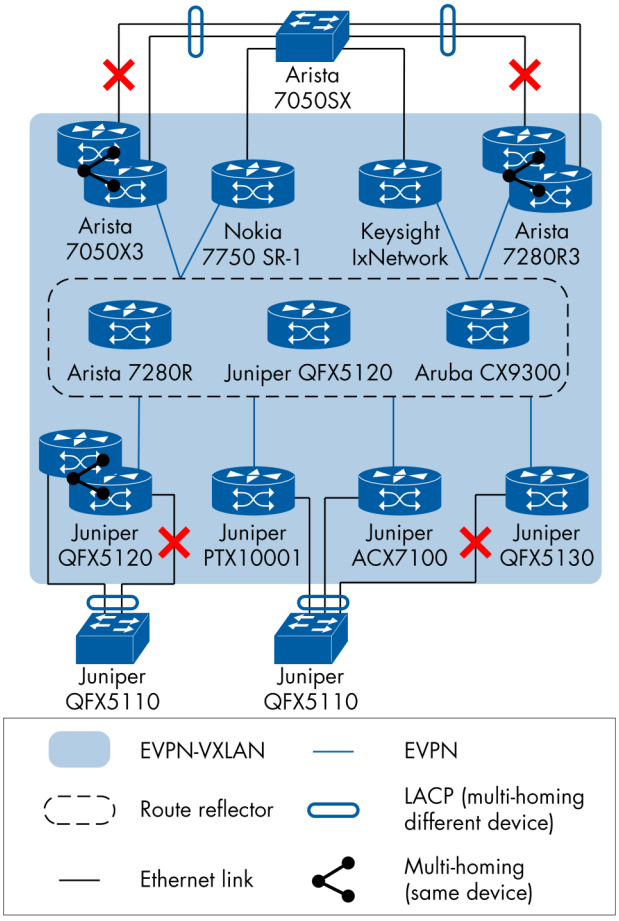

Figure 4: VLAN-based Symmetric IRB—multi-homed

The following Vendors participated successfully in this test case :

Multi-homed PEs: Arista 7050X3, Arista 7280R3, Juniper ACX7100-32C, Juniper PTX10001, Juniper QFX5120, Juniper QFX5130

CE: Arista 7050SX, Juniper QFX5110

Traffic generator: Keysight IxNetwork

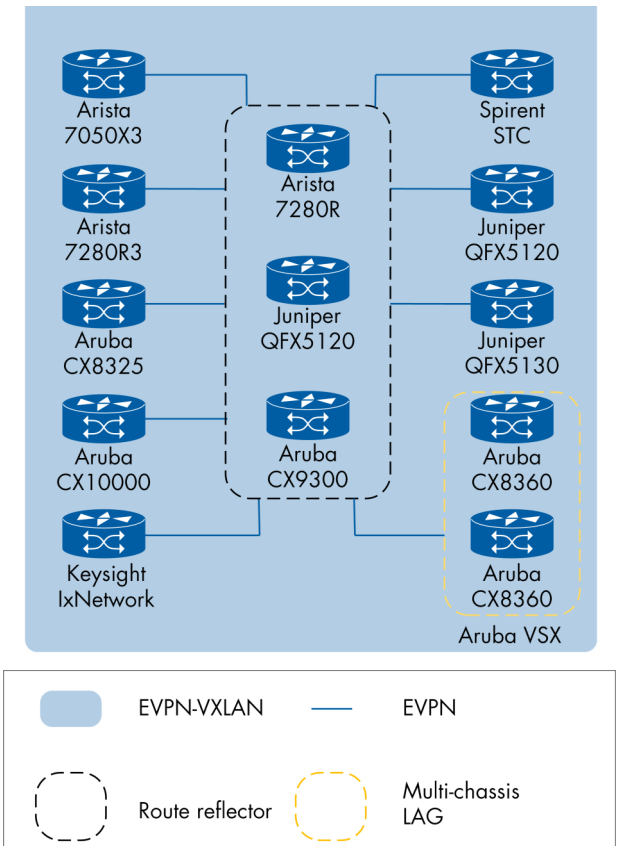

Figure 5: VLAN-aware-bundle Symmetric IRB—single-homed

The following Vendors participated successfully in this test case :

Single-homed PEs: Arista 7050X3, Arista 7280R3, Aruba CX8325, Aruba CX8360, Aruba CX10000, Juniper QFX5120, Juniper QFX5130, Keysight IxNetwork, Spirent-STC

CE: Arista 7050SX, Juniper QFX5110, Traffic generator: Keysight IxNetwork, Spirent-STC

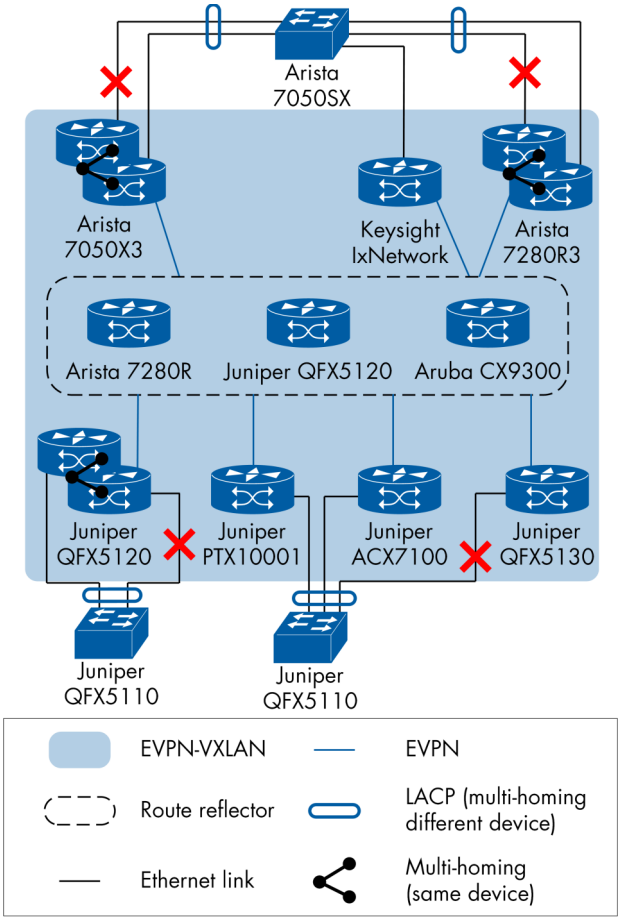

Figure 6: VLAN-aware-bundle Symmetric IRB—multi-homed

The following Vendors participated successfully in this test case :

Multi-homed PEs: Arista 7050X3, Arista 7280R3, Juniper ACX7100-32C, Juniper PTX10001, Juniper QFX5120, Juniper QFX5130

CE: Arista 7050SX, Juniper QFX5110, Traffic generator: Keysight IxNetwork, Spirent-STC

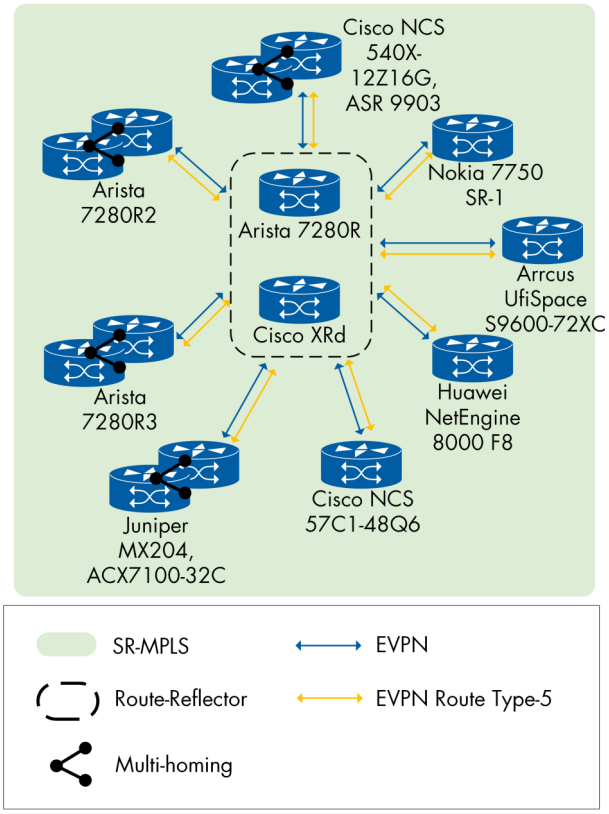

Figure 7: VLAN-Based Symmetric IRB Route Type-5

The following Vendors participated successfully in this test case :

Multi-homed PE devices: Arista 7280R2, Arista 7280R3, Cisco NCS 540X-12Z16G with Cisco ASR 9903, and Juniper MX204 with Juniper ACX7100-32C

Single-homed PE devices: Arrcus UfiSpace S9600-72XC, Cisco NCS 57C1-48Q6, Huawei NetEngine 8000 F8, and Nokia 7750 SR-1

CE: Arista 7050SX3, Traffic Generator: Spirent-STC

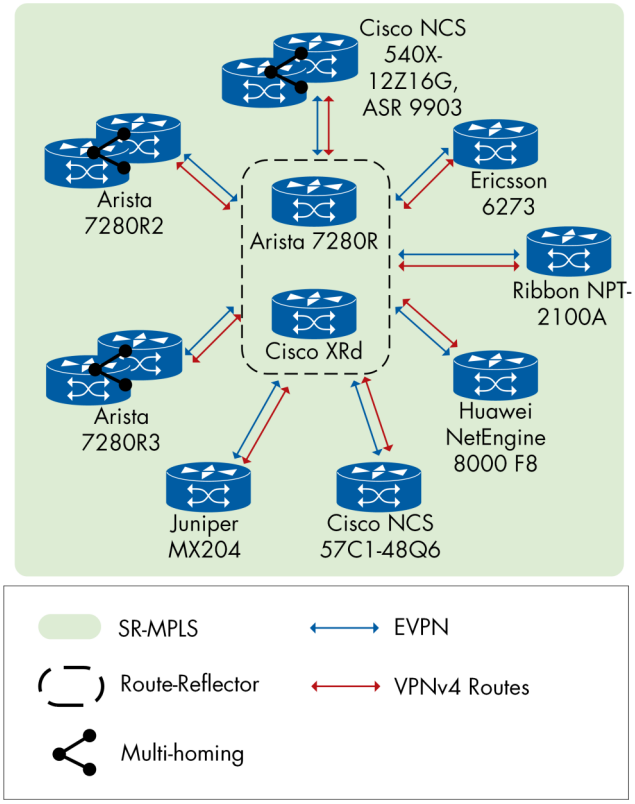

Figure 8: VLAN-Based Symmetric IRB VPNv4 Route

The following Vendors participated successfully in this test case :

Multi-homed PE devices: Arista 7280R2, Arista 7280R3, and Cisco NCS 540X-12Z16G with Cisco ASR 9903

Single-homed PE devices: Cisco NCS 57C1-48Q6, Huawei NetEngine 8000 F8, Ericsson 6273, Juniper MX204, and Ribbon NPT-2100A

CE: Arista 7050SX3, Traffic Generator: Spirent-STC

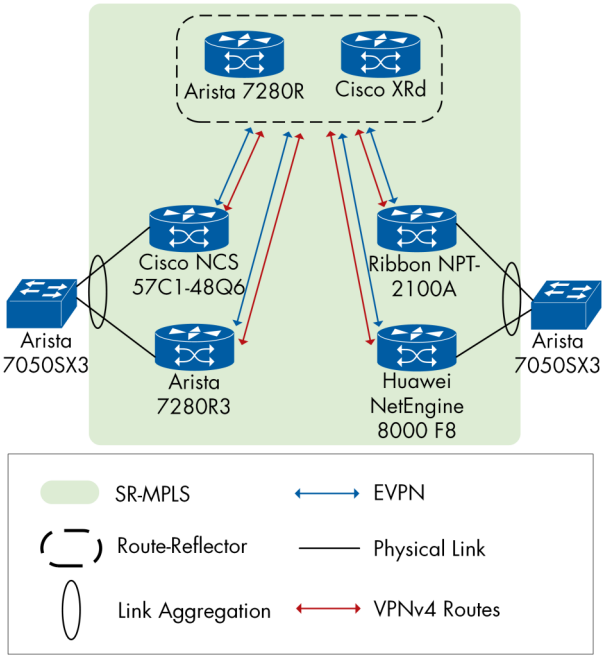

Figure 9: VLAN-Based Symmetric IRB Route VPNv4 Multi-Vendor

Multi-homed PE devices: Arista 7280R3 with Cisco NCS 57C1-48Q6 in Site A and Ribbon NPT-2100A with Huawei NetEngine 8000 F8 in Site B

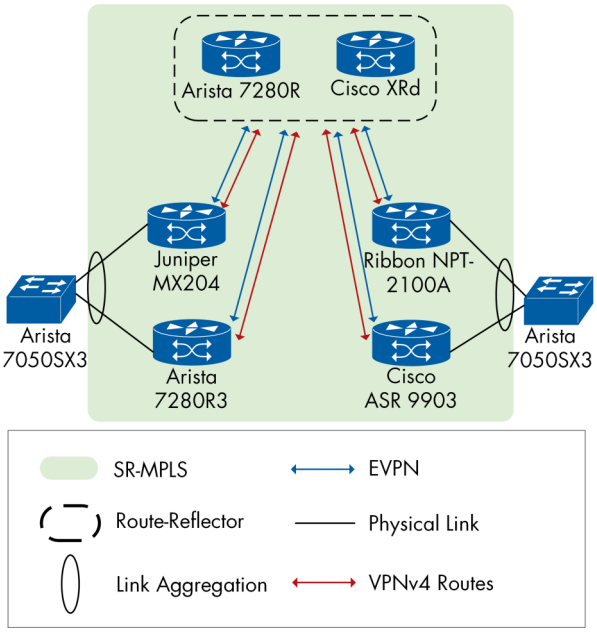

Figure 10: VLAN-Based Symmetric IRB Route VPNv4 Multi-Vendor 2

Multi-homed PE devices: Arista 7280R3 with Juniper MX204 in Site A and Ribbon NPT-2100A with Cisco ASR 9903 in Site B

CEs: Arista 7050SX3, Traffic Generator: Spirent-STC

MAC Mobility

Nowadays, data centers' standard maintenance includes but is not limited to VM creation, mitigation, deletion, etc. However, VM mitigation may move from one Ethernet segment to another, which causes the VM to be out of service. It causes high out-of-service time if everything relies on the manual prevision from network administrators. Therefore, RFC 7432 introduces a sequence-number-based BGP EVPN MAC mobility extended community to solve the issue. Once a MAC address appears in the network, the sequence number is 0. And when it moves to a new Ethernet segment, the sequence number will increase by 1 and send along with RT-2. All the RT-2 receivers will update their route accordingly and keep only the highest sequence number as the final target of the MAC address.

The test tool simulated a fixed IP and MAC addresses combination for the first DUT, then moved to all DUTs individually under the same VLAN. We verified that once the MAC address was transferred to a new DUT, the sequence number was increased by 1, and the RT-2 update was sent out. All other DUTs had their route table updated accordingly.

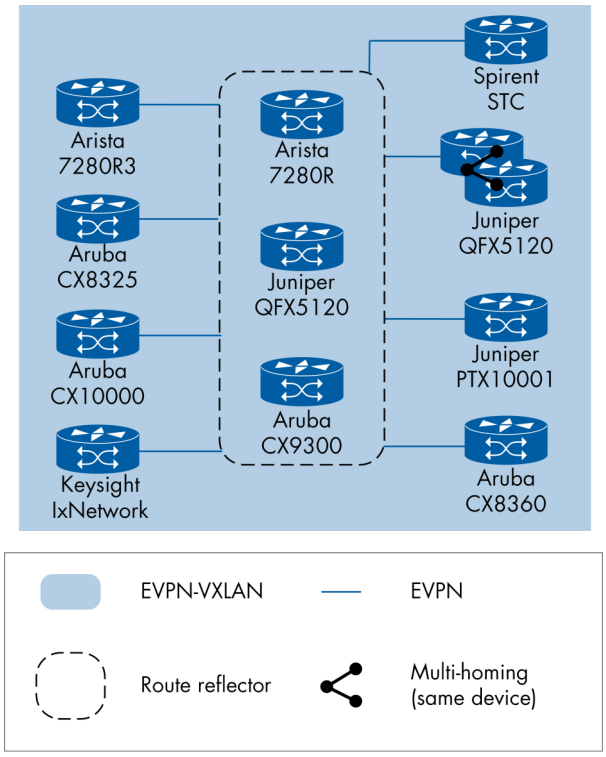

Figure 11: EVPN-VXLAN MAC Mobility

The following Vendors participated successfully in this test case :

PEs: Arista 7280R3, Aruba CX8325, Aruba CX8360, Aruba CX10000, Juniper PTX10001, Juniper QFX5120, Keysight IxNetwork, Spirent-STC

CE: Arista 7050SX, Juniper QFX5110

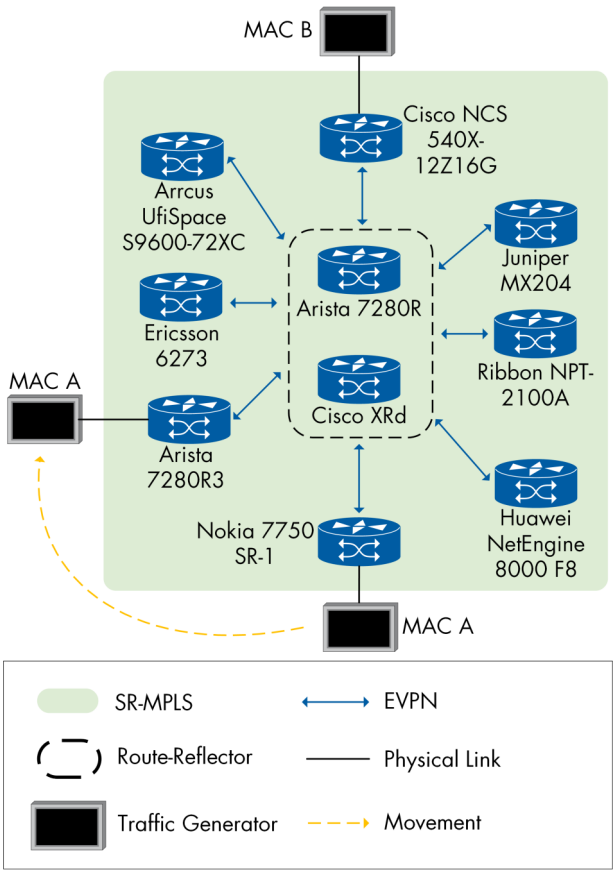

Figure 12: EVPN SR-MPLS MAC Mobility

Arista 7280R3, Arrcus UfiSpace S9600-72XC, Cisco NCS 540X-12Z16G, Huawei NetEngine 8000 F8, Juniper MX204, Ericsson 6273, Ribbon NPT-2100A, and Nokia 7750 SR-1 as PE devices. Spirent-STC participated as traffic generator.

In the SR-MPLS area, one vendor had an issue with the sequence numbering resolved by an OS upgrade.

Centralized L3 Gateway

The previous IRB test case was a distributed Layer 3 (L3) gateway deployment, meaning all the DUTs were Layer 2 (L2) and 3 gateways of the EVPN. In this test, we tested a centralized L3 gateway deployment, which implies that L2 bridgings and L3 routings are split into different DUTs, as shown in figure 13. Once the L2 VTEP receives known bridging unicast traffic, the L2 VTEP will establish VXLAN tunnel directly to the destination L2 VTEP instead of forwarding the traffic to centralized L3 gateway. The centralized L3 gateway handles all the ARP/ND and routing functions.

We have performed 3 runs of the tests. For each run, we had a single centralized L3 gateway and multiple L2 VTEPs. We sent intra- and inter-subnet traffic simultaneously and verified that L2 VTEPs forwarded intra-subnet traffic, and only inter-subnet traffic was forwarded to the centralized L3 gateway.

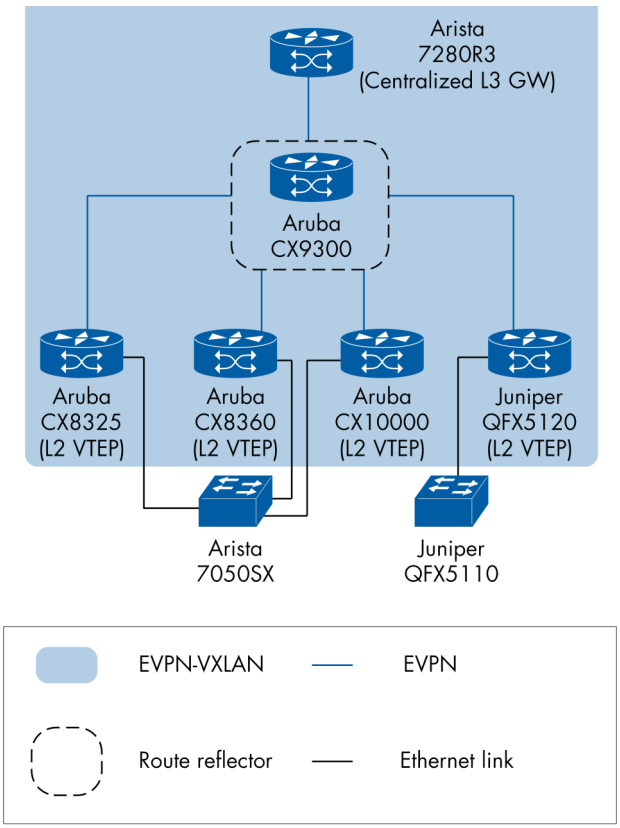

Figure 13: Centralized L3 Gateway—Run 1

The following vendors participated successfully in this test case:

Centralized L3 Gateway: Arista 7280R3, Aruba CX8325, Juniper PTX10001

L2 VTEP: Arista 7280R3, Aruba CX8325, Aruba CX8360, Aruba CX10000, Juniper QFX5120

CE: Arista 7050SX, Juniper QFX5110

Traffic generator: Keysight IxNetwork

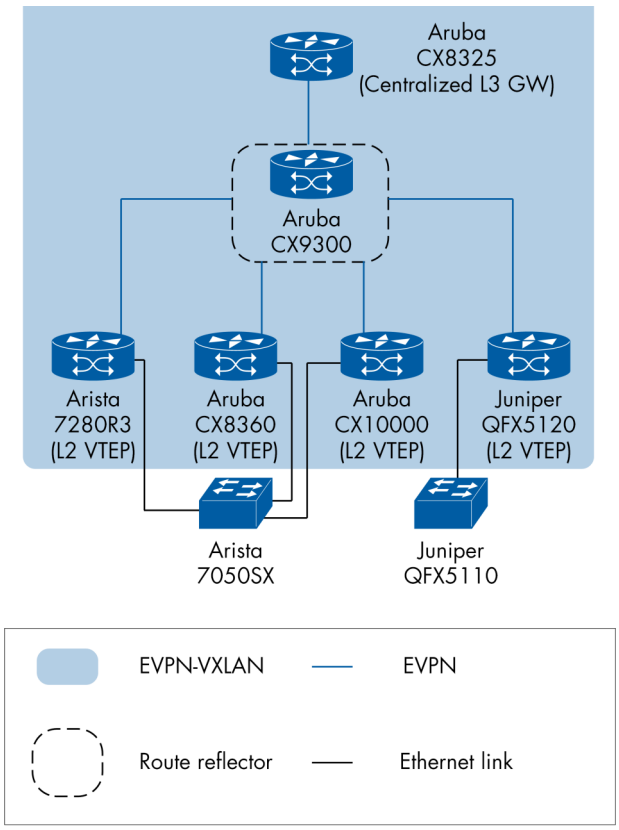

Figure 14: Centralized L3 Gateway—Run 2

The following vendors participated successfully in this test case:

Centralized L3 Gateway: Arista 7280R3, Aruba CX8325, Juniper PTX10001

L2 VTEP: Arista 7280R3, Aruba CX8325, Aruba CX8360, Aruba CX10000, Juniper QFX5120

CE: Arista 7050SX, Juniper QFX5110

Traffic generator: Keysight IxNetwork

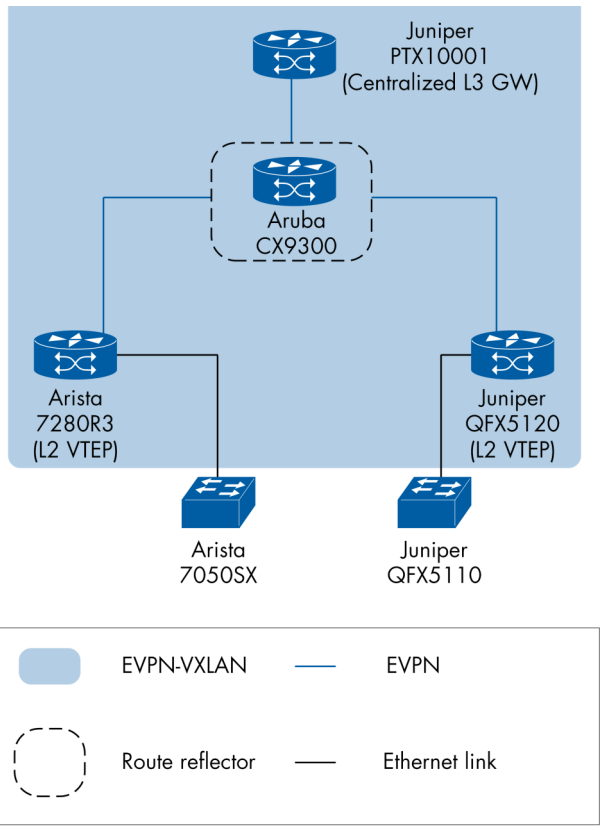

Figure 15: Centralized L3 Gateway—Run 3

During this test, one of the combinations had issues where the L3 GW responsible for providing the ARP response did not respond to the ARP request of one of the vendors, which caused the traffic not to flow. We left it for further investigation.

The following vendors participated successfully in this test case:

Centralized L3 Gateway: Arista 7280R3, Aruba CX8325, Juniper PTX10001

L2 VTEP: Arista 7280R3, Aruba CX8325, Aruba CX8360, Aruba CX10000, Juniper QFX5120

CE: Arista 7050SX, Juniper QFX5110

Traffic generator: Keysight IxNetwork

EVPN-VXLAN to EVPN VXLAN Tunnel Stitching for DCI

With the data center's scaling, the number of VXLAN tunnels will also increase dramatically. It will burden the DC gateways between the DC and WAN networks. Therefore, an demand for optimizing the VXLAN tunnel between DC and WAN networks is present. VXLAN tunnel stitching is a solution for it. VXLAN stitching stitches together specific VXLAN Virtual Network Identifiers (VNIs) to provide Layer 2 stretch between data centers on a granular basis.

We simulated 2 DC and EVPN domains. eBGP was used to build EVPN-VXLAN inside the same DC/EVPN domain. iBGP was used to create EVPN-VXLAN between 2 DC through the WAN.

VXLAN stitching was enabled on the iBGP node to optimize the VXLAN tunnel number between 2 DCs. We had three runs for the test case: VLAN-based scenario, VLAN-aware bundle scenario, and L3 gateway scenario. We verified that bridging traffic (VLAN-based and VLAN-aware bundle) and routing traffic (L3 gateway) worked well without packet loss.

The following vendors participated successfully in this test case:

VXLAN tunnel stitching gateway: Arista 7280R3, Juniper QFX5130, Nokia 7750 SR-1 (VLAN-based only)

PEs: Arista 7050X3, Juniper ACX7100-32C, Juniper PTX10001, Juniper QFX 5120, Spirent-STC

CE: Arista 7050SX, Juniper QFX5110

Traffic generator: Spirent-STC

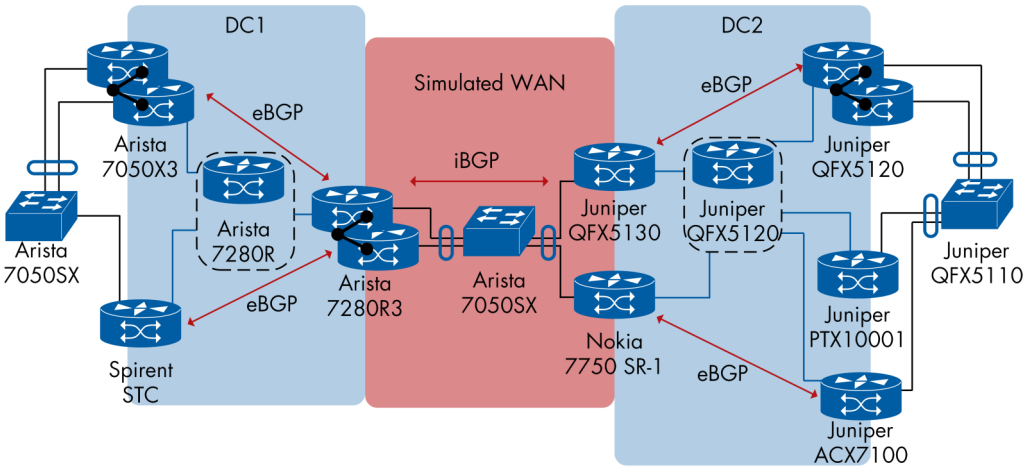

Figure 16: EVPN-VXLAN to EVPN VXLAN Tunnel Stitching for DCI—VLAN-based

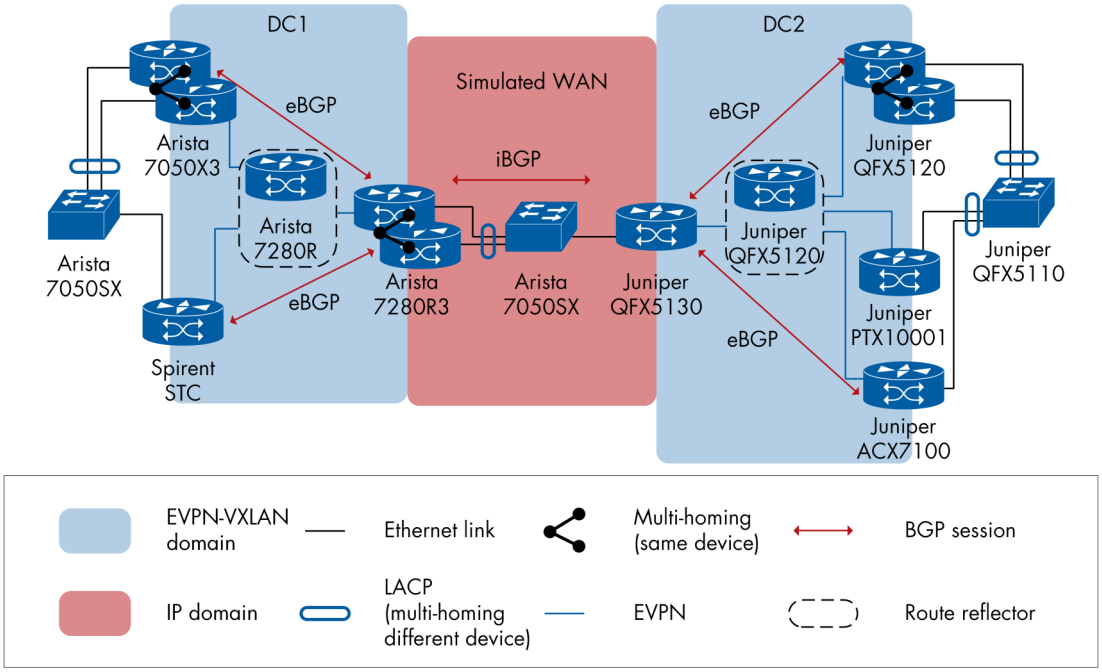

Figure 17: EVPN-VXLAN to EVPN VXLAN Tunnel Stitching for DCI—VLAN-aware-bundle + L3 GW

OISM-based Multicast Forwarding Inside DC with IR

OISM, or Optimized Inter-Subnet Multicast forwarding, is an EVPN-based technology for the integrated bridging and routing of IP multicast traffic, and it is specified in draft-ietf-bess-evpn-irb-mcast. Due to the popularity of IPTV, video surveillance, live broadcast, and application needs, multicast is a topic we cannot avoid in any network. Therefore, we have a good reason to test the EVPN-based OISM features for the DC networks.

We performed two runs of the test case in SBD and BD scenarios. Bidirectional multicast streams were set up in these two tests. The receivers were the simulated hosts that issued IGMPv3 join messages to the connected leaf routers. The leaf routers enabled the IGMP proxy. Upon receiving the IGMP join messages, the leaf routers triggered the SMET routes (RT-6) advertisement to pull the traffic for the multicast groups. The streams were successfully forwarded using ingress replication to the simulated hosts. Since each of the two flows originated from different subnets, the multicast packets were routed at the egress leaf in the SBD scenario. The multicast packets were bridged at the BD scenario's egress leaf.

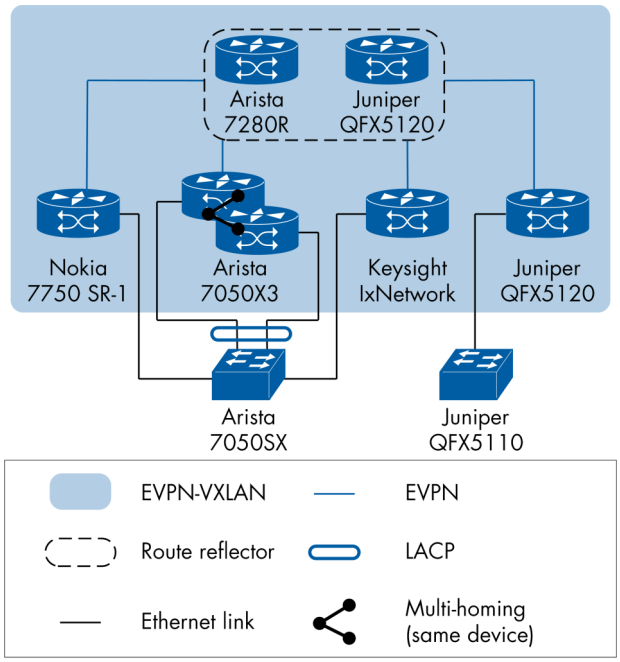

Figure 18: OISM-based Multicast Forwarding Inside DC with IR—BD

The following vendors participated successfully in this test case:

OISM-based PEs (IR): Arista 7050X3 (All-active multi-homing), Juniper QFX5120, Keysight IxNetwork, Nokia 7750 SR-1

CE: Arista 7050SX, Juniper QFX5110, Traffic generator: Keysight IxNetwork

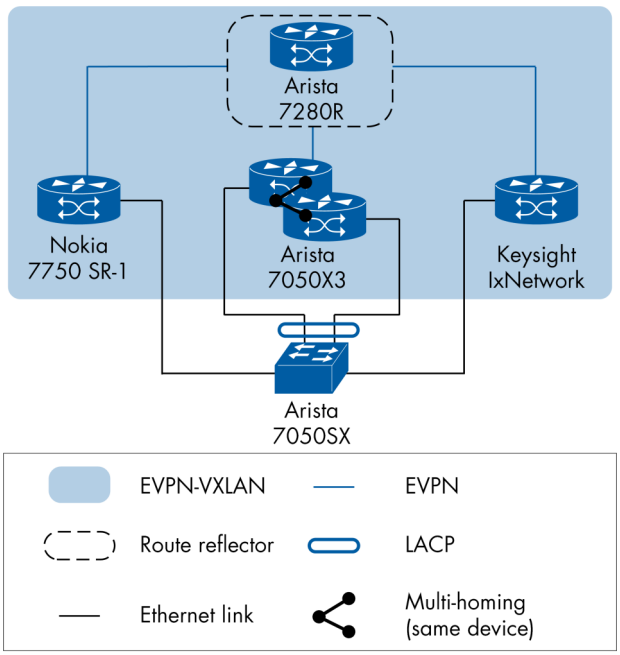

Figure 19: OISM-based Multicast Forwarding Inside DC with IR—SBD

The following vendors participated successfully in this test case:

OISM-based PEs (IR): Arista 7050X3 (All-active multi-homing), Keysight IxNetwork, Nokia 7750 SR-1

CE: Arista 7050SX, Traffic generator: Keysight IxNetwork

One vendor did not deduct the TTL even though the traffic was routed. The draft defined that the TTL should be decreased by 1 once it's been routed. Two vendors had interop issues on the SMET route. Due to one vendor's missing flag in the SMET route, another vendor recognized the remote vendor as a non-OISM-based router and sent extra copies.

OISM-based L3 Multicast IR with PEG

In the previous test, we verified that OISM-based multicast IR worked well inside DC. Therefore, we move to another essential part of the DC network: Redundancy. The DC network has a very high SLA requirement as their customer requires their data and resource to be available 99.999% of the time or even higher in different verticals. Hence, we tested the PIM-SM DR election (modular-based) on multiple PIM/EVPN Gateways (PEG).

Three PEGs were sat on the border between EVPN and PIM domains. Bidirectional multicast traffic was sent between two simulated hosts in different SBD. Modular-based PIM-SM DR election was performed, and one of the PEG was the PIM-SM DR. We verified that there was no packet loss before we shut down the link of DR.

Then we shut down the DR link to simulate the link failure scenario in the real world.

The other two PEGs performed a DR election and chose a new DR to continue forwarding the multicast traffic. And then shut down the link of the second DR and force the last PEG to be the DR.

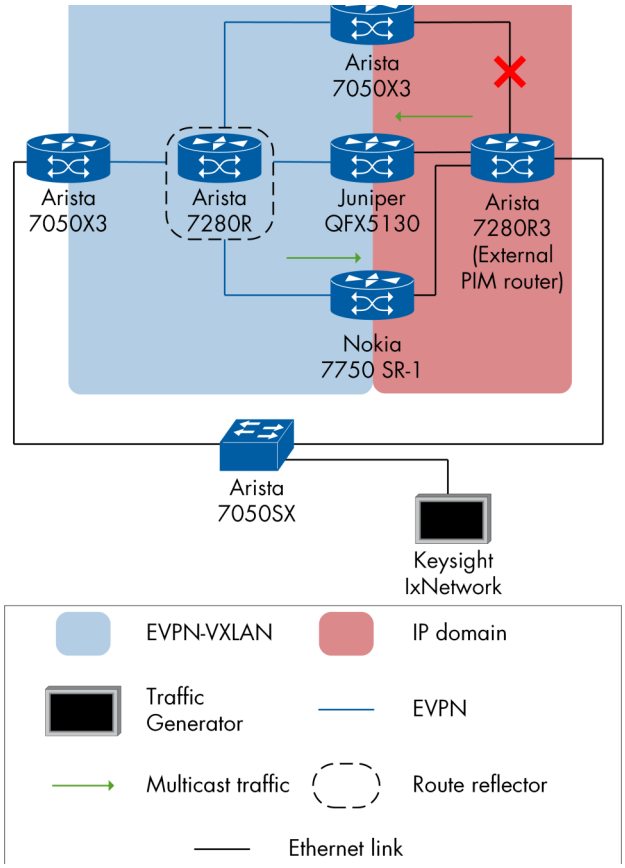

Figure 20: OISM-based L3 Multicast IR with PEG—SBD + PEG Election

The following vendors participated successfully in this test case:

PIM router: Arista 7280R3

PEGs: Arista 7050X3, Juniper QFX5130, Nokia 7750 SR-1

PE: Arista 7050X3

CE: Arista 7050SX

Traffic generator: Keysight IxNetwork

When one vendor became the last DR, it dropped 100% of one direction's multicast traffic.

E-Tree Service

The EVPN E-Tree service is a rooted-multipoint only over MPLS running core defined by MEF (Metro-Ethernet Forum). In this service, each customer site has a label as a Root or Leaf site. A Leaf AC can only send and receive traffic only from Root ACs and a Root AC can send traffic to another root or any other Leaves. To achieve ingress filtering, the ingress PE should color the ingress MAC addresses with a Root or Leaf indication before advertising them to the other PEs.

We observed firstly, the network status including the IGP sessions and BGP EVPN sessions establishment and Leaf/Root tags. Secondly, Unicast/Broadcast traffic from Roots to Roots, Roots to Leaves, and Leaves to Roots is generated without packet loss. Finally, we verified filtered Unicast/Broadcast traffic from Leaves to Leaves.

In this test, Cisco ASR 9903 and Nokia 7750 SR-1 participated as PE devices with both Root and Leaf ACs. Huawei NetEngine 8000 F8 and Juniper MX204 participated as PE devices with Leaf ACs, and Arista 7280R3 participated as a PE with Root AC. Also, we had Arista 7280R and Cisco XRd as Route Reflectors and Arista 7050SX3 as CE device. Spirent-STC participated as Traffic Generator.

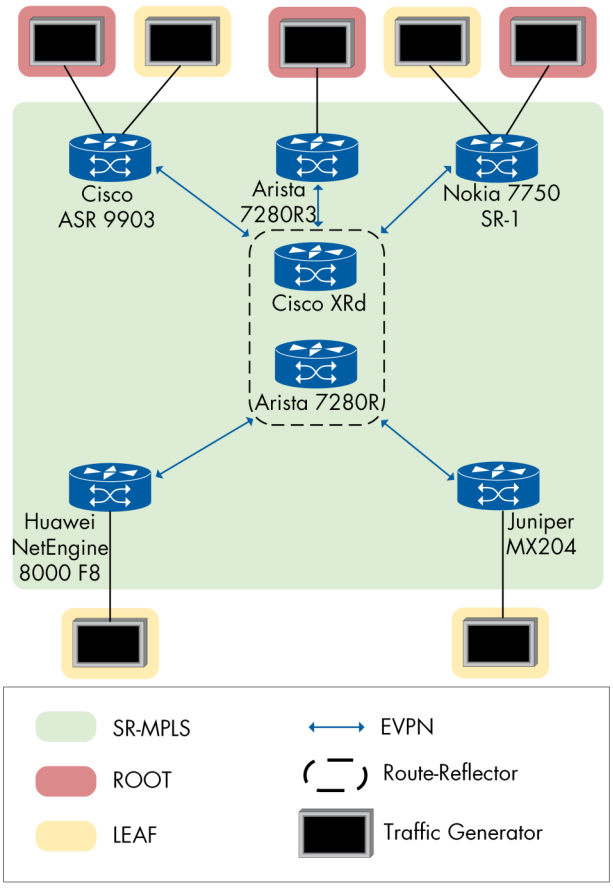

Figure 21: E-Tree Service

IPv6 BGP Unnumbered Underlay, Overlay and VTEP

The IPv4 addresses running out is not news anymore. The world starts to move to IPv6 slowly, and the DC network cannot avoid moving to IPv6. RFC 5549 offers the solution to forward IPv4 overlay traffic in an IPv6 underlay network. It also uses the IPv6 address stateless autoconfiguration to reduce the DC network deployment process within the same DC. Once the underlay is IPv6-ready, the VTEP should also move to IPv6 eventually. However, the simulated end host in the test was still IPv4, and we plan to test a dual-stack and purely IPv6 network next year to demo the path to IPv6.

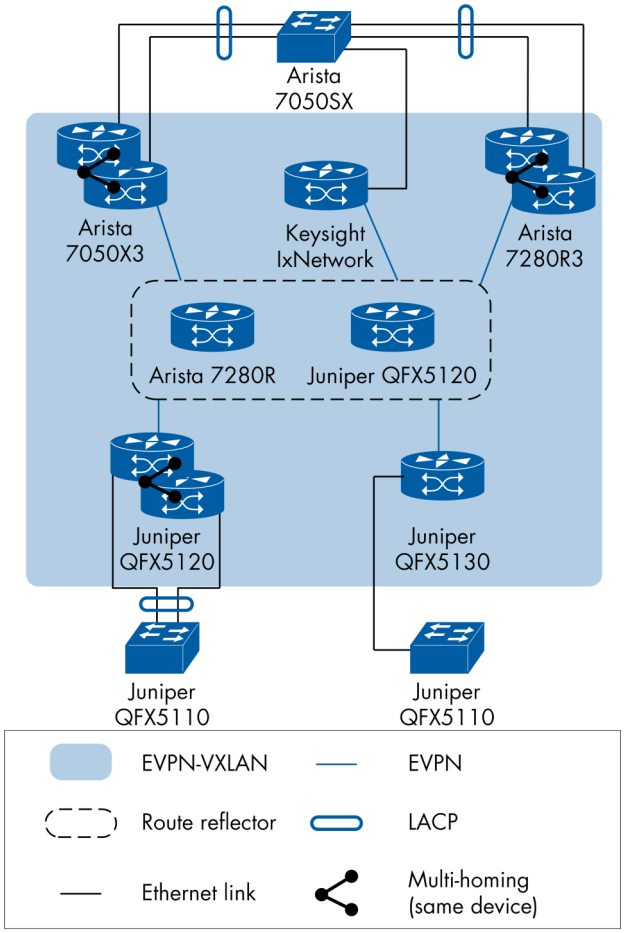

Figure 22: IPv6 BGP Unnumbered Underlay, Overlay and VTEP

In our test, we built an IPv6 underlay with BGP unnumbered feature. We verified the DUT interface IP was the IPv6 link-local address assigned by IPv6 stateless autoconfiguration.

And then, the BGP underlay next-hop address was the IPv6 link-local address of the remote end. Then we pinged the remote IPv4 host to confirm the reachability of IPv6 underlay and IPv4 overlay networks. Then we moved the overlay to IPv6 afterward. Ultimately, we sent bidirectional intra- and inter-subnet traffic from simulated IPv4 hosts through IPv6 underlay and overlay networks. We saw no packet loss during the test.

The following vendors participated successfully in this test case:

IPv6 PEs: Arista 7050X3, Arista 7280R3, Juniper QFX5120, Juniper QFX5130, Keysight IxNetwork

CE: Arista 7050SX, Juniper QFX5110

Traffic generator: Keysight IxNetwork

E-Line Service

The EVPN VPWS (E-Line) is a point-to-point service model with a BGP control plane architecture. It provides Layer 2 connectivity between two or more customer sites over the provider's MPLS/IP core network and forwards traffic without MAC address lookup. In addition, this service supports single-active or all-active multi-homing capabilities.

We created a mix of multi-homing and single-homing PEs for the E-Line service verification. The test steps included verifying IGP and MP-BGP sessions and VPWS signaling, DF election for single-active multi-homing, or traffic load balancing for all-active multi-homing ESs. We also monitored how the service behaves both when a link failure occurs and when it restores.

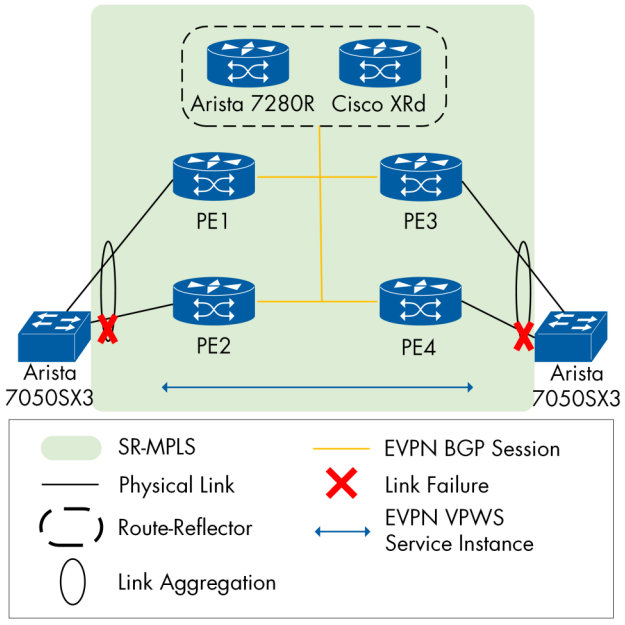

Figure 23: E-Line Service

Spirent-STC participated as Traffic Generator and Arista 7050SX3 as CE device.

One vendor could not interop with multi-homing topology in the remote site. Link failover did not work for one vendor and VPWS instance never recovered.

#	Sita A Peerings	Sita B Peerings
1	Huawei NetEngine 8000 F8 with Cisco ASR 9903	Arista 7280R3 with Cisco NCS 57C1-48Q6
2	Huawei NetEngine 8000 F8	Juniper MX204 with Ericsson 6273
3	Cisco ASR 9903 with Ciena 5166	Juniper MX204
4	Huawei NetEngine 8000 F8 with Arista 7280R2	Juniper MX204 with Cisco NCS 57C1-48Q6
5	Arista 7280R2 with Nokia 7750 SR-1	Juniper MX204 with Cisco NCS 57C1-48Q6
6	Ribbon NPT-2100A with Nokia 7750 SR-1	Juniper MX204 with Cisco NCS 57C1-48Q6
7	Cisco ASR 9903 with Ribbon NPT-2100A	Juniper ACX7100-32C with Cisco NCS 57C1-48Q6
8	Ribbon NPT-2100A with Huawei NetEngine 8000 F8	Juniper ACX7100-32C with Cisco NCS 57C1-48Q6
9	Nokia 7750 SR-1 with Juniper ACX7100	Cisco NCS 57C1-48Q6 with Arista 7280R3
10	Cisco ASR 9903 with Nokia 7750 SR-1	Cisco NCS 57C1-48Q6 with Arista 7280R3

Table 3: E-Line Service

Flexible Cross-Connect

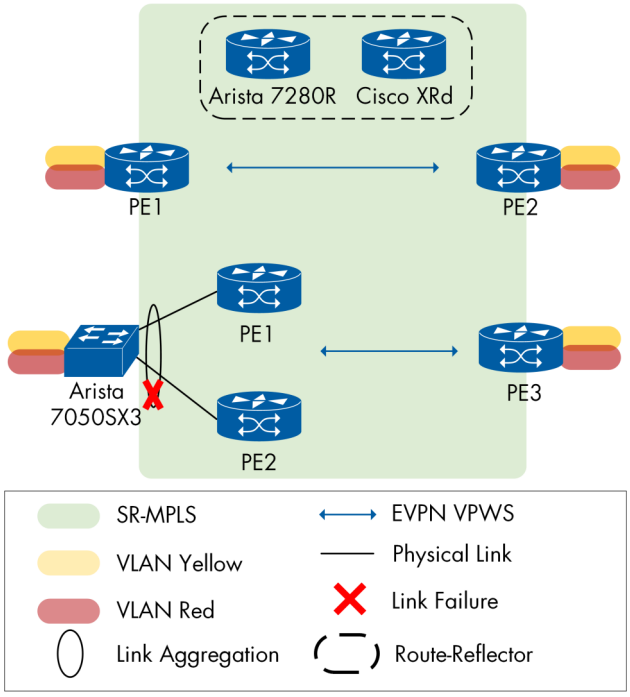

Figure 24: Flexible Cross-Connect

The Flexible Cross-Connect (FXC) service is introduced to aid service providers with a large number of ACs that require backhauling across their MPLS/IP core network. It achieves this by multiplexing multiple ACs into a single EVPN VPWS service tunnel associated with a VPWS service ID, thereby reducing the EVPN BGP signaling and associated EVPN labels to VPWS tunnels. These optimizations are particularly useful for those who use low-end access routers that may face label resource challenges.

We performed six test runs with single-homed, multi-homed, VLAN-Unaware and VLAN-Aware configurations. We verified IGP and MP-BGP sessions and VPWS signaling. Then we generated bidirectional unicast traffic toward the DUTs without any packet loss.

CE Device: Arista 7050SX3

Traffic Generator: Spirent-STC

We observed 1% packet loss while two vendors were pairing with each other.

#	Type	Sita A Peerings	Sita B Peerings
1	Single-Homed VLAN Unaware	Cisco NCS 57C1-48Q6	Arista 7280R2
2	Single-Homed VLAN Unaware	Cisco NCS 57C1-48Q6	Juniper ACX7100-32C
3	Single-Homed VLAN Unaware	Arista 7280R3	Juniper ACX7100-32C
4	Single-Homed VLAN Aware	Cisco NCS 57C1-48Q6	Juniper ACX7100-32C
5	Single-Homed VLAN Aware	Ciena 5166	Juniper ACX7100-32C
6	Multi-Homed VLAN Aware All-Active	Ciena 5166 and Juniper ACX7100	Nokia 7750 SR-1

Table 4: Flexible Cross-Connect

EVPN-VPWS Seamless Integration

The Virtual Private Wire Service (VPWS) is a Layer 2 service model that enables the establishment of point-to-point Layer 2 connections between multiple customer sites using the provider's MPLS/IP network. The EVPN-VPWS service model utilizes the benefits of EVPN and incorporates features such as all-active multi-homing, fast convergence, load balancing, and mass withdrawal functions to the legacy VPWS service. As a result, service providers are inclined towards transitioning to the EVPN-VPWS service due to its enhanced capabilities.

Upon confirmation of the IGP and BGP session statuses, unicast traffic was initiated between the hosts. Following this, we commenced the migration process from L2VPN VPWS to EVPN-VPWS. As soon as the migration was successfully executed, the BGP VPWS A-D and BGP EVPN A-D routes were advertised by the EVPN-VPWS PEs, and the EVPN-VPWS service superseded the legacy VPWS service in terms of priority.

Then we observed the PE router signaled the remote PE to bring down the legacy VPWS tunnel and use the EVPN-VPWS tunnel to forward traffic.

Arista 7280R2 and Cisco NCS 57C1-48Q6 participated as multi-homed PEs where Cisco NCS 57C1-48Q6 ran L2VPN-VPWS while Arista 7280R2 had EVPN-VPWS service.

Cisco ASR 9903, Huawei NetEngine 8000 F8, and Nokia 7750 SR-1 participated as PE devices that executed the migration from legacy VPWS to EVPN-VPWS service. Spirent-STC participated as Traffic Generator.

During the testing process, we observed one vendor could not signal the remote PE to terminate the legacy VPWS service automatically and the operator must do it manually.

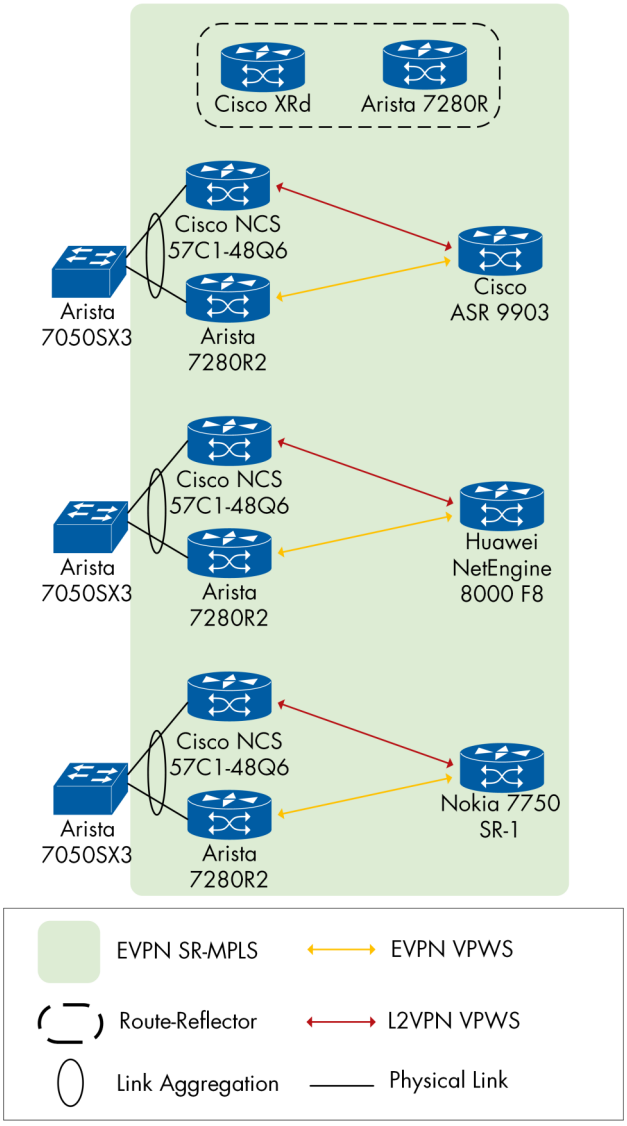

Figure 25: EVPN-VPWS Seamless Integration

EVPN-VPLS Seamless Integration

VPLS is a point-to-multipoint Layer 2 VPN service that provides Layer 2 connectivity between geographically separated data centers or customer sites across a provider's MPLS/IP backbone. VPLS is a widely deployed Layer 2 service worldwide. EVPN, on the other hand, can provide features including scalability, resiliency, control plane MAC/IP learning, all-active multi-homing, and MAC mobility. Some service providers prefer to integrate their existing VPLS network with the new EVPN running network without any changes to the existing VPLS. The seamless migration can be done on a site-by-site basis per VPN instance and must allow the coexistence of VPLS and EVPN simultaneously. A PE device may serve some customers using VPLS, while others might have been migrated to EVPN.

In two runs, we conducted this test with both LDP and BGP signaling for VPLS. After verifying both tests' IGP and BGP status, we started sending traffic toward the PE devices and proved that all the PEs were using VPLS PWs to forward traffic. Then we enabled EVPN on EVPN/VPLS PEs. As soon as the EVPN service came up, PEs advertised EVPN Inclusive multicast routes and route type-2 and discovered each other through EVPN routes. As a result, EVPN-enabled PEs shut down the PWs between each other and forwarded traffic using EVPN service; however, they kept forwarding traffic to VPLS PEs using VPLS pseudowires.

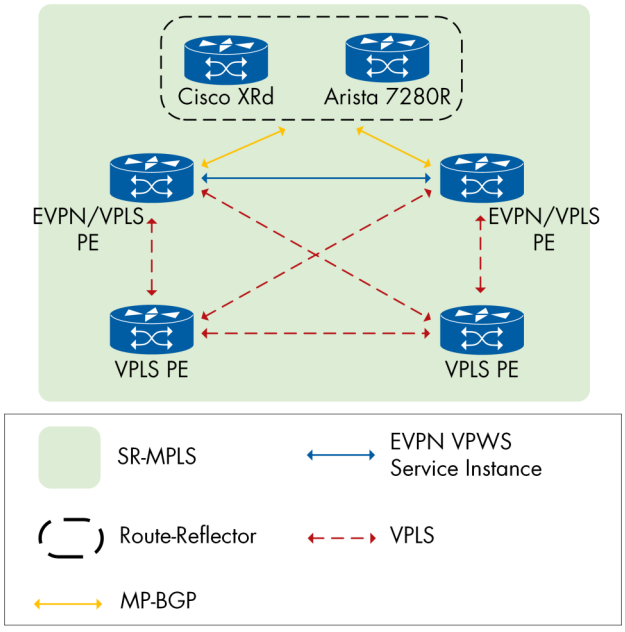

Figure 26: EVPN-VPLS Seamless Integration with LDP Signaling

In the first run, we used LDP signaling for VPLS; Cisco ASR 9903, Nokia 7750 SR-1, and Huawei NetEngine 8000 F8 participated as EVPN/VPLS PEs, while Arista 7280R3 and Juniper MX204 participated as VPLS PEs.

In the second run, where we used VPLS with BGP signaling, Cisco ASR 9903 and Nokia 7750 SR-1 participated as EVPN/VPLS PEs. Juniper MX204 participated as VPLS PE. Spirent-STC participated as Traffic Generator.

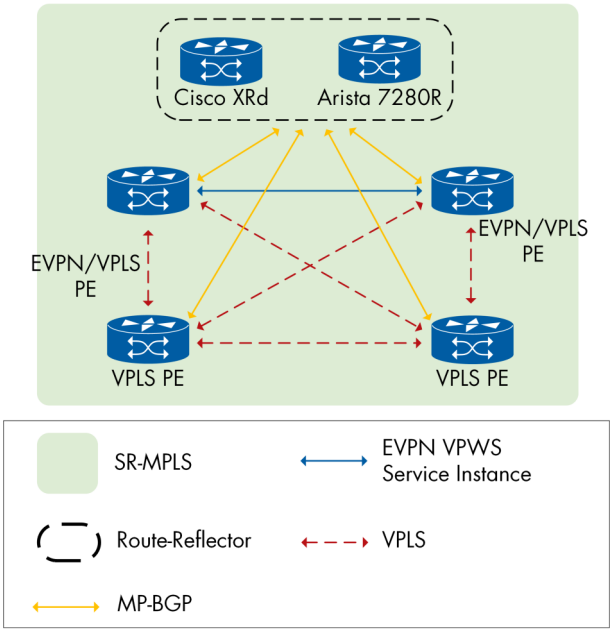

Figure 27: EVPN-VPLS Seamless Integration with BGP Signaling

EVPN PW Headend Multi-homed EVPN-VPWS Access to L3VPNs

EVPN is a control plane technology to distribute Layer 2 and Layer 3 routing information across the MPLS/IP network, while BGP is used for interconnecting the different VPNs. This allows for flexible and scalable connectivity between all types of VPNs, including Layer 2 VPNs (VPWS) and Layer 3 VPNs (L3VPN). When a tenant network stretches over EVPN and BGP VPN-IP domains, the interworking between different BGP families is crucial for inter-subnet forwarding.

Arista 7059SX3 as CE in all runs. Arista 7280R2 and 7280R3 as PE devices in the EVPN domain in all runs. Juniper ACX7100-32C as PE device in the VPN-IP domain in all runs. Spirent-STC participated as Traffic Generator.

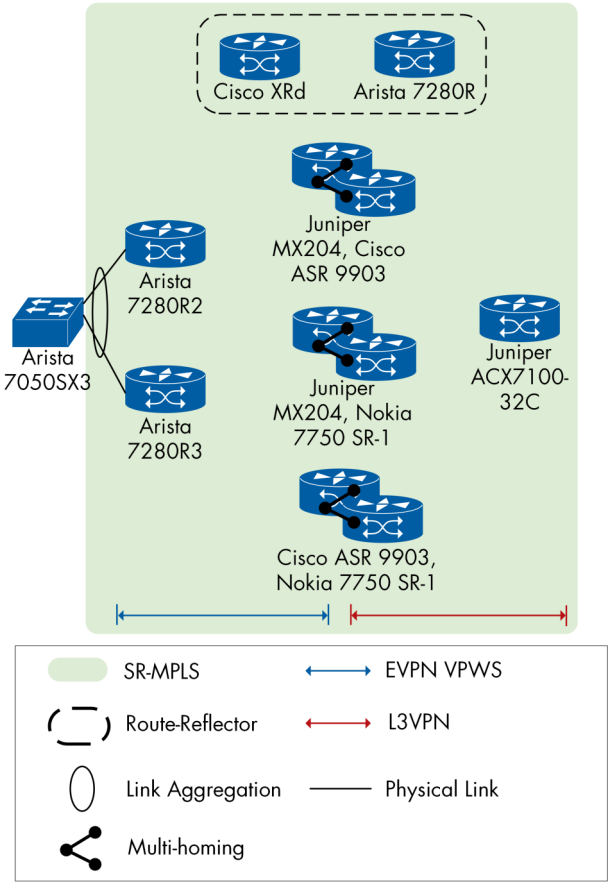

Figure 28: PW Headend Multi-homed EVPN-VPWS Access to L3VPNs

EVPN Loop Detection

EVPN loop detection is a mechanism to detect and prevent loops in an EVPN environment. The EVPN control plane uses the EVPN MAC Duplication procedure to prevent infinite MAC/IP route advertisement when a loop between two or more PE devices attached to the same Bridge-Domain is detected. The loop detection mechanism can apply a loop protection action on the duplicate MAC address to protect the data plane against endlessly looped BUM traffic. The PE puts the duplicate MAC address as a Black-Hole in its switching Table and discards the frames with the source address (and optionally with the destination address) of the Black-Holed MAC.

Arista 7280R3, Arrcus UfiSpace S9600-72XC, Cisco NCS 540X-12Z16G, Huawei NetEngine 8000 F8, Juniper MX204, and Nokia 7750 SR-1. Spirent-STC participated as the traffic generator.

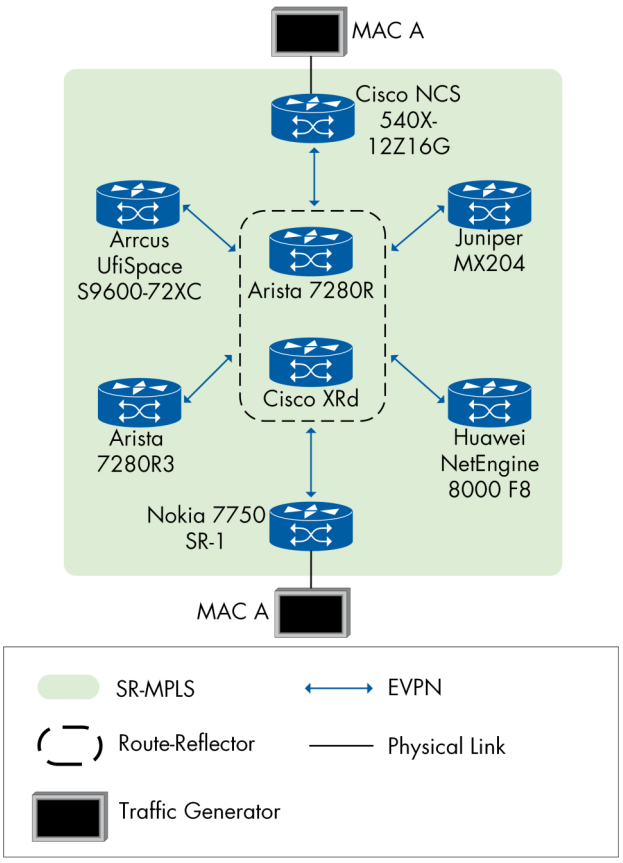

Figure 29: Loop Detection

Run	GW Peering
1	Cisco ASR 9903 multi-homed with Juniper MX204
2	Cisco ASR 9903 multi-homed with Nokia 7750 SR-1
3	Juniper MX204 multi-homed with Nokia 7750 SR-1

Table 5: PW Headend Multi-homed EVPN-VPWS Access to L3VPNs

EVPN Gateway Interworking

One scenario that can arise when integrating various technologies is advertising prefixes for different services between the access and core layers. The success of this interworking depends on the gateways' capability between different domains to receive IPVPN/EVPN prefixes from one side and subsequently advertise them to the other. This test validates the interworking between different domains.

In the first scenario we verified the interworking of EVPN SR-MPLS with EVPN SRv6 and the usage of Route Type-5.

The second run verified the interworking of EVPN SR-MPLS with SRv6 IP VPN. In the EVPN area we used Route Type-5 and in the SRv6 area we used IP VPN Routes.

In the third scenario, we verified the interworking of SR-MPLS with IP VPN Routes and EVPN VXLAN v6 (IPv6 VTEPs).

The fourth and fifth scenarios verified the interworking of SR-MPLS and EVPN VXLAN. In scenario four, we had IP VPN Routes in the SR-MPLS domain, while we changed it to EVPN Route Type-5 in the fifth scenario.

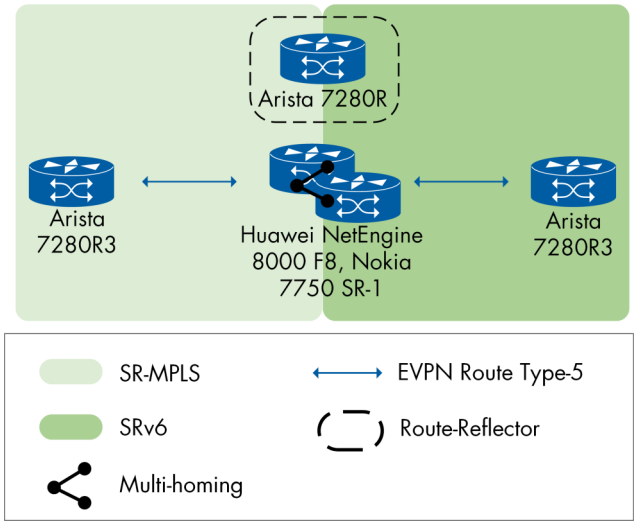

Figure 30: Gateway Interworking Scenario 1

The following vendors participated successfully in this test case:

Gateways: Huawei NetEngine 8000 F8 multi-homed with Nokia 7750 SR-1

PE devices: Arista 7280R3

Traffic Generator: Spirent

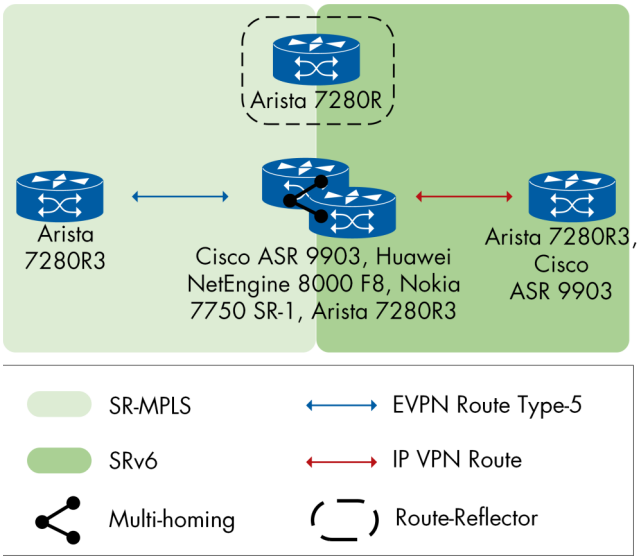

Figure 31: Gateway Interworking Scenario 2

The following vendors participated successfully in this test case:

Gateways: Huawei NetEngine 8000 F8, Nokia 7750 SR-1, Arista 7280R3, and Cisco ASR 9903

SR-MPLS PE devices: Arista 7280R3, SRv6

PE devices: Arista 7280R3, Cisco ASR 9903

Traffic Generator: Spirent-STC

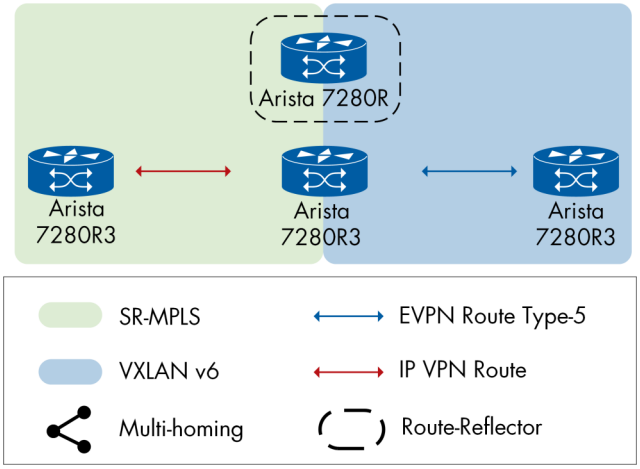

Figure 32: Gateway Interworking Scenario 3—Run 1

Gateway and PE device: Arista 7280R3

Traffic Generator: Spirent-STC

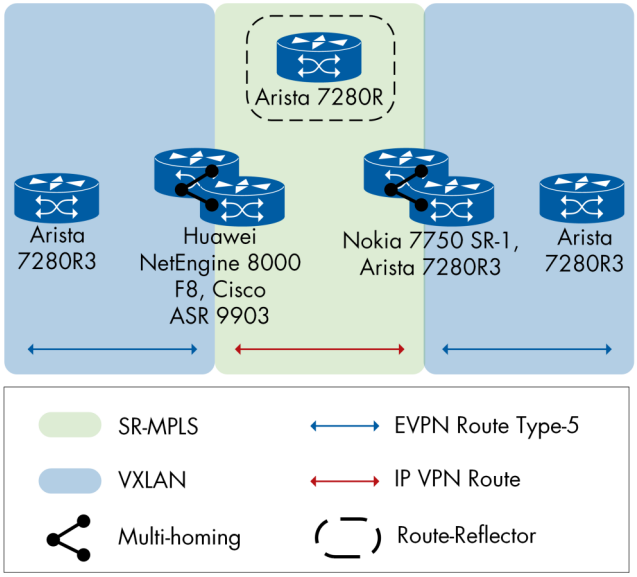

Figure 33: Gateway Interworking Scenario 4

The following vendors participated successfully in this test case:

Gateways: Huawei NetEngine 8000 F8 multi-homed with Cisco ASR 9903. Nokia 7750 SR-1 multi-homed with Arista 7280R3.

PE device: Arista 7280R3

Traffic Generator: Spirent-STC

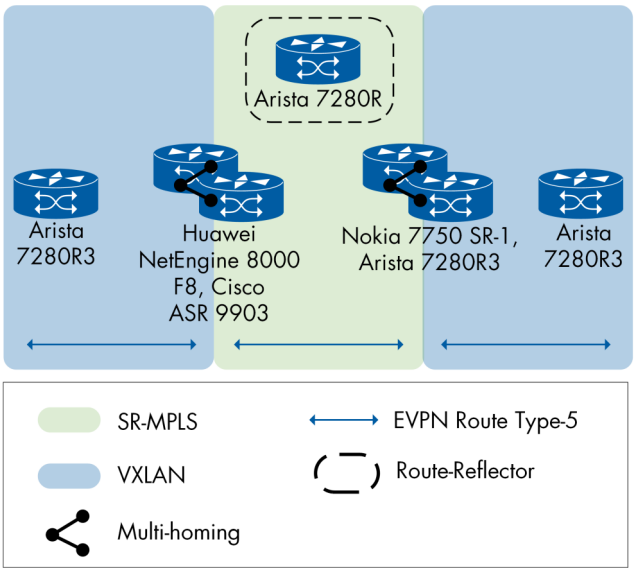

Figure 34: Gateway Interworking Scenario 5

The following vendors participated successfully in this test case:

Gateways: Huawei NetEngine 8000 F8 multi-homed with Cisco ASR9903. Nokia 7750 SR-1 multi-homed with Arista 7280R3.

PE device: Arista 7280R3

Traffic Generator: Spirent-STC

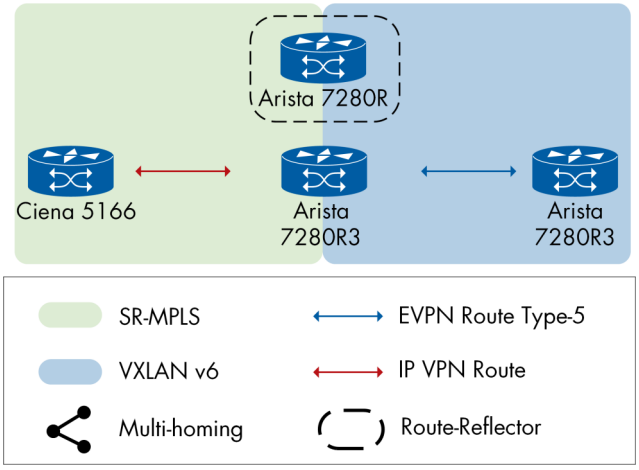

Figure 35: Gateway Interworking Scenario 3—Run 2

The following vendors participated successfully in this test case:

Gateways: Arista 7280R3

PE device: Arista 7280R3 and Ciena 5166

Traffic Generator: Spirent-STC

Segment Routing—SRv6

Segment Routing version 6 (SRv6)(RFC 8754, RFC 8986) has become a powerful option for meeting the changing needs of modern networks. It enables network administrators to set up and control network routes in a more granular and flexible manner, allowing the development of customized network services to satisfy the demands of particular applications and user groups. It also makes network operations easier by minimizing the number of protocols and control planes required to run the network.

For the first time in our annual interoperability event, we conducted tests on multi-vendor SRv6 with micro segment (µSID for short). The µSID solution is an extension to the SRv6 Network Programming model (RFC 8986)which allows the expression of SRv6 segments with a very compact and efficient representation. It is defined as the NEXT Compressed-SID flavor in IETF draft draft-ietf-spring-srv6-srh-compression.

These tests covered SRv6 BGP based Overlay services (RFC 9252), including: L3VPN, EVPN VPWS, EVPN RT5, and EVPN LAN.

Additionally, we confirmed several underlay test cases, including TI-LFA (Topology-Independent Loop-Free Alternate) Flex Algo, summarization, UPA (Unreachable prefix Announcement), and SR-TE policies while implementing µSID.

Additionally, we verified most of the previous tests using the full SID.

L3VPN over SRv6

The interoperability for IPv4/v6 BGP-based L3VPN service between vendors was the starting point of the SRv6 tests. This year, EANTC conducted the test with both Full SID and µSID.

The chosen IGP was ISIS and the physical topology was spine-leaf architecture. VRF with both VPNv4/v6 address families was configured on each PE then they were advertised via BGP to the router reflector using the locator and the SID function. We verified the control plane by checking the received routes in the routing table for vpnv4/v6 and their next hop (SID function). We used traffic generation between all the PEs to verify the data plane.

A compatibility problem was encountered by a pair in which the head-end device failed to recognize the END.DT46 message sent by the egress device. However, as per section 5 of RFC 9252, the head-end should recognize the message.

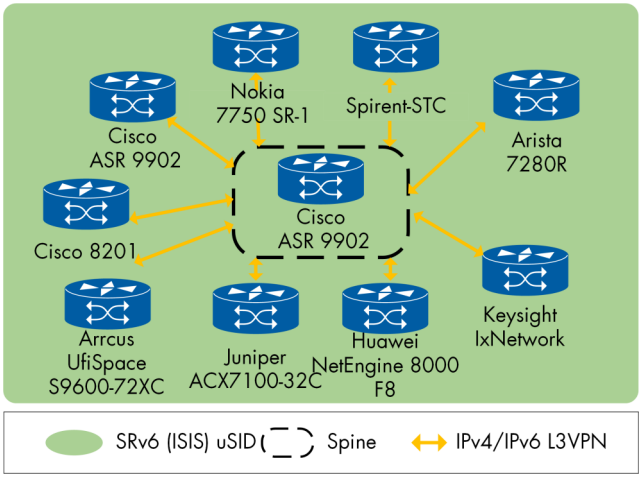

Figure 36: L3VPN over SRv6 (µSID)

These devices participated successfully as:

PE: Arista 7280R, Arrcus UfiSpace S9600, Cisco 8201, Cisco ASR 9902, Huawei NetEngine 8000 F8, Juniper ACX7100-32C, Keysight IxNetwork, Nokia 7750 SR-1, Spirent-STC

Router Reflector: Cisco ASR 9902

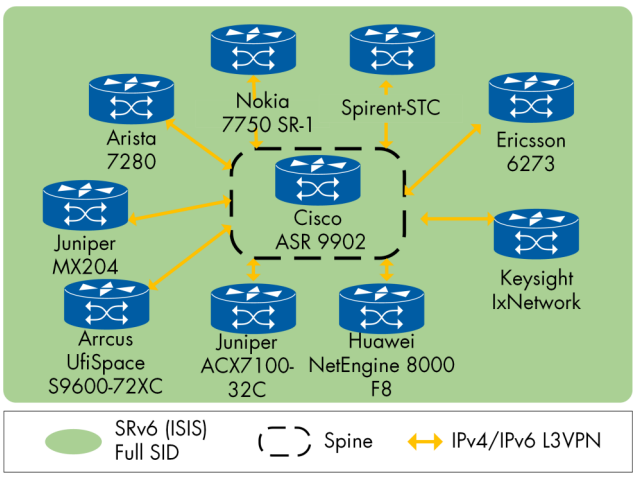

Figure 37: L3VPN over SRv6 (Full SID)

These devices participated successfully as:

PE: Arista 7280R, Arrcus UfiSpace S9600, Ericsson 6273, Huawei NetEngine 8000 F8, Juniper MX204, Juniper ACX7100-32C, Keysight IxNetwork, Nokia 7750 SR-1, Spirent-STC

Router Reflector: Cisco ASR 9902

SRv6 OAM

We conducted a ping test on the locator and endpoint, and performed a trace route between all vendors.

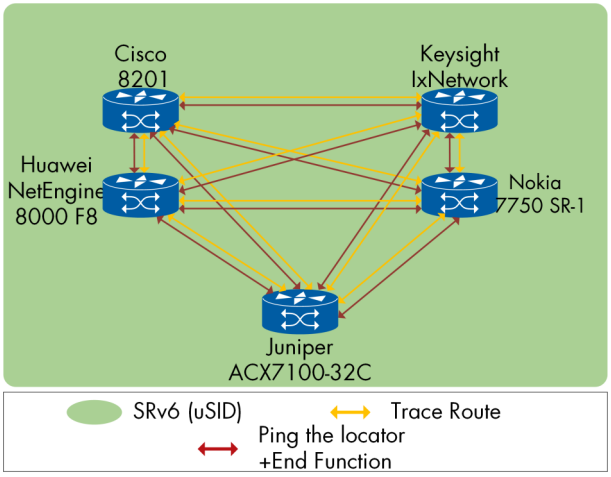

Figure 38: Ping-Trace Route µSID

These devices participated successfully in the test:

Cisco ASR 9902, Huawei NetEngine 8000 F8, Juniper ACX7100-32C, Nokia 7750 SR-1, Keysight IxNetwork

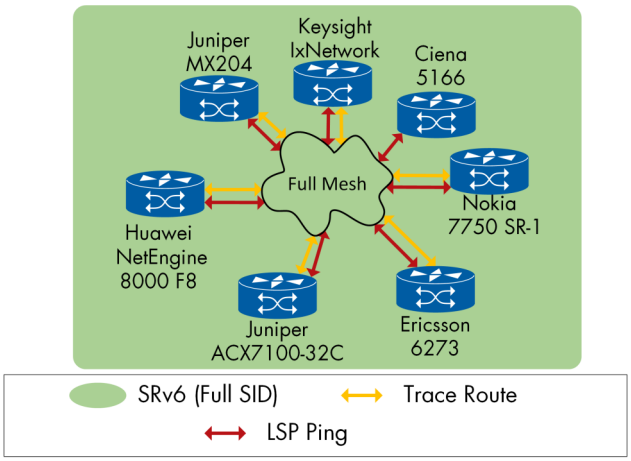

Figure 39: Ping-Trace Route Full SID

Ciena 5166, Ericsson 6273, Huawei NetEngine 8000 F8, Juniper MX204, Juniper ACX7100-32C, Nokia 7750 SR-1, Keysight IxNetwork

One vendor supported only ping over SRv6 (Full SID).

BGP IPv4/IPv6 Global Routing Table (µSID)

In order to confirm the delivery of IPv4/v6 traffic over an SRv6 core network, the process involved encapsulating packets within IPv6 packets. As per RFC 9252, the SRv6 Endpoint Behavior should be one of the following: End.DX4/6, End.DT4/6, or End.DT46.

Using SRv6 instructions uDT4 (End.DT4 with NEXT-CSID) and uDT6 (End.DT6 with NEXT-CSID) the PEs decapsulated the packet (popped the outer IPv6 header from the packet) and performed lookup in the global routing table for the IPv4 or IPv6 destination address of the inner packet.

The traffic between all participant PEs flowed with no packet loss.

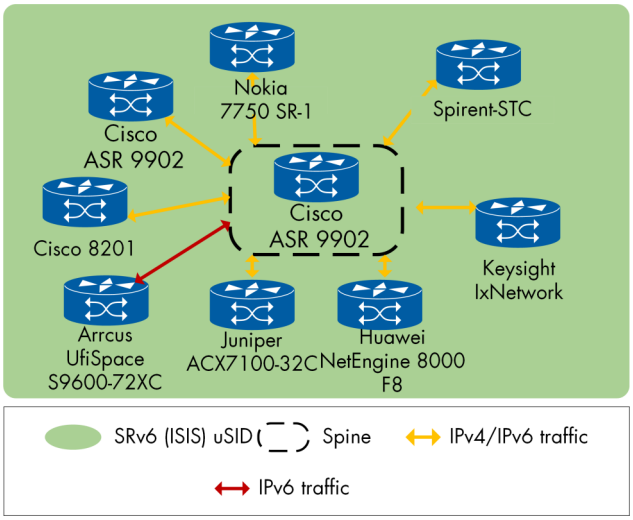

Figure 40: BGP IPv4/v6 Global Routing Table (µSID)

With a full mesh of traffic streams set up between all the listed devices, the test was successful: Arrcus UfiSpace S9600-72XC, Cisco 8201, Cisco ASR 9902, Huawei NetEngine 8000 F8, Juniper ACX7100-32C, Keysight IxNetwork, Nokia 7750 SR-1, Spirent-STC

One vendor only supported uEnd.DT6.

EVPN Services over SRv6

By leveraging the capabilities of SRv6 and EVPN, we can achieve a highly flexible and powerful solution for data center interconnects. RFC 9252 introduces an expansion to BGP that facilitates the dissemination of L2 and L3 reachability data throughout the network, while SRv6 offers an efficient forwarding plane for effective packet delivery and EVPN provides the control plane functionalities necessary for virtual network overlays.

EVPN VPWS over SRv6

Single Homing

EANTC verified VPWS service over SRv6 using both µSID and Full SID.

The Ethernet Auto-Discovery route (Route Type 1) was used to advertise point-to-point service IDs while configuring the locator to support END.DX2 (that specifies endpoint decapsulation and L2 cross-connect behavior).

For a single home scenario, each node had configured the same EVPN Instance (EVI) route and enabled BGP protocol to advertise and accept the EVPN NLRI for SRv6 services.

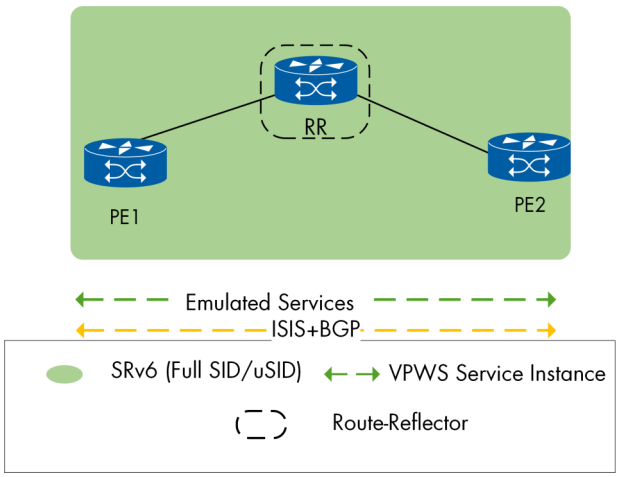

Figure 41: EVPN VPWS Over SRv6 (Full SID/µSID) Single homing

PE 1	PE 2
Juniper MX204	Nokia 7750 SR-1
Juniper MX204	Ericsson 6273
Juniper MX204	Huawei NetEngine 8000 F8
Ericsson 6273	Huawei NetEngine 8000 F8
Ericsson 6273	Keysight IxNetwork

Table 6: Full SID Pairs

PE 1	PE 2
Cisco NCS 540-28Z4C	Nokia 7750 SR-1
Cisco NCS 540-28Z4C	Huawei NetEngine 8000 F8
Nokia 7750 SR-1	Huawei NetEngine 8000 F8
Huawei NetEngine 8000 F8	Spirent-STC
Nokia 7750 SR-1	Keysight IxNetwork

Table 7: µSID Pairs

Multi Homing (All Active/Single Active)

In many network environments, network availability is crucial to the success of the business. Multi-homing in VPWS adds a layer of resilience and failover capacity.

We conducted a verification test on the CE node that had two PEs connected to it via Ethernet links, and found that all the multi-homed PEs were able to forward traffic to and from that Ethernet segment for a specified VLAN. The traffic flow was load-balanced to both PE1 and PE2, and we observed no loss of traffic.

To ensure redundancy, we had configured Link Aggregation Control Protocol (LACP) on the multi-homed CE. We then simulated a link failure on one of the links and observed that the traffic continued to flow through the second PE.

We conducted tests on both µSID and Full SID configurations. To ensure a comprehensive assessment, we also tested all active Multi-homing and Single-Active Multi-homing scenarios.

In the first setup: Juniper MX204 and Nokia 7750 SR-1 tested successfully for Multi homing all active and single active with Spirent as the remote node.

Second setup: Ericsson 6273, Huawei NetEngine 8000 F8, Juniper MX204, and Nokia 7750 SR-1 tested for Multi homing all active.

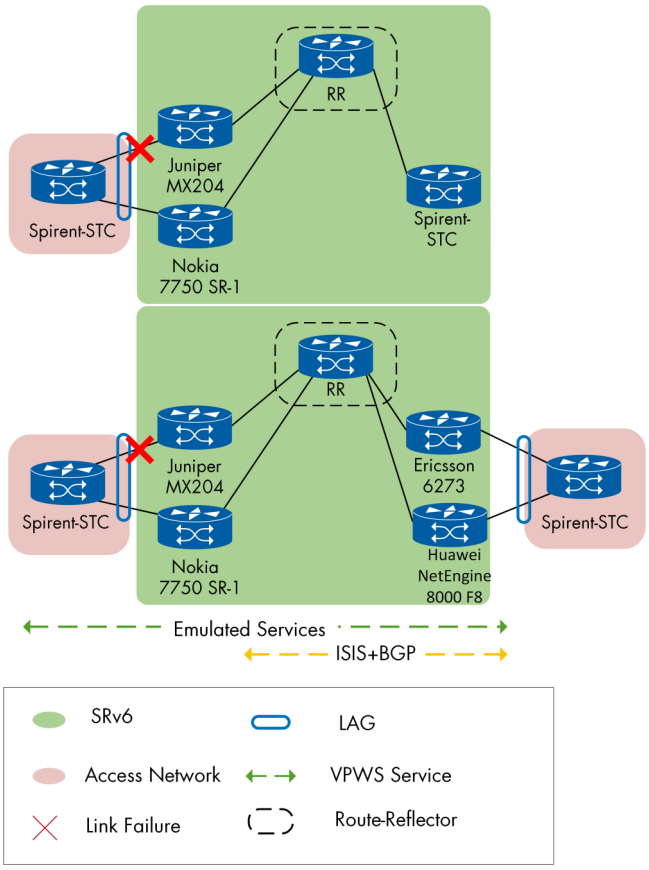

Figure 42: EVPN VPWS over SRv6 (Full SID) Multi homing

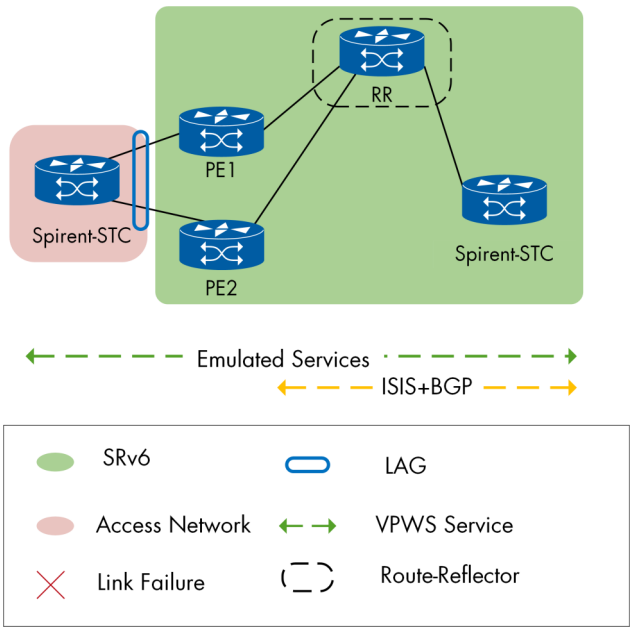

Figure 43: EVPN VPWS µSID Multi homing

PE 1	PE 2
Cisco NCS 540-28Z4C	Nokia 7750 SR-1
Cisco NCS 540-28Z4C	Huawei NetEngine 8000 F8
Huawei NetEngine 8000 F8	Nokia 7750 SR-1

Table 8: EVPN VPWS µSID Multi homing

EVPN Route Type-5 using µSID

EVPN Route Type 5 is used to advertise IP address reachability through MP-BGP to all other PEs in a given EVPN instance.

We verified that PEs can advertise VPN routes as EVPN routes to a peer in an EVPN L3VPN over the SRv6 network using µSID.

During the control plane evaluation, we confirmed that the EVPN type 5 route prefix advertised by the Provider Edge (PE) routers contains an SRv6 Layer 3 Service TLV with an SRV6 Segment Identifier (SID) sub-TLV. The IPv4/IPv6 routing table information on the PEs included a route intended for the remote Customer Edge (CE).

As we generated traffic between the PEs, no packet loss was observed.

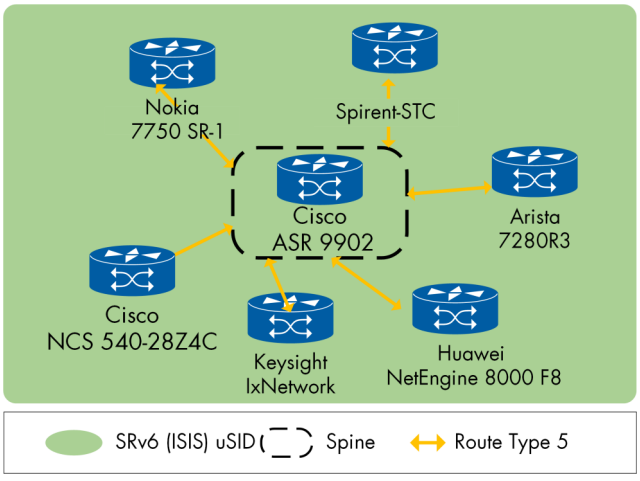

Figure 44: EVPN RT5 using µSID

The IP Prefix routes are being associated with the correct RDs on the following devices:

Arista 7280R3, Cisco NCS 540-28Z4C, Huawei NetEngine 8000 F8, Keysight IxNetwork, Nokia 7750 SR-1, Spirent-STC

EVPN E-LAN using µSID

We verified multipoint-to-multipoint Ethernet services over an SRv6-based network using µSID.

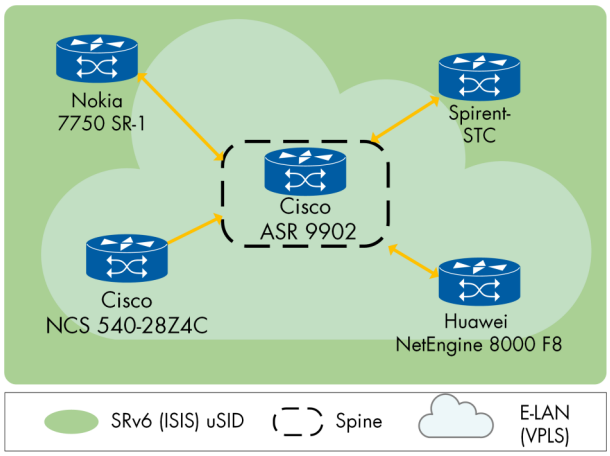

Figure 45: EVPN E-LAN over SRv6 using µSID

We verified that CE sites are able to communicate with each other over the E-LAN service through the following devices: Cisco NCS 540-28Z4C, Huawei NetEngine 8000 F8, Nokia 7750 SR-1, Spirent-STC

SRv6 Locator Summarization with µSID

Building large-scale networks is well-suited for SRv6 and its summarization features. At the boundary of each domain, IPv6 locator blocks can be summarized and distributed to neighboring domains. This end-to-end redistribution of summary prefixes allows any two nodes on the network to achieve reachability by performing a longest-prefix match on the destination address.

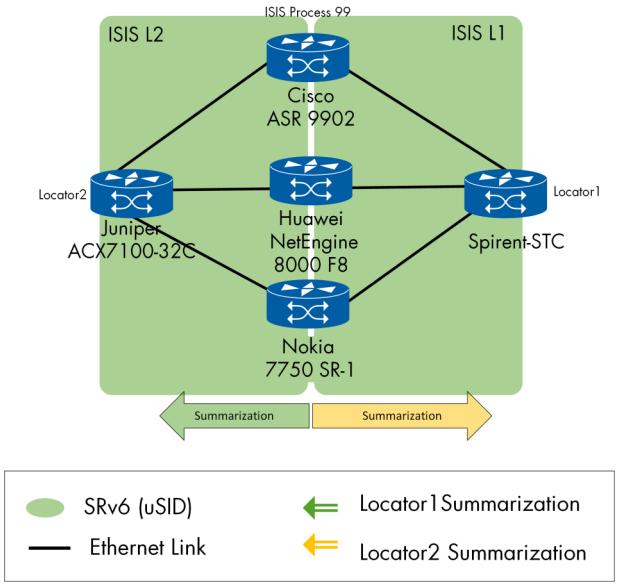

Figure 46: Locator Summarization with µSID

The setup consisted of three PE routers that worked as ABRs between two ISIS areas. The locator summary was configured in both directions using the three ABRs. Of these, two sent Locator summary advertisements through both the IP Prefix Reachability TLV and the SRv6 Locator TLV, while one sent them only through the IP Prefix Reachability TLV. We verified the received prefixes by checking the ISIS database.

This resulted (SRv6 Locator TLV is missing) for the PE in L2 not being able to resolve the service route, and traffic from L2 to L1 used the remaining ABRs.

The Summarization was accomplished by the following ABRs: Cisco ASR 9902, Huawei NetEngine 8000 F8, Nokia 7750 SR-1, as PEs Juniper ACX7100-32C and Spirent-STC

SRv6 FA Locators Summarization using µSID

Flex Algo is a powerful mechanism that offers network operators the ability to influence how the IGP calculates the least cost path for each prefix segment. By doing so, each prefix segment can traverse a unique path to reach its destination.

One of the key use cases for flexible algorithms is in the creation of multi-plane networks. These networks can be configured with multiple parallel planes that enable different types of traffic to be routed separately. By utilizing flexible algorithms and based on real-time measurements and network conditions, the network operators can ensure that each plane is able to leverage the most efficient paths for its specific type of traffic, resulting in optimal network performance.

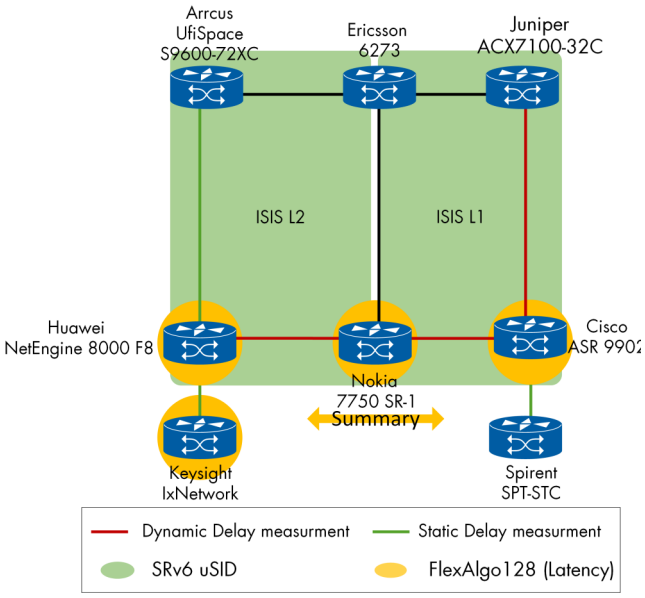

Figure 47: Summarization of Flex Algo Locators over SRv6 using µSID

We conducted the test where we evaluated the performance of the Flex algorithm using delay metrics. Some participant nodes utilized dynamic link delay measurement (TWAMP) for path calculation and incorporated the resulting latency values in ISIS (RFC 7810).

Additionally, the summarized locators in the flex algorithm were displayed with aggregated metrics.

We verified the data plane with the correct traffic passing through FA 128 plane with no issues.

Unreachable Prefix Announcement

According to "draft-ppsenak-lsr-igp-ureach-prefix-announce-02", when summarization is used, it is important to notify the network of a loss of reachability to a specific prefix that is included in the summary. This enables quick convergence away from paths that lead to the node which can no longer be reached.

In this test, we verified the process advertise such a loss of prefix reachability using the Unreachable Prefix Announcement (UPA).

The setup included an ABR that was in charge of the summary, an Ingress PE, and two egress PEs. When the ABR loses connectivity to one of the nodes in domain 2, it identified that the node's locator is included in the summary prefix and created an Unreachable Prefix Attribute (UPA) for that locator and distributed it in domain 1. After receiving the UPA via IGP, the Ingress PE switched to the backup path, and this transition took 134 ms for convergence.

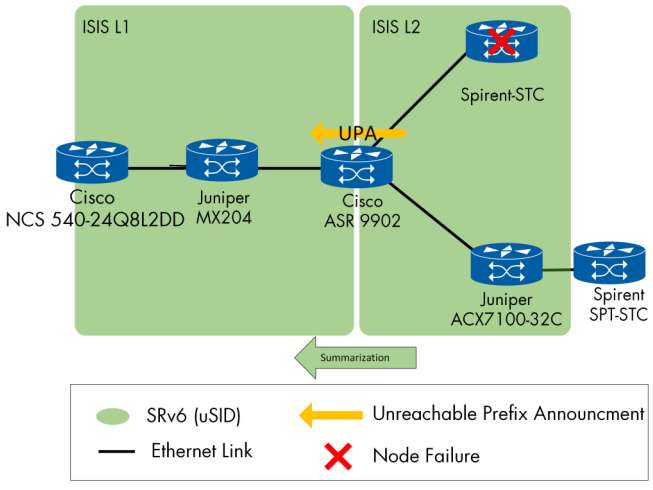

Figure 48: Unreachable Prefix Announcement

The UPA was advertised by Cisco ASR 9902 the ABR. The egress PEs were: Juniper ACX7100-32C, and Spirent-STC.

The Ingress node is Cisco NCS 540-24Q8L2D, and Juniper MX204 as P-node. Traffic generator and CEs: Spirent-STC

SRv6 TE SR Policies with Explicit Paths

SRv6 traffic engineering (SRv6 TE) utilizes the concept of source routing, where the origin calculates the route and encodes it in the packet header as a sequence of segments. This sequence of segments is added to the incoming packet via the SRv6 Segment Routing Header (SRH).

To control the flow of traffic through the network, SRv6 traffic engineering utilizes a policy that contains groups of segments.

An explicit policy in SRv6 traffic engineering is a collection of IPv6 addresses that represents an ordered list of segment IDs. The policy path is predetermined as the operator defines the segment list statically.

To create an explicit policy, vendors established a segment list(s), provided a policy name, endpoint, and color, and then linked it to a segment list from the policy. This test was completed using both µSID and Full SID.

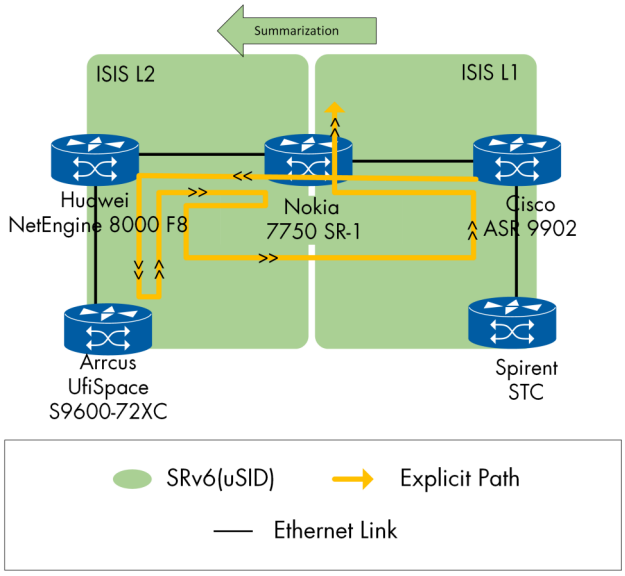

Figure 49: SRv6 TE SR Policies with Explicit Paths using µSID

The SR-TE policy was configured in the ingress node Cisco ASR 9902, the other devices acted as P nodes including:

Arrcus UfiSpace S9600-72XC, Cisco ASR 9902, Huawei NetEngine 8000 F8, Nokia 7750 SR-1, Spirent-STC

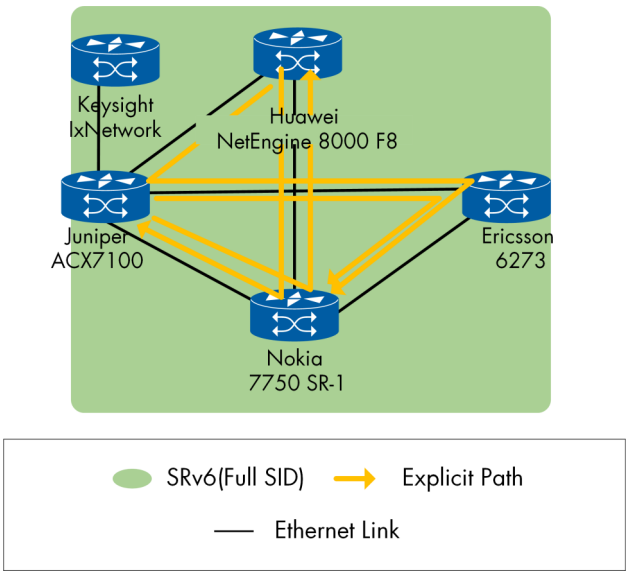

Figure 50: SRv6 TE SR Policies with Explicit Paths using Full SID

SR TE policy was configured on the following devices: Huawei NetEngine 8000 F8, Juniper ACX7100-32C,Keysight IxNetwork, Nokia 7750 SR-1, while Ericsson 6273 as a P node.

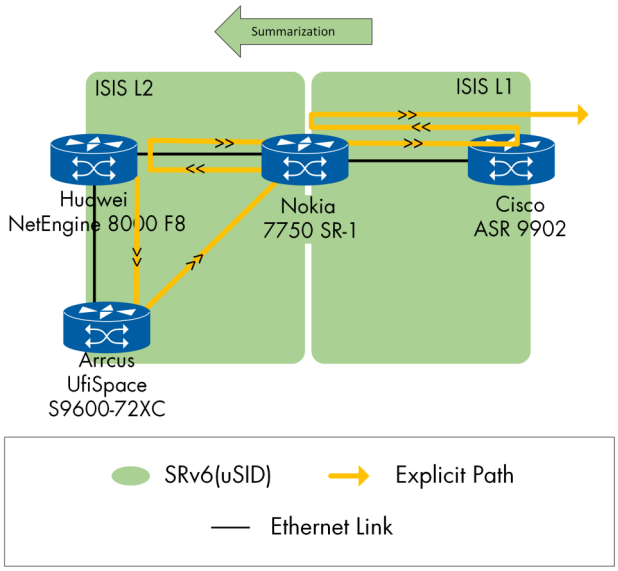

Figure 51: SRv6 TE SR Policies with Explicit Paths using µSID, Setup 2

The SR-TE policy was configured in the ingress node Huawei NetEngine 8000 F8.

The remaining devices acted as P nodes including: Arrcus UfiSpace S9600-72XC, Cisco ASR 9902, Huawei NetEngine 8000 F8, Nokia 7750 SR-1, Spirent-STC.

Topology Independent Loop Free Alternative over SRv6

For rapid protection and recovery from link or node failures, TI-LFA is used method in Segment Routing networks. It enables rapid rerouting of traffic around a failed node or connection, guaranteeing high availability and short recovery periods.

TI-LFA is especially helpful in large-scale networks where conventional protection methods are slow and can cause network-wide disruptions.

In our testing of TI-LFA, we evaluated its performance for both µSID and full SID scenarios. For µSID, we tested TI LFA with u-loop prevention and observed that the out-of-service time ranged between 2 to 33 milliseconds.

A test was carried out on the local Shared Risk Link Group (SRLG), in which two links were configured to belong to the same SRLG. Then we simulated a failure on one of the links in the SRLG and observed that the traffic correctly avoided those links. The SRLG proved effective, and the total time that the links were out of service during the test was only 3 milliseconds.

For Full SID, we conducted tests and found that the out-of-service time ranged from 5 to 4 milliseconds.

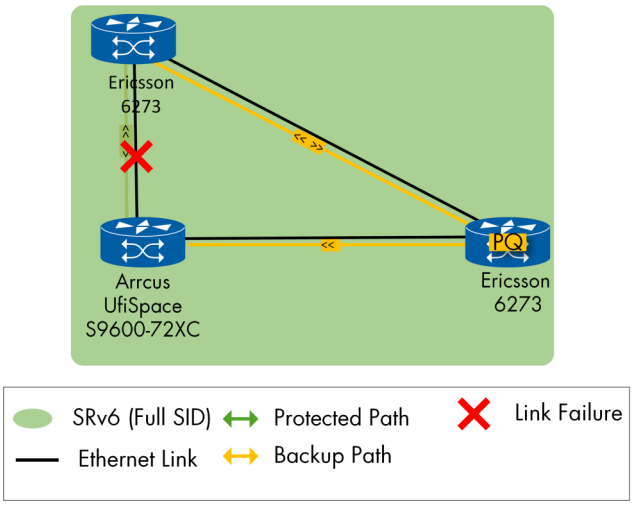

Figure 52: TI-LFA over SRv6 with Full SID

PLR nodes: Arrcus UfiSpace S9600-72XC, Ericsson 6273

PQ: Ericsson 6273

Traffic Generator: Keysight IxNetwork

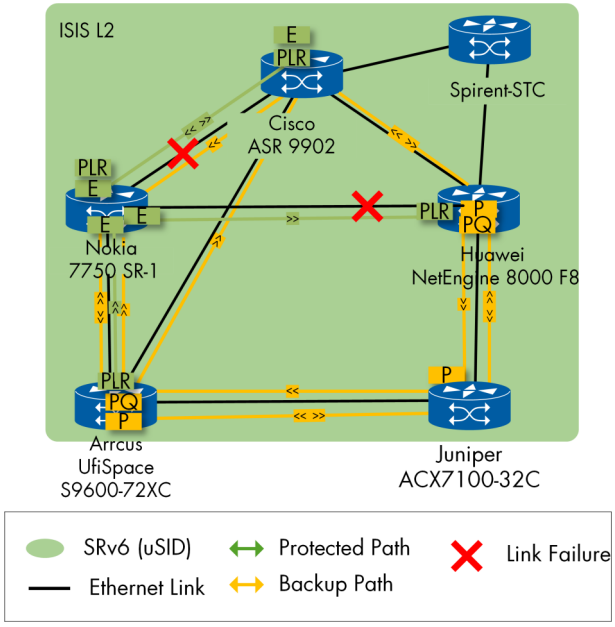

Figure 53: TI-LFA over SRv6 using µSID

PLR node: Arrcus UfiSpace S9600-72XC, Cisco ASR 9902, Huawei NetEngine 8000 F8, Nokia 7750 SR-1.

PQ nodes: Arrcus UfiSpace S9600-72XC, Huawei NetEngine 8000 F8

P nodes: Arrcus UfiSpace S9600-72XC, Juniper ACX7100-32C

Traffic Generator: Spirent-STC

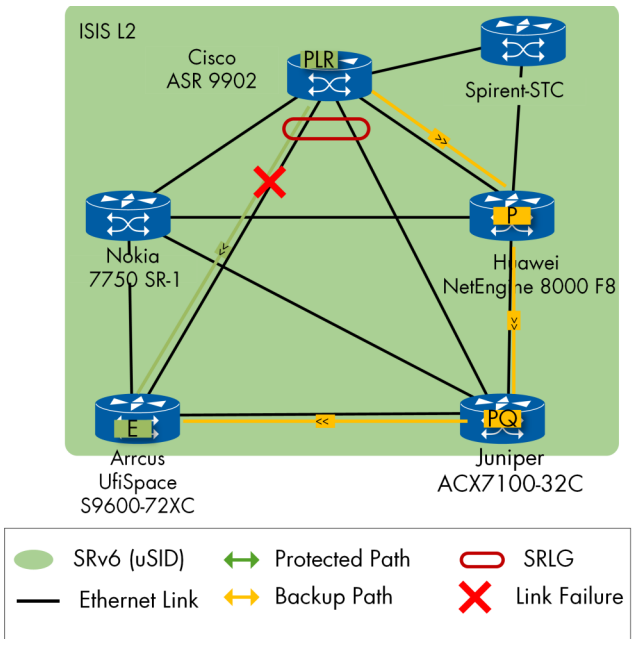

Figure 54: TI-LFA over SRv6 with Local SRLG (µSID)

SRLG was configured on Cisco ASR 9902.

PQ nodes: Juniper ACX7100-32C

P nodes: Huawei NetEngine 8000 F8

Traffic Generator: Spirent-STC

Segment Routing—SR-MPLS

Segment Routing Multi-Protocol Label Switching (SR-MPLS) has emerged de-facto industry standard to meet the requirements of modern networks as the world becomes more interconnected and reliant on high-speed data transfer. Lately, there were many efforts to establish an end-to-end intent-aware path across multi-domains of service provider environments. In EANTC, we tested a BGP-based routing solution dedicated to this goal, called the BGP Classful Transport Planes. As SR-MPLS is particularly popular in Inter-AS (Autonomous System) networks, which connect numerous network domains owned by different enterprises, we tested several mechanisms, such as BGP LS (Link State), Flexible Algorithm Prefix Metric (FAPM), and chaining ASs options. This year, we placed significant emphasis on OSPF segment routing, which included the implementation of Flex Algo and FAPM (which was done for the first time in our interop event), as well as covering mechanisms for fast re-routing, performance measurement, failure discovery, and SR traffic steering.

L3VPN Services

As a preliminary test for interoperability in SR-MPLS, we utilized L3VPN services. The participating nodes were interconnected in a spine-leaf topology and established ISIS/OSPF sessions with one another. The routing tables contained the loopback addresses and corresponding SIDs of the involved PEs. After verifying that all VPNv4 and VPNv6 services were operational on the vendor devices, we proceeded to generate IPv4 and IPv6 traffic between each pair of PEs, which did not result in any packet loss.

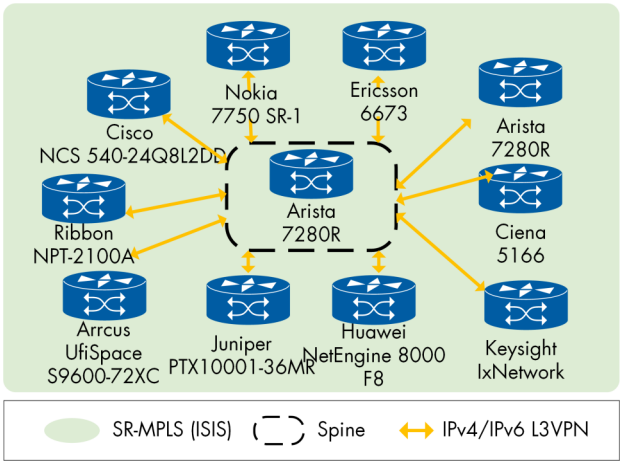

Figure 55: L3VPN over SR-MPLS (ISIS)

Arrcus UfiSpace S9600-72XC, Arista 7280R, Cisco NCS 540-24Q8L2DD, Ericsson 6673, Ciena 5166, Huawei NetEngine 8000 F8, Juniper PTX10001-36MR, Keysight IxNetwork, Nokia 7750 SR-1, Ribbon NPT-2100A

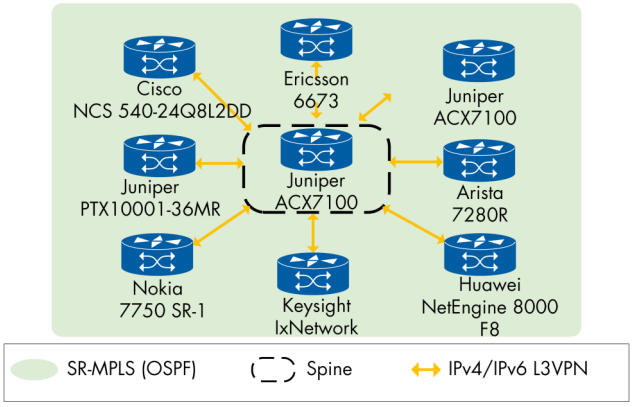

Figure 56: L3VPN over SR-MPLS (OSPF)

The following devices passed the test as PEs: Arista 7280R, Cisco NCS 540-24Q8L2DD, Ericsson 6673, Huawei NetEngine 8000 F8, Juniper ACX-7100, Juniper PTX10001-36MR, Nokia 7750 SR-1

Traffic Generator for both: Keysight IxNetwork

SR-MPLS OAM

The ability to quickly identify and troubleshoot network failures is essential for network operators. To help with this task, RFC 8287 defines a set of tools for detecting and diagnosing network issues, including Label Switched Path (LSP) Ping/Traceroute for Segment Routing IGP-Prefix Segment Identifiers (SIDs) with MPLS Data Plane. These tools are widely used in networks to test connectivity, measure latency, and identify the location of network faults. In this context, we conducted a verification of the troubleshooting and failure detection tool. All participants successfully performed the test except for one vendor that supports only ping over SR-MPLS.

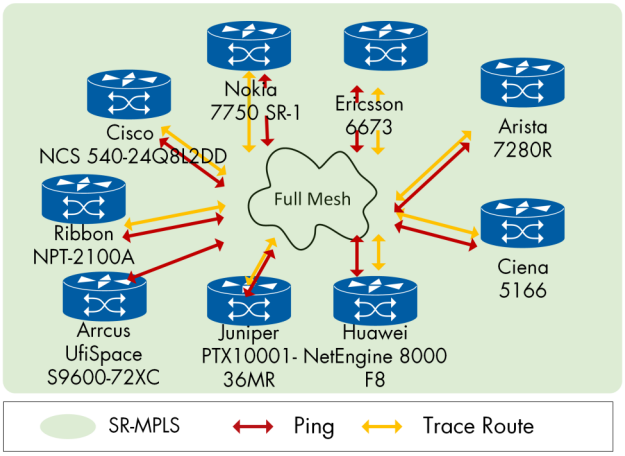

Figure 57: Ping/Trace Route over SR-MPLS

The following devices passed the test as PEs: Arrcus UfiSpace S9600-72XC, Arista 7280R, Cisco NCS 540-24Q8L2DD, Ericsson 6673, Ciena 5166, Huawei NetEngine 8000 F8, Juniper PTX10001-36MR, Keysight IxNetwork, Nokia 7750 SR-1, Ribbon NPT-2100A

Flexible Algorithm

Flexible Algorithm over ISIS

IGP protocols historically compute the best paths over the network based on the IGP metric assigned to the links. IGP Flexible Algorithm (RFC 9350) enhances IGPs to compute the best paths based on a given combination of calculation-type, metric-type, and constraints. With Flexible Algorithm an operator can associate one or more SR-MPLS Prefix-SIDs or SRv6 locators with a particular Flex-Algorithm. Each such Prefix-SID or SRv6 locator then represents a path that is computed according to the identified Flex-Algorithm.

In our test, we confirmed the generation of multiple network planes utilizing the flex algorithm based on ISIS. We employed three different flex algorithms, FA 128 was based on the minimum delay metric, FA 129 was based on IGP metric and exclusion of interfaces with a given link administrative group (green affinity) and FA 130 was relying on the TE metric. All participants had TE attributes advertised in Flex-Algo specific Application-Specific Link Attribute (ASLA) sub-TLVs. One vendor could not generate SID per Flex-Algo with a single loopback IP, so they did not participate. Two vendors had to utilize a knob to prevent fallback to the native algorithm 0 LSP. With that fallback Flex-Algo 129 with "exclude green" worked. Since Ribbon was one of the vendors that only supported Flex Algo Legacy and not ASLA, we conducted a test with the Legacy Flag enabled. This test included two different constraints: the use of a manually delayed metric with algorithm 128 and a TE metric with algorithm 130.

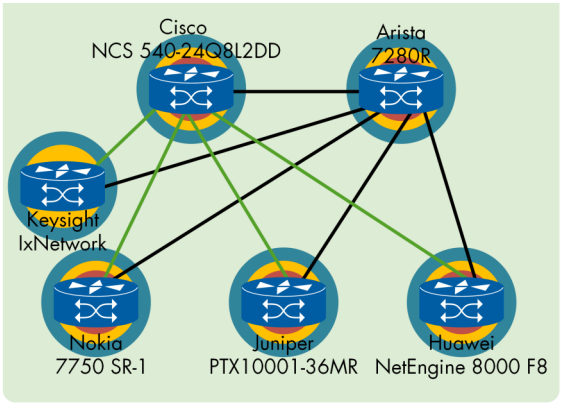

Figure 58a: Flex Algo ASLA (ISIS)

The participated nodes for ASLA: Arista 7280R, Cisco NCS 540-24Q8L2DD, Huawei NetEngine 8000 F8, Juniper PTX10001-36MR, Nokia 7750 SR-1, Keysight IxNetwork

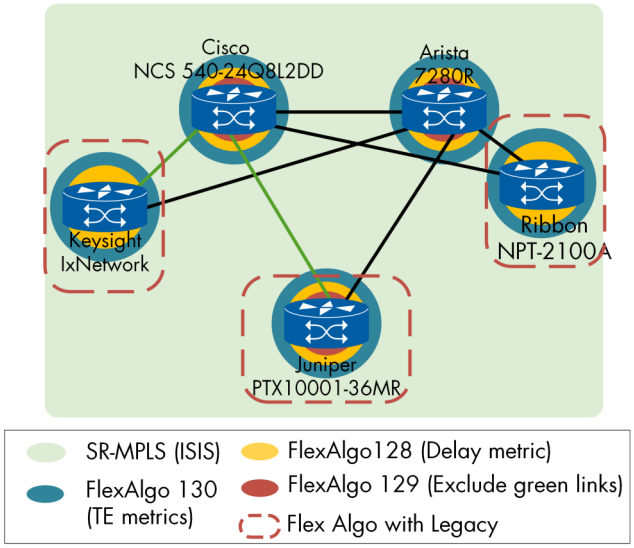

Figure 58b: Flex Algo with Legacy (ISIS)

For Legacy test: Juniper PTX10001-36MR, Keysight IxNetwork, Ribbon NPT-2100

Flexible Algorithms over OSPF

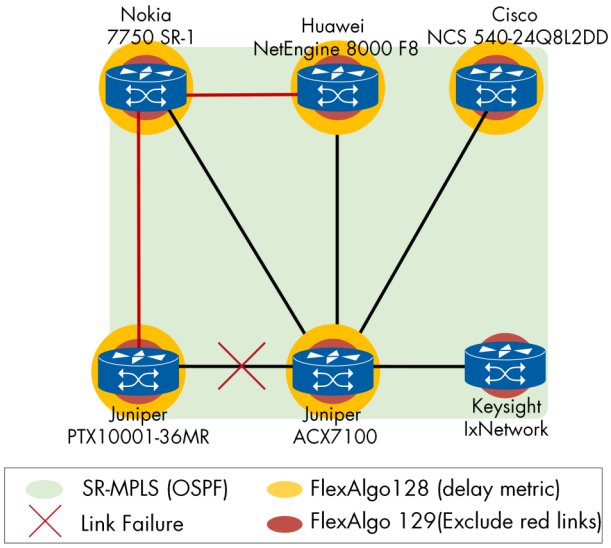

Figure 59: Flexible Algorithm (OSPF)

We conducted a test to implement Flex Algo using OSPF for the first time. Our objective was to enable multi-planes and isolation for the existing OSPF underlay network, in order to fully utilize its potential. To achieve this, we defined two Flex Algo configurations, which the OSPF underlay network calculated by identifying the nodes and links for each Flex Algo. Using the performance metrics, two different Flex Algo planes were formed. We then created L3VPNs within each of the Flex Algo configurations and introduced traffic in different parts of the network, expecting each VPN to follow its respective Flex Algo path. To simulate impairments, we induced delays while devices employed dynamic delay measurement on the link and observed the traffic taking different paths within the same Flex Algo plane. During our testing, we simulated a failure on the only non-red link on a node. The available paths were limited to the link with the admin group "RED," which should have been excluded from Flex-Algo 129. In normal circumstances, traffic would have fallen back to algo 0. However, the vendors had implemented a knob to prevent this fallback mechanism. And as expected when algo 0 fallback was disabled, we observed complete loss of traffic for that particular Flex Algo after the failure event.

The participating devices in the Flex Algos: Cisco NCS 540-24Q8L2DD, Huawei NetEngine 8000 F8, Juniper PTX10001-36MR, Juniper ACX7100-32C, Keysight IxNetwork, Nokia 7750 SR-1.

Traffic Generator: Keysight IxNetwork

Flex Algo Prefix Metric over OSPF

Flexible Algorithm can provide the optimal path to a destination in a remote area or IGP domain. The RFC9350 outlines a sub-TLV for OSPF Flexible Algorithm Prefix Metric (FAPM) so the calculation of the best path across multiple areas will take into account the constraints used for Flexible Algorithm paths.

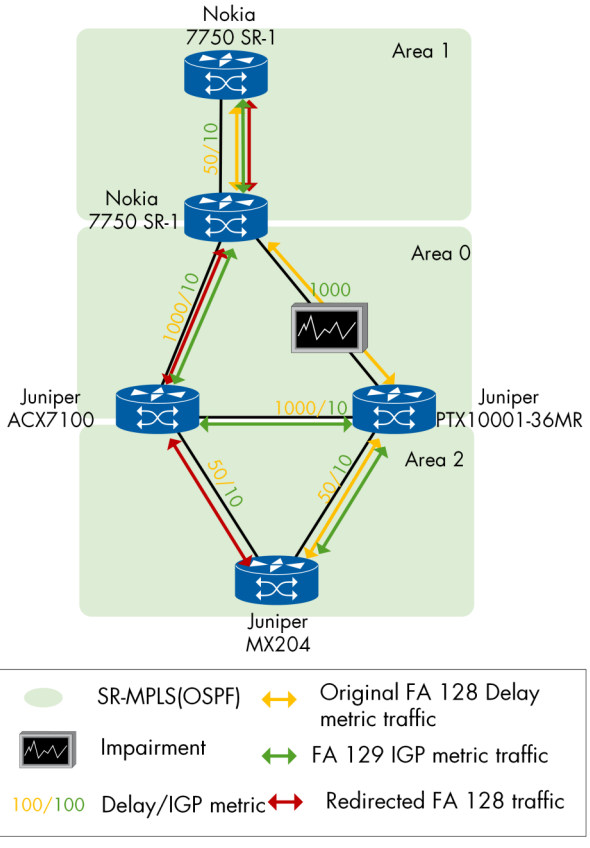

Figure 60: FAPM over OSPF

The test was carried out successfully with the following devices: Juniper MX204, Juniper PTX10001-36MR, Juniper ACX7100-32C, Nokia 7750 SR-1

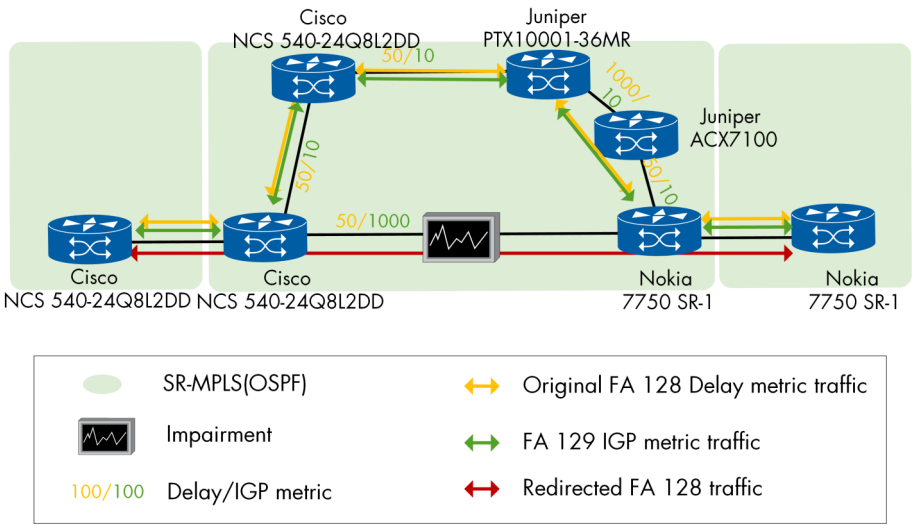

Figure 61: FAPM over OSPF

Traffic Generator: Keysight IxNetwork, Impairment device: Calnex SNE

FAPM to allow optimal end-to-end path for an interarea prefix. The area border router (ABR) must include the FAPM when advertising the prefix between areas that is reachable in that given Flexible Algorithm. The testing was conducted in a setup comprising three OSPF areas, with both Flex Algo 128 (based on delay metric) and Flex Algo 129 (based on IGP metric) configured in all areas. In order to advertise a prefix between areas, the area border router (ABR) included the FAPM for the corresponding Flex Algo. To establish a tunnel between the endpoints (PEs), the delay-metric FA128 was utilized to select the path with the least delay through all three domains. When there was a change in delay within the middle area, the tunnel was switched to ensure that the delay-metric FA128 path with the least delay was maintained.

The test was carried out successfully with the following devices: Cisco NCS 540-24Q8L2DD, Juniper MX204, Juniper PTX10001-36MR, Juniper ACX7100-32C, Nokia 7750 SR-1, Traffic Generator: Keysight IxNetwork, Impairment device: Calnex SNE

In order to enable FAPM, the FAD flags-TLV requires the M-flag to be set when advertising to ensure OSPF routers use Flex-Algorithm aware metrics for inter-area routing. During testing, one vendor had to correct their M-flag implementation to successfully complete the test.

SR-MPLS—SR-TE Traffic Steering

Traffic engineering involves the ability to manipulate the routing path of specific traffic from the network edge to its destination. This can be particularly useful in scenarios where network congestion needs to be managed and the quality of service for specific traffic, such as that of gold customers, needs to be prioritized. At the network edge, traffic can be directed along a path that ensures sufficient bandwidth or minimal delay for these high-priority traffic flows.

RFC 9256 (SR Policy Architecture) details the concept of an SR Policy and its associated steering mechanisms.

A headend can steer a packet flow into a valid SR Policy in various ways:

- Binding SID Steering: Incoming packets have an active SID matching a local BSID at the headend.
- Per-Destination Steering: incoming packets match a BGP/Service route, which recurses on an SR Policy
- Per-Flow Steering: incoming packets match or recurse on a forwarding array of which some of the entries are SR Policies.
- Policy-Based Steering: incoming packets match a routing policy that directs them on an SR Policy.

Initially, SR TE was implemented in SR-MPLS using binding SIDs then prefix-based steering.

Finally, traffic flow characteristics, such as the DSCP value, were employed to steer traffic along a specific path in order to optimize network performance. Keysight IxNetwork was used as a traffic generator for these tests.

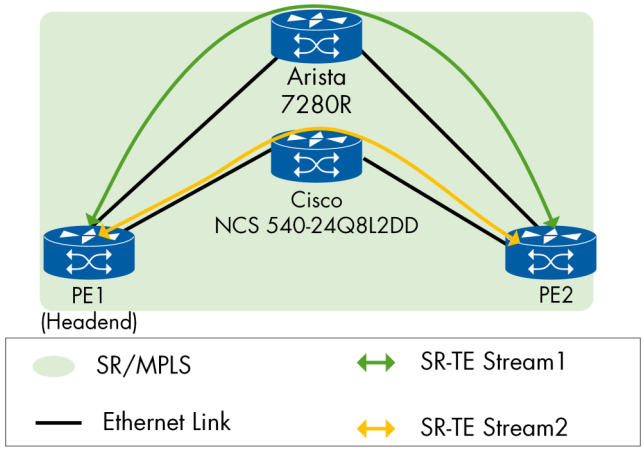

Figure 62: SR-MPLS—SR-TE Traffic Steering

PE 1 (Head End)	PE 2
Nokia 7750 SR-1	Juniper PTX10001-36MR
Juniper PTX10001-36MR	Nokia 7750 SR-1
Huawei NetEngine 8000 F8	Nokia 7750 SR-1
Cisco NCS 540-24Q8L2DD	Nokia 7750 SR-1
Ciena 5166	Juniper PTX10001-36MR
Nokia 7750 SR-1	Huawei NetEngine 8000 F8
Huawei NetEngine 8000 F8	Cisco NCS 540-24Q8L2DD
Cisco NCS 540-24Q8L2DD	Huawei NetEngine 8000 F8
Ribbon NPT 2100A	Huawei NetEngine 8000 F8
Nokia 7750 SR-1	Cisco NCS 540-24Q8L2DD

Table 9: Binding SID-based steering

PE 1 (Head End)	PE 2
Nokia 7750 SR-1	Juniper PTX10001-36MR
Juniper PTX10001-36MR	Nokia 7750 SR-1
Huawei NetEngine 8000 F8	Ribbon NPT-2100A
Cisco NCS 540-24Q8L2DD	Nokia 7750 SR-1
Nokia 7750 SR-1	Cisco NCS 540-24Q8L2DD
Cisco NCS 540-24Q8L2DD	Arista 7280R
Arista 7280R	Cisco NCS 540-24Q8L2DD
Ericsson 6673	Juniper PTX10001-36MR

Table 10: Flow-based steering

PE 1 (Head End)	PE 2
Nokia 7750 SR-1	Juniper PTX10001-36MR
Juniper PTX10001-36MR	Nokia 7750 SR-1
Cisco NCS 540-24Q8L2DD	Nokia 7750 SR-1
Ciena 5166	Juniper PTX10001-36MR
Nokia 7750 SR-1	Huawei NetEngine 8000 F8
Huawei NetEngine 8000 F8	Cisco NCS 540-24Q8L2DD
Cisco NCS 540-24Q8L2DD	Huawei NetEngine 8000 F8
Ribbon NPT-2100A	Huawei NetEngine 8000 F8
Nokia 7750 SR-1	Cisco NCS 540-24Q8L2DD
Huawei NetEngine 8000 F8	Ribbon NPT-2100A
Keysight IxNetwork	all previous devices

Table 11: Destination/Prefix based steering

SR-MPLS Inter-Domain SR-TE Assisted by BGP-LS

In large-scale networks, it is common to have multiple domains or Autonomous Systems (AS) due to the need to control the scale of Interior Gateway Protocols (IGP). By dividing a network into smaller domains, IGP can be kept under control, and better control can be maintained over network performance. The issue of end-to-end tunnels was ultimately encountered, and one of the proposed solutions to this problem is the use of BGP LS.

Usually, BGP-LS information is used by an SDN controller to steer SR-TE, SRv6, or RSVP LSPs in multi-area or multi-AS provider networks. In this use case, two different roles have been tested: first is the ability to translate IGP-TE information into BGP-LS messages (Ericsson and Arista nodes), and second is the capability to digest BGP-LS messages for building own Traffic-Engineering Database, which is used for constructing inter-AS (or inter-level, inter-area) SR-TE LSPs (Juniper node acting as head-end). This allows using such head-end nodes for distributed ingress PE calculation of paths, as an alternative to the use of centralized SDN.

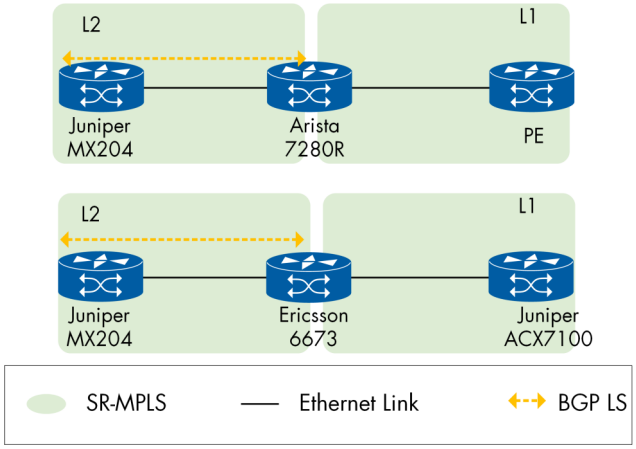

Figure 63: SR-MPLS Inter-Domain SR-TE Assisted by BGP-LS

Head end: Juniper MX204

ABR: Arista 7280R, Ericsson 6673

Inter AS BGP Classful Transport

draft-ietf-idr-bgp-ct-02 describes the service mapping to express the association of overlay routes with underlay routes satisfying a certain SLA using BGP.

The objective of this technology is to maintain end-to-end service intent across AS boundaries. This is accomplished by classifying tunnels with similar intents into transport classes and propagating this information across domains using BGP. As a result, when a service passes through each domain, it is directed to a specific transport tunnel class that aligns with its intended purpose.

The test setup consisted of two domains, where each PE node did the mapping of service prefixes with color communities to transport routes associated with transport-target communities.

In one domain, the prefixes were mapped to specific Flex-Algo values (128 and 129), while in the other domain, they were mapped to RSVP-TE paths (gold and bronze).

BGP-CT sessions were established between PEs and ASBRs, as well as between ASBRs themselves.

PE1 associated the received prefixes with their corresponding color BGP-CT label and local tunnel for ASBR1, based on the color community. Consequently, ASBR2 routed traffic to the Flex-Algo for PE2 based on the color of the incoming BGP-CT label.

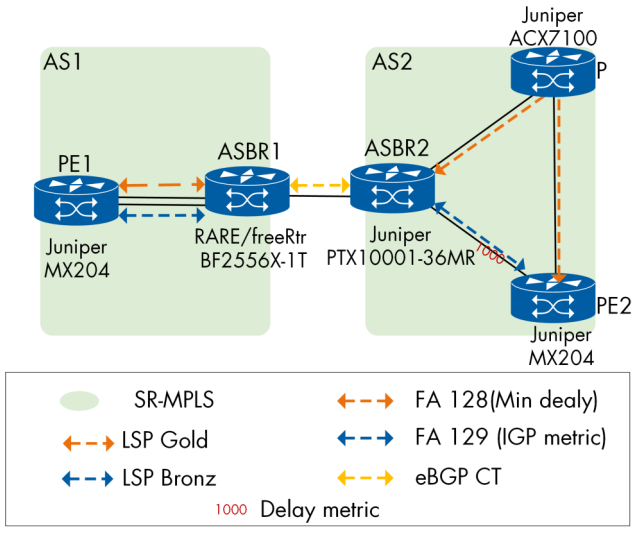

Figure 64: Inter As BGP Classful Transport

As ASBR: Juniper PTX10001-36MR, RARE/freeRtr BF2556X-1T

As PE: Juniper MX204, as P: Juniper ACX7100-32C

Traffic Generator: Spirent-STC

The data plane was tested by sending traffic between PE1 and PE1. Both PE1 and PE2 advertised service routes with color:0:128 (Gold SLA), and it was observed that the data traffic was correctly forwarded according to the intended specifications in both AS1 (using Gold RSVP-TE LSPs) and AS2 (using FA-128).

Inter AS SR-MPLS

Inter-AS connectivity is an essential aspect of modern network design that enables Service Providers to offer end-to-end services across multiple autonomous systems (AS). RFC 4364 describes two widely deployed methods for achieving inter-AS connectivity: Inter-AS Option B and Inter-AS Option C. These methods provide Service Providers with the flexibility to extend their networks while maintaining control over their own routing policies.

We have conducted tests for both inter-AS options B and C according to RFC 4364. However, during our testing, we encountered an issue between two Autonomous Boundary Routers (ABRs), as each ABR supported different SR Global Block (SRGB) ranges. To resolve this issue, we introduced a third ABR which is capable of stitching labels between inconsistent SRGB and/or dynamic label ranges. This validated the co-existence of BGP-SR with domains with heterogenous SRGBs and/or non-SR domains.

Arista 7280R and Ericsson 6673 tested as the ABRs and Ciena 5166 and Huawei NetEngine 8000 F8 tested as PEs.

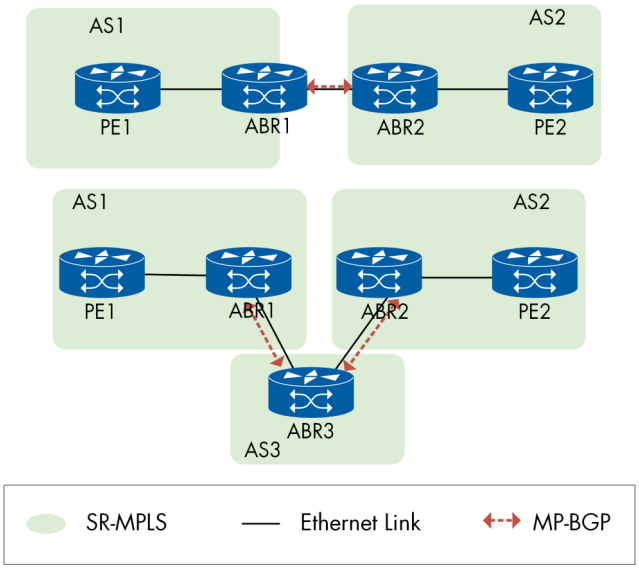

Figure 65: Inter AS Option C

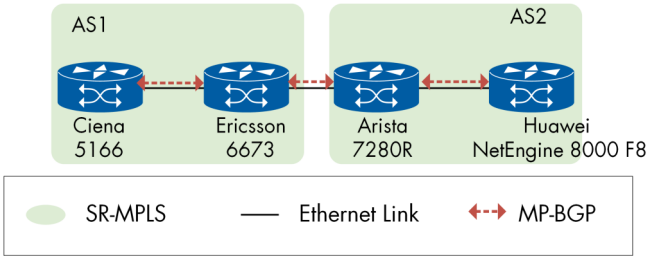

Figure 66: Inter AS Option B

PE1	ABR1	ABR2	PE2
Ribbon NPT-2100A	Arista 7280R	Juniper PTX10001-36MR	Juniper MX204
Ciena 5166	Cisco NCS 540-24Q8L2D	Juniper PTX10001-36MR	Juniper MX204
Ciena 5166	Cisco NCS 540-24Q8L2D	Juniper PTX10001-36MR	Arista 7280R

Table 12: Inter AS SR-MPLS—First setup

PE1	ABR1	ABR3	ABR2	PE2
Ribbon NPT-2100A	Arista 7280R	Juniper PTX10001-36MR	Cisco NCS 540-24Q8L2D	Ciena 5166
Huawei NetEngine 8000 F8	Arista 7280R	Juniper PTX10001-36MR	Cisco NCS 540-24Q8L2D	Ribbon NPT-2100A

Table 13: Inter AS SR-MPLS—Second setup

LDP and SR Interworking

Service providers often use a combination of MPLS transport and the LDP signaling protocol at the edges of their networks. Although LDP is simple to implement, it does not offer advanced traffic engineering and path repair features that are often necessary in the core of the network. To address this, we tested the use of an SR mapping server to enable interoperability between SR and LDP networks.

The configuration of prefix-to-SID mappings was done on the mapping server, which is then advertised in the ISIS on behalf of non-SR-capable nodes.

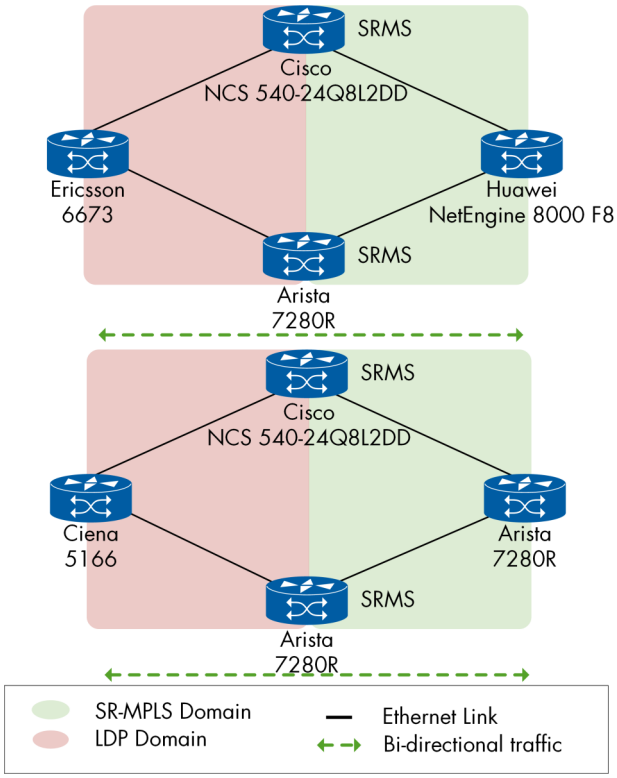

Figure 67: LDP and SR Interworking

SRMS: Arista 7280R, Cisco NCS 540-24Q8L2DD

SR mapping client: Arista 7280R, Huawei NetEngine 8000 F8

LDP only nodes: Ciena 5166, Ericsson 6673

Traffic Generator: Keysight IxNetwork

SR-MPLS Performance Measurement

RFC 5357 defines the "Two-Way Active Measurement Protocol" (TWAMP), which is utilized to evaluate network performance and troubleshoot network issues. The testing process begins with the control endpoint, which initiates the test by transmitting control packets to the sender. The sender generates the test traffic and returns it to the control endpoint, where performance metrics are computed based on the analysis of the test traffic.

Initially, we validated the device's capability to measure and identify changes in link delay that we introduced on the link using an impairment device. Subsequently, we confirmed the propagation of these measurements across the network using IGP Traffic Engineering (TE) Metric Extensions (IS-IS (RFC 7810) / OSPF (RFC 7471)).

In each combination, PE1 and PE2 were the sender and reflector. One vendor does not support advertising the delay in ISIS TLV using SR Policy.

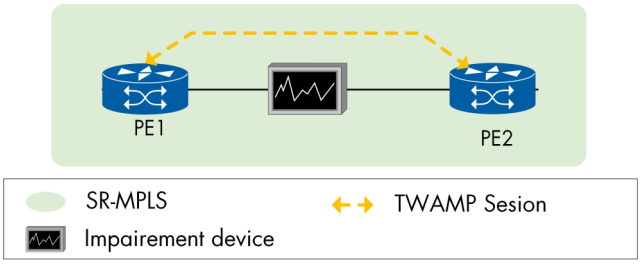

Figure 68: SR-MPLS Delay Measurement using TWAMP

PE1	PE2
Ciena 5166	Nokia 7750 SR-1
Ericsson 6673	Huawei NetEngine 8000 F8
Huawei NetEngine 8000 F8	Nokia 7750 SR-1
Cisco NCS 540-24Q8L2DD	Nokia 7750 SR-1
Cisco NCS 540-24Q8L2DD	Juniper MX204
Cisco NCS 540-24Q8L2DD	Ericsson 6673
Nokia 7750 SR-1	Juniper PTX10001-36MR
Cisco NCS 540-24Q8L2DD	Huawei NetEngine 8000 F8

Table 14: SR-MPLS Performance Measurement

Topology Independent Loop Free Alternative over SR-MPLS

To test the link and SRLG TI-LFA over an SR-MPLS network, we created a topology consisting of four nodes, with each participating vendor configuring the network nodes for an L3VPN service. Prior to the link failure, traffic was forwarded from the ingress PE (PLR) to the directly connected egress PE. To simulate a link failure, we asked the egress PE vendor to disconnect the protected link between the egress and ingress nodes while traffic continued to flow from the generator toward the ingress PE.

Out-of-service times ranged from 3ms to 34ms. For local SRLG, PLR nodes used a port to repair the link fault, regardless of the cost, because it shared the same SRLG as the failed port. Failover times ranged from 3ms to 15ms for the two combinations we tested.

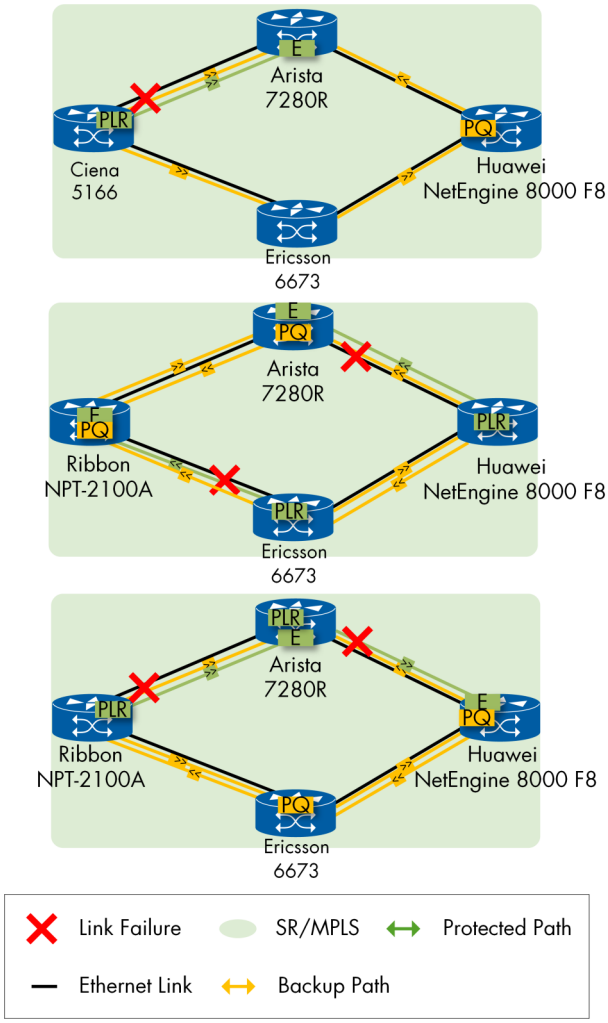

Figure 69: TI-LFA over SR MPLS

The following devices successfully participated in the test: Arista 7280R, Ciena 5166, Ericsson 6673, Huawei NetEngine 8000 F8, Ribbon NPT-2100A

Traffic Generator: Keysight IxNetwork

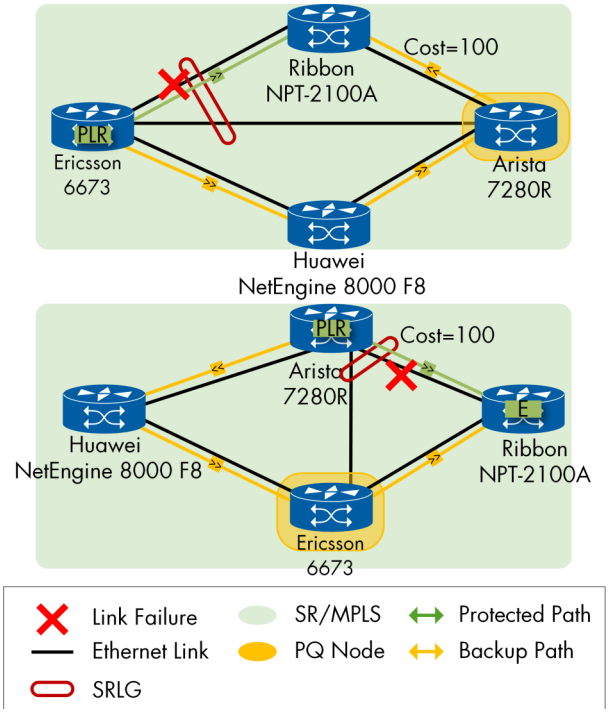

Figure 70: TI-LFA over SR-MPLS with Local SRLG

The following devices successfully participated in the test:

AS PLRs: Arista 7280R, Ericsson 6673

AS PQ and P nodes: Huawei NetEngine 8000 F8, Ribbon NPT-2100A

Traffic Generator: Keysight IxNetwork

Seamless BFD

Seamless BFD, or S-BFD, simplifies BFD usage by eliminating a large proportion of negotiation aspects, which leads to quick provisioning and improved control and flexibility for network nodes initiating path monitoring. In a test, we verified the ability of an SR policy to steer traffic into an SR-TE tunnel, and how S-BFD can detect link failures and trigger SR-TE hot standby protection.

To perform the test, we created a triangle topology consisting of an egress PE, ingress PE, and one P router. Each pair of PEs was configured with two SR-TE policies: a primary path and a backup. The initiator interval was set to 20 ms, allowing an acceptable out of service time between 40-100ms. We generated traffic between the Initiator and Reflector through the P node, which served as the longer primary path. To demonstrate S-BFD's role in network convergence, we emulated a tear-down session by shutting down a remote port and observed the traffic switch to the backup SR MPLS TE path in 75ms.

We also verified S-BFD sessions established between different vendors and configured an ACL to filter BFD packets, which resulted in the sessions being down and the traffic switching to the backup path.

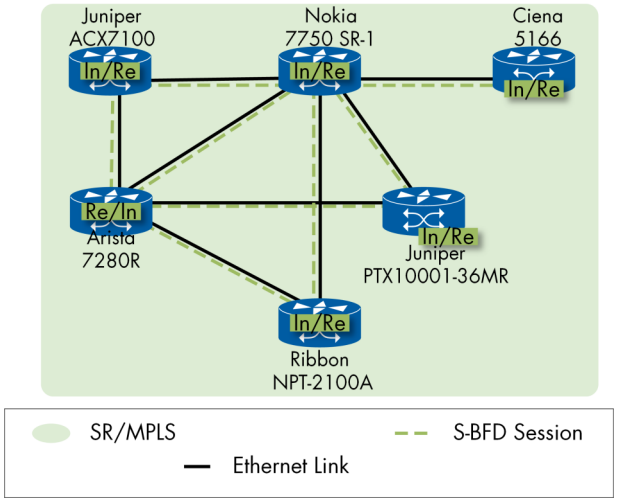

Figure 71: Seamless BFD

The following devices participated successfully as Initiator and reflectors:

Arista 7280R, Ciena 5166, Juniper PTX10001-36MR, Juniper ACX7100-32C, Nokia 7750 SR-1, Ribbon NPT-2100A

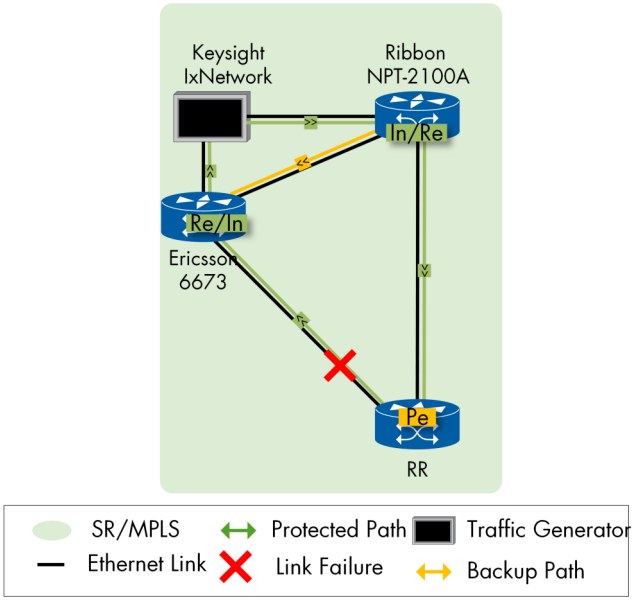

Figure 72: Seamless BFD

The following devices participated successfully as Initiator and reflectors: Ericsson 6673, Ribbon NPT-2100A

IPv6 BGP-LU

BGP-LU (Labeled Unicast) is used to provide connectivity between regions by advertising PE loopbacks and label bindings.

In this test, we verified using BGP to exchange reachability information among the routers in the network, including the IPv6 prefixes and the next-hop information.

The Spine node established BGP peering with two neighbors and activated labeled-unicast for IPv6 address family, allowing the router to forward traffic using MPLS labels. The Leaf nodes were configured to advertise the BGP prefix SID attribute in the BGP LU NLRI.

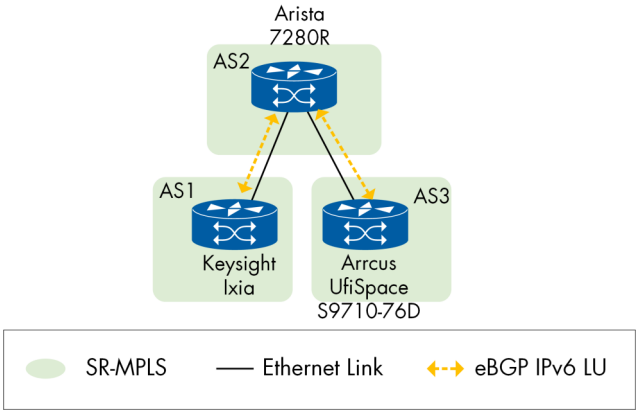

Figure 73: IPv6 BGP-LU

Arrcus UfiSpace S9710-76D, Arista 7280R, Keysight IxNetwork

400GE ZR+

The 400GE ZR+ specification is expected to play a key role in enabling extended reach in packet transport for metro and regional applications, which will be required to support the growing demand for cloud services, 5G wireless networks, and other bandwidth-intensive applications.

During the test, we conducted interoperability verification of 400ZR+ implementations across various vendors' routers and pluggable modules (Ciena/Cisco/Juniper).

A control plane using SR-MPLS (OSPF) was established between multiple nodes, and we monitored 400Gb/s bidirectional passing through three nodes, all of which were connected using 400ZR+ without any problems.

Additionally, we successfully tested channelization to allocate different bandwidths to different applications by dividing the link into four 100G channels between (Cisco/Juniper).

Juniper's streaming telemetry for 400ZR+ allowed us to verify important 400G-ZR+ standard parameters such as chromatic dispersion, oSNR, eSNR, module temperature, carrier offset frequency, wavelength, input and output power, pre-FEC bit error rate. By monitoring these parameters, we were able to ensure that the equipment was functioning as expected and delivering the necessary quality of service.

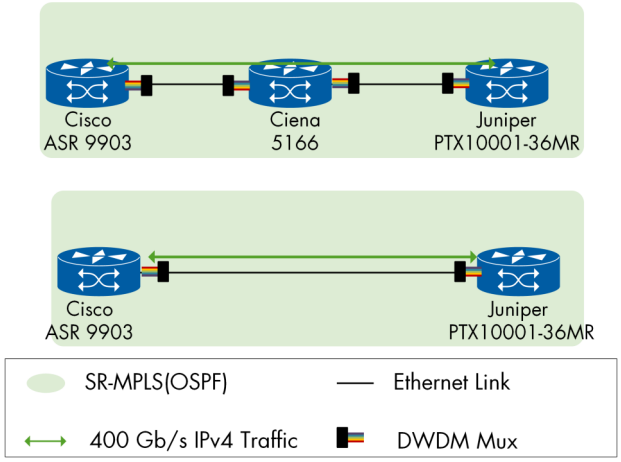

Figure 74: 400GE ZR+

400GE+ ZR was verified between: Cisco ASR 9903, Ciena 5166, Juniper PTX10001-36MR

Traffic Generator: Keysight IxNetwork

Segment Routing Anycast

Anycast-SID is an important component in Segment Routing, providing improved node resiliency, traffic load-sharing, and the ability to create separate network planes for different types of traffic. An Anycast-SID is a Node Prefix-SID that is advertised by multiple nodes, typically two, forming an anycast set. By including the Anycast-SID in the SID list of an SR policy path, traffic load-sharing, and resiliency can be improved.

Our network architecture consisted of five devices in the anycast set. After configuring an SR policy on the head end node, traffic was steered and the Anycast-SID was included in the segment list. As a result, traffic could be load-balanced and directed towards the remote end, utilizing the anycast set as next-hops.

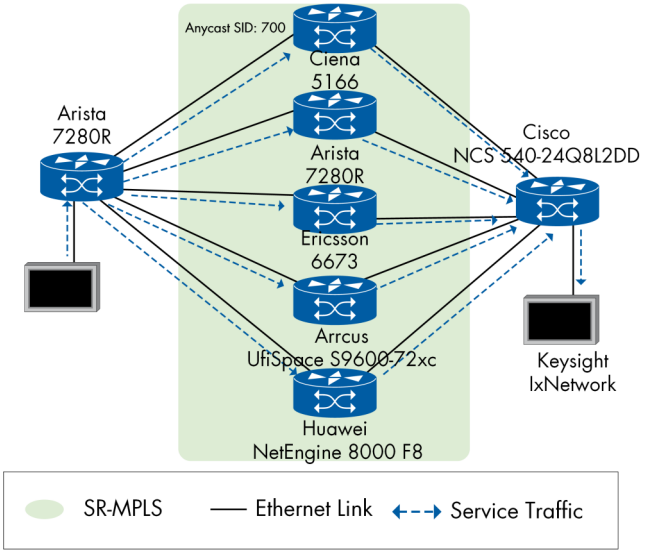

Figure 75: Segment Routing Anycast

The following DUTs participated successfully in this test:

Arrcus UfiSpace S9600-72XC, Arista 7280R, Ciena 5166, Ericsson 6673, Huawei NetEngine 8000 F8.

PE nodes: Arista 7280R, Cisco NCS 540-24Q8L2DD

Traffic Generator: Keysight IxNetwork

SDN

The need for centralized management of network devices through the use of a controller or path computation element is experiencing rapid growth due to the expansion of networks and the increasing complexity of networks' protocols in today's world. In response, Software Defined Networks (SDN) have gained significant attention and undergone notable advancements in recent years.

SDN is an innovative technology that separates network control from the data plane, enabling network administrators to manage their networks through software applications easily. This year's SDN testing efforts focused on key protocols such as PCEP and NETCONF, including BGP-LS, and explored the usage of SRv6 and SR-MPLS as data planes. The tests specifically explored the interoperability between Path Computation Clients (PCC) and Path Computation Elements (PCE), as well as between NETCONF capable network elements and NETCONF controllers. To prove interoperability, the tests focused on combinations of devices in which the controller and the routers are from different vendors. Additionally, the testing covered areas such as path computation, service provisioning, telemetry, and inventory use cases.

PCE Path Computation

The path computation test aims to verify the PCE's ability to provide network traffic paths in a multivendor environment, which is crucial to meet the network's application requirements. The test aims to ensure the interoperability between the PCE and each PCC independently, as the instantiation of LSP on one PCC should not depend on the other PCC.

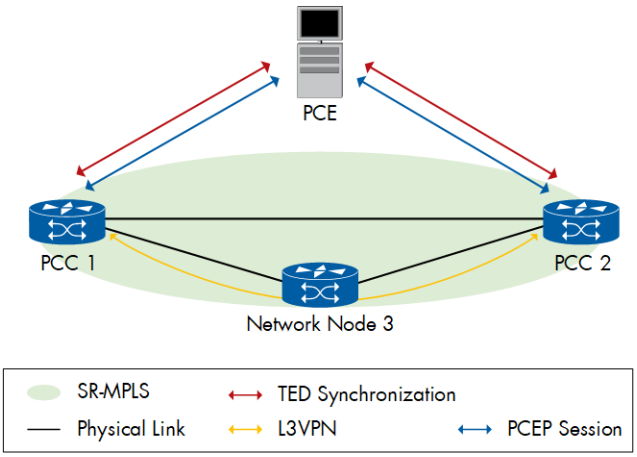

Figure 76: Path Computation with SR-MPLS

Table 15 lists the combinations that interoperate seamlessly over SR-MPLS data plane and PCE-initiated LSP without any known issues.

PCE	PCC
Juniper Paragon Pathfinder	Cisco NCS 540-24Q8L2DD
Juniper Paragon Pathfinder	Huawei NetEngine 8000 M4
Juniper Paragon Pathfinder	Nokia 7750 SR-1
Juniper Paragon Pathfinder	Ribbon NPT-2100A
Juniper Paragon Pathfinder	Ericsson 6673
Juniper Paragon Pathfinder	Ciena 5166
Cisco Crosswork	Juniper MX204
Cisco Crosswork	Ribbon NPT-2100A
Cisco Crosswork	Nokia 7750 SR-1
Cisco Crosswork	Ericsson 6673
Cisco Crosswork	Huawei NetEngine 8000 M4
Huawei NCE-IP	Nokia 7750 SR-1
Huawei NCE-IP	Juniper MX204
Huawei NCE-IP	Ericsson 6673
Huawei NCE-IP	Keysight IxNetwork
Huawei NCE-IP	Cisco NCS 540-24Q8L2DD
Keysight IxNetwork	Ericsson 6673
Keysight IxNetwork	Huawei NetEngine 8000 M4
Keysight IxNetwork	Juniper MX204

Table 15: PCE-Initiated Path Computation over SR-MPLS Data Plane

In this test, we used SR-MPLS as the data plane. We executed two tests, one in which the PCE initiated the instantiation of LSP paths and another in which the PCC initiated it.

- The DUTs initiated the IGP adjacencies, and we confirmed that the connection was established. IS-IS was used as the IGP for this test.
- We validated the Stateful PCEP session, PCE path instantiation, and LSP state synchronization.
- Instead of creating VPN services, we used LSP-pings to ensure that transport paths were installed correctly for this test.

The combinations that completed the test with SR-MPLS as data plane and PCC-initiated path are listed in Table 16.

PCE	PCC
Juniper Paragon Pathfinder	Cisco NCS 540-24Q8L2DD
Juniper Paragon Pathfinder	Ribbon NPT-2100A
Cisco Crosswork	Juniper MX204
Cisco Crosswork	Nokia 7750 SR-1
Cisco Crosswork	Ciena 5166

Table 16: PCC-Initiated Path Computation over SR-MPLS Data Plane

While conducting the test, we encountered an issue wherein not all vendors could support both PCC and PCE-initiated paths. One router faced an issue updating the IGP metric as the PCE sent an update message without a symbolic path name, and the router rejected the creation.

Signaling a Segment Routing Policy via PCEP

As more and more networks move from traditional IP / MPLS to segment routing, signaling Segment Routing policies from controllers to routers becomes critical.

An SR Policy (RFC 9256) is made up of a set of SR Candidate Paths that all share the same <headend, color, endpoint> tuple.

IETF draft "PCEP extension to support Segment Routing Policy Candidate Paths" (draft-ietf-pce-segment-routing-policy-cp) extends [RFC8664] to fully support the SR Policy construct.

With it, an SR Policy is modeled in PCEP as an Association of one or more SR Candidate Paths. PCEP extensions are defined to signal additional attributes of an SR Policy which were not covered by [RFC8664].

This test confirmed that the Path Computation Element with PCEP can effectively signal Segment Routing policies to Path Computation clients in a multi-vendor environment. The main goal of the test was to verify that the PCE and each PCC could operate together seamlessly without one PCC's colored policy signaling being reliant on the other PCC.

To conduct the test, the following steps were performed. First, the DUTs established IGP adjacencies using IS-IS, which were confirmed to be established. Second, the PCEP session, PCE path instantiation, and LSP state synchronization were validated. Third, LSP-pings were used instead of creating VPN services to verify proper path initiation. Finally, the PCE signaled a colored SR-Policy to the PCC, and the test was conducted using SR-MPLS and SRv6 as the data plan.

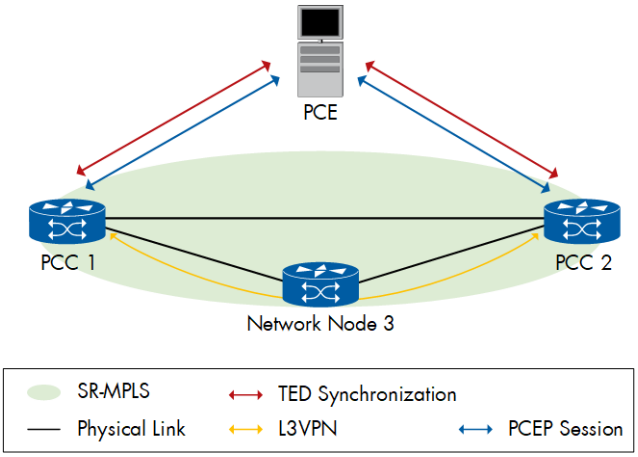

Figure 77: Segment Routing Policy Signaling with SR-MPLS

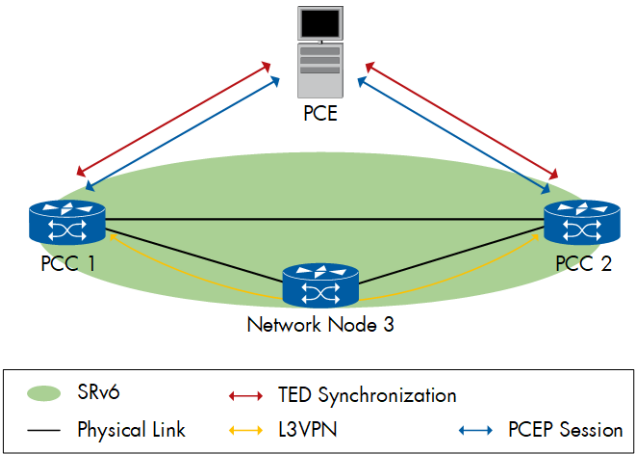

Figure 78: Segment Routing Policy Signaling with SRv6

Table 17 shows the combinations that interoperate seamlessly over SR-MPLS data plane.

PCE	PCC
Juniper Paragon Pathfinder	Cisco NCS 540-24Q8L2DD
Juniper Paragon Pathfinder	Ribbon NPT-2100A
Cisco Crosswork	Juniper MX204
Cisco Crosswork	Ribbon NPT-2100A
Cisco Crosswork	Huawei NetEngine 8000 M4
Cisco Crosswork	Ciena 5166
Huawei NCE-IP	Juniper MX204
Huawei NCE-IP	Keysight IxNetwork
Huawei NCE-IP	Cisco NCS 540-24Q8L2DD
Keysight IxNetwork	Juniper MX204
Keysight IxNetwork	Huawei NetEngine 8000 M4

Table 17: Colored Segment Routing Policy Signaling over SR-MPLS Data Plane

The combination that interoperated with SRv6 as a data plane and no known issues are listed in table 18.

PCE	PCC
NCE-IP	Juniper MX204
NCE-IP	Keysight IxNetwork
Keysight IxNetwork	Juniper MX204
Keysight IxNetwork	Huawei NetEngine 8000 M4

Table 18: Colored Segment Routing Policy Signaling over SRv6 Data Plane

Distribution of TE Policies and State via BGP-LS/SR

Effective distribution of Traffic Engineering (TE) policy and state is critical to network management. The PCE is responsible for computing paths according to TE policies and state. Meanwhile, the PCC requests and receives the calculated path from the PCE.

Border Gateway Protocol-Link State (BGP-LS) is a protocol for exchanging network topology and TE information between routers on a network. BGP-LS can distribute TE policy and state information to PCC and PCE at the same time, realizing efficient traffic forwarding and delivery.

The test aims to verify the interoperability of different vendor solutions for distributing TE policy and status via BGP-LS/SR between PCC and PCE. Tests were performed using SR-MPLS and SRv6 data planes.

- The DUTs initiated the IGP adjacencies, and we confirmed that the connection was established.
- We validated the BGP-LS sessions.
- The policy was delivered to the PCCs over BGP-SR
- The PCCs reported the status back to the PCE over BGP-LS

Instead of creating VPN services, we used LSP-pings to ensure that transport paths were instantiated.

The test combinations that interoperated with no observed problems over SR-MPLS data plan are shown in Figure 79, while Figure 80 shows the test combination with over SRv6 data plane.

We encountered an issue during the test where not all vendors supported BGP-LS for the distribution of TE policies.

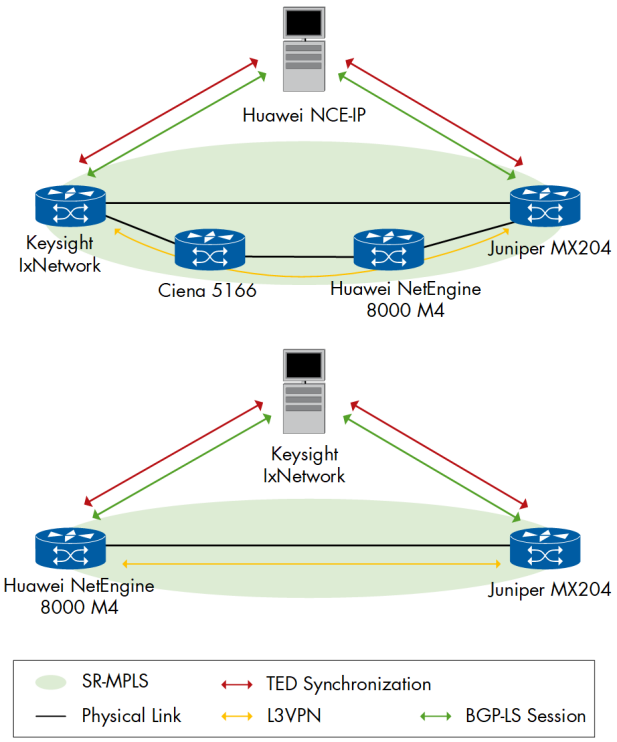

Figure 79: Distribution of TE Policies and State using BGP-LS with SR-MPLS

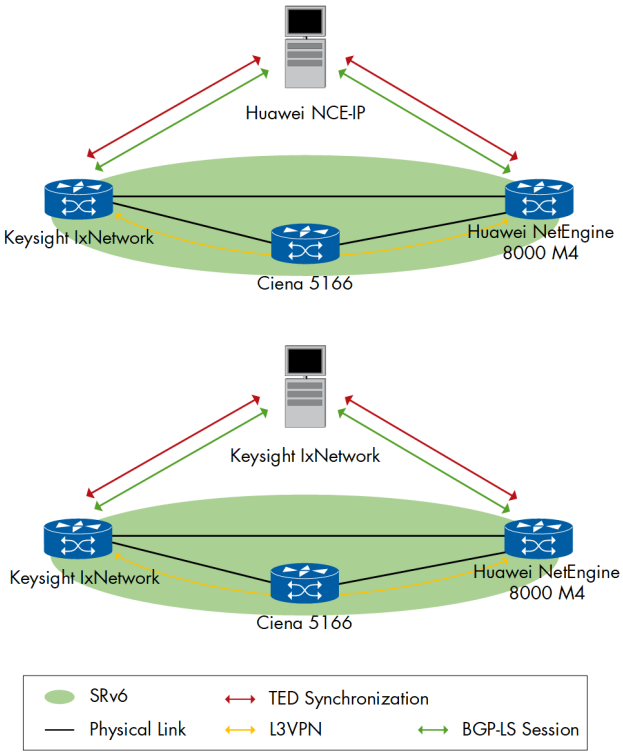

Figure 80: Distribution of TE Policies and State using BGP-LS with SRv6

PCC—Dynamic Paths Instantiation

Dynamically instantiating a Segment Routing policy at the headend PCC enables the VPN traffic to be automatically steered through the SDN network along the best available path based on real-time network conditions. This allows for better traffic optimization and faster packet delivery, ensuring that SLA requirements are met.

In the test, we advertised L3VPN routes from the PE router to the PCC, marking them with a specific color extended community. This was done to trigger on-demand segment routing policies calculated on the headend PCC itself to minimize latency towards the BGP next hop. Once the policy was calculated, the PCC reported it to the PCE.

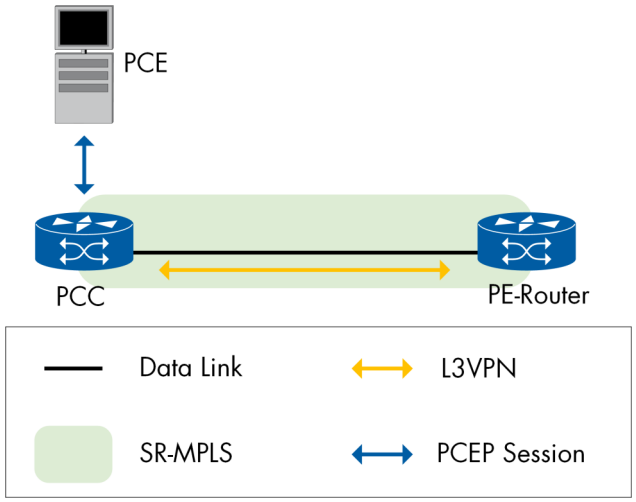

Figure 81: PCC—Dynamic Paths Instantiation

The combinations shown in Table 19 represent the devices that successfully participated in the test.

During the test, one PCC could not compute the policy itself. The PCC advertised its link state information (LS) over BGP, establishing a peering connection with the PCE. The PCC then initiated the SR policy, and the PCE computed a path based on the IGP metric and delegated it back to the PCC.

PCE	PCC	PR Router
Juniper Paragon Pathfinder	Juniper MX204	Cisco NCS 540-24Q8L2DD
Juniper Paragon Pathfinder	Cisco NCS 540-24Q8L2DD	Juniper MX204
Juniper Paragon Pathfinder	Juniper MX204	Keysight IxNetwork
Cisco Crosswork	Juniper MX204	Cisco NCS 540-24Q8L2DD
Cisco Crosswork	Cisco NCS 540-24Q8L2DD	Juniper MX204
Cisco Crosswork	Cisco NCS 540-24Q8L2DD	Keysight IxNetwork
Nokia Network Service Platform - NSP	Ciena 5166	Cisco NCS 540-24Q8L2DD

Table 19: PCC—Dynamic Paths Instantiation

L3/L2 VPN Service Provisioning

L2 and L3 VPN provisioning is a common requirement in modern networks. However, configuring and deploying VPNs can be complex and time-consuming, especially when dealing with multiple vendors and protocols. NETCONF is a standardized protocol that allows programmatic management of network devices and simplifies the VPN creation and configuration process. This test ensures that L2 and L3 VPNs can be created and configured seamlessly across multivendor controllers and devices using NETCONF, ensuring interoperability and compliance with standards. The test has been conducted using both Standardized OpenConfig and vendor-specific YANG models.

Provisioning the VPNs, we ensured that the NETCONF sessions between the controllers and the routers were established and stable. For the L3VPN, we pushed the VRF configurations to both devices and verified connectivity using ping tests.

For the L2VPN, we pushed the EVPN VPWS configurations from the controller to the routers and verified connectivity using ping tests as well.

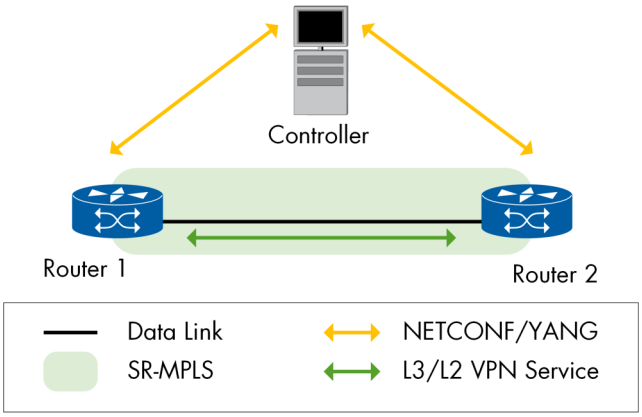

Figure 82: L3/L2 Service Provisioning

Table 20 shows the combinations that participated successfully in the L3VPN test. Meanwhile, table 21 presents the combinations successfully interoperated in the L2VPN test.

As OpenConfig and vendor-specific data models have been used for this test, we marked in the tables below the devices that are configured using OpenConfig. Otherwise, vendor-specific models have been used.

During the L3VPN test, one router could not receive the route target from the controller, which had to be manually set to complete the test.

In the L2VPN test, in two combinations, the configurations were pushed successfully to the routers. However, the service was down.

Controller	Router 1	Router 2
Keysight IxNetwork	Ciena 5166 (Open-Config)	Nokia 7750 SR-1 (Open-Config)
Keysight IxNetwork	Juniper MX204	Ericsson 6673 (Open-Config)
Nokia Network Service Platform - NSP	Ciena 5166	Cisco NCS 540-24Q8L2DD
Nokia Network Service Platform - NSP	Ciena 5166 (Open-Config)	Juniper MX204
Nokia Network Service Platform - NSP	Cisco NCS 540-24Q8L2DD	Juniper MX204
Nokia Network Service Platform - NSP	Juniper MX204	Huawei NetEngine 8000 M4
Huawei NCE-IP	Cisco NCS 540-24Q8L2DD	Nokia 7750 SR-1
Huawei NCE-IP	Juniper MX204	Huawei NetEngine 8000 M4
Cisco Crosswork	Ericsson 6673	Cisco NCS 540-24Q8L2DD
Cisco Crosswork	Nokia 7750 SR-1 (Open-Config)	Juniper MX204
Cisco Crosswork	Cisco NCS 540-24Q8L2DD	Juniper MX204
Cisco Crosswork	Cisco NCS 540-24Q8L2DD	Nokia 7750 SR-1 (Open-Config)

Table 20: L3VPN Provisioning

Controller	Router 1	Router 2
Cisco Crosswork	Cisco NCS 540-24Q8L2DD	Juniper MX204
Cisco Crosswork	Cisco NCS 540-24Q8L2DD	Ericsson 6673
Cisco Crosswork	Cisco NCS 540-24Q8L2DD	Nokia 7750 SR-1
Nokia Network Service Platform - NSP	Nokia 7750 SR-1	Ciena 5166
Nokia Network Service Platform - NSP	Nokia 7750 SR-1	Huawei NetEngine 8000 M4
Nokia Network Service Platform - NSP	Nokia 7750 SR-1	Juniper MX204
Huawei NCE-IP	Nokia 7750 SR-1	Huawei NetEngine 8000 M4

Table 21: L2VPN Provisioning

Routing Policies Configurations

Routing policies determine how network traffic is routed through the network. These policies control traffic and ensure that it is routed along the best paths according to predefined criteria.

This test verified that the Routing policies could be configured in a multivendor environment using NETCONF.

We validated the connectivity between the routers using ping. The NETCONF controller pushed a configuration to the routers, which included an IP-Filtering rule. The configurations were applied successfully to the routers, and as expected, the ping test indicated that the connectivity between them was lost.

The following devices successfully participated in the test. Huawei NCE-IP as NETCONF controller, and Nokia 7750 SR-1 and Huawei NetEngine 8000 M4 as routers.

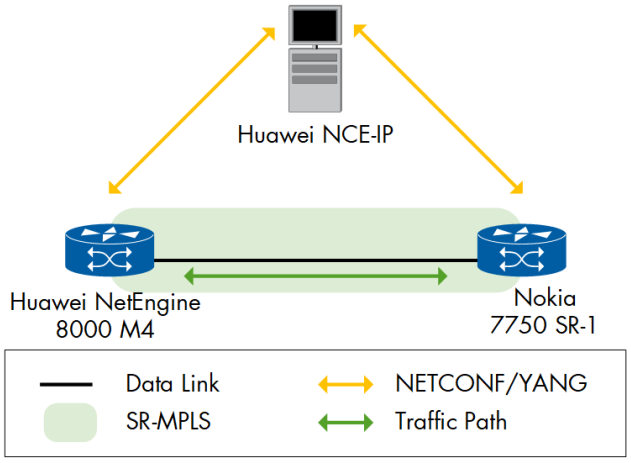

Figure 83: Routing's Polices Configuration

System Inventory

Retrieving network devices' information using NETCONF is beneficial for network administrators. This allows quick and easy access to network device information, such as configuration, status, and performance metrics.

In this test, we verified that using NETCONF, the controller can perform a device inventory in a multivendor environment.

After verifying the NETCONF session between the controller and the router, the controller retrieved the operating systems, software versions, system time-of-day, and configuration of the DNS resolver from the routers.

The pairs that participated in the test successfully are shown in table 22.

Due to an interoperability issue, one controller encountered a problem retrieving the configuration of the DNS resolver from one router.

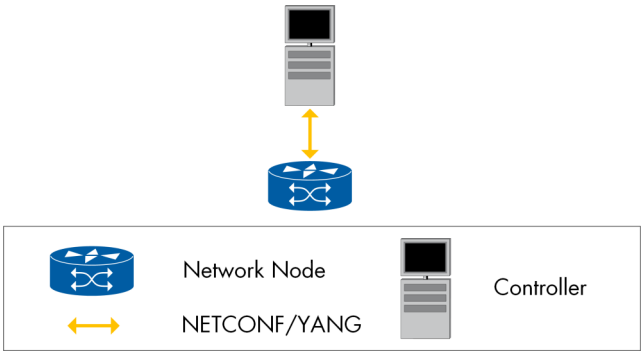

Figure 84: System Inventory

Controller	Network Node
Nokia Network Service Platform - NSP	Huawei NetEngine 8000 M4
Nokia Network Service Platform - NSP	Ciena 5166
Nokia Network Service Platform - NSP	Cisco NCS 540-24Q8L2DD
Nokia Network Service Platform - NSP	Juniper MX204
Huawei NCE-IP	Nokia 7750 SR-1
Huawei NCE-IP	Juniper MX204

Table 22: System Inventory

Telemetry—gNMI

gNMI, or gRPC Network Management Interface, is a remote procedure call-based protocol used for managing and monitoring network devices. In this test, we focused on monitoring, which involved retrieving telemetry data from the network devices. After verifying the NETCONF sessions status, the controller subscribed to the router using OpenConfig-interface YANG model. Table 23 shows the devices that participated successfully in the test.

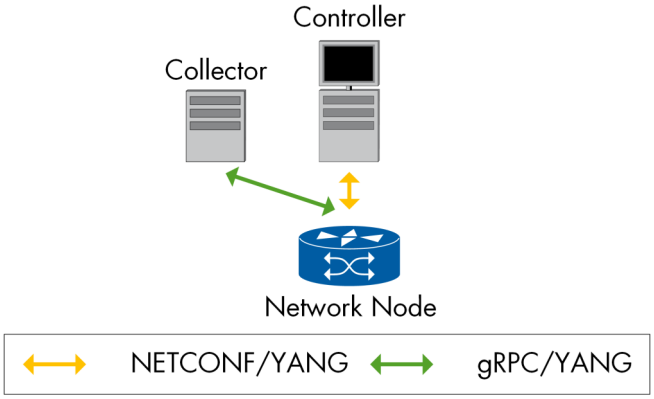

Figure 85: gNMI-Telemetry via the OpenConfig model

Controller/Collector	Network Node
Cisco Crosswork	Juniper MX204

Table 23: gNMI-Telemetry via OpenConfig Data Model

NETCONF Transport Slicing Controller

Network slicing is a crucial aspect of modern networking that enables the creation of multiple virtual networks on top of a single physical network infrastructure. The ability to create network slices allows service providers to offer customized services to their customers, each with specific requirements for performance, security, and other network attributes.

The YANG model provided in the draft 'draft-ietf-teas-ietf-network-slice-nbi-yang' was used for the northbound interface, while vendor-specific models were used for the southbound interface.

The following devices successfully interoperated: Nokia Network Service Platform - NSP as NETCONF controller and orchestrator, Huawei NetEngine 8000 M4, Ciena 5166, Juniper MX204 and Nokia 7750 SR-1 as routers.

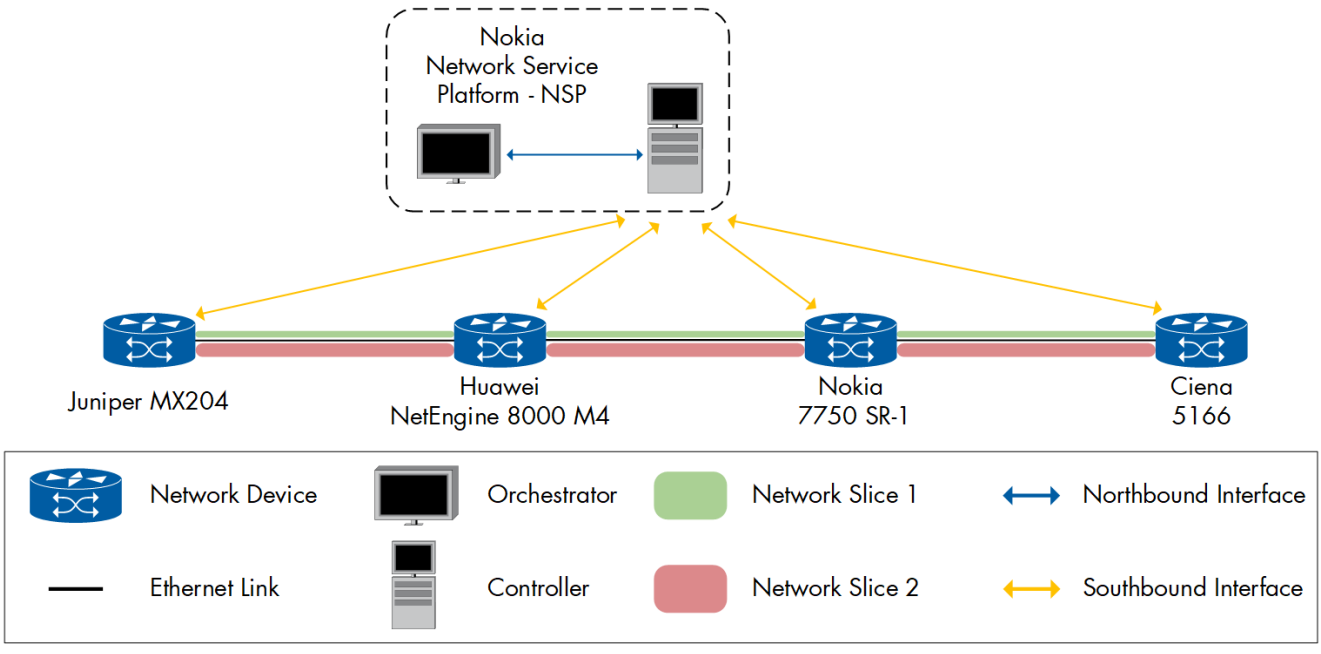

Figure 86: Network Slices

Time Synchronization

Time synchronization is a critical aspect of networking that is essential for any type of network, whether it be an enterprise, data center, service provider front/backhaul network, or otherwise. Implementing a robust time synchronization infrastructure is challenging and complex, requiring careful attention to detail.

At the EANTC MPLS/SDN Interoperability Testing event, we strive to create realistic and extensive tests with our partners, who are among the leading vendors in the industry. This year, we have focused on meeting the most recent industry's time synchronization requirements, which include hot topics such as general 5G and its synchronization needs, as well as Open RAN.

We have tightened our requirements to match the conditions of 5G and have also tested, for the first time, virtualized devices. Intel Ethernet Network Adapters E810-XXVDA4T and E810-CQDA2T network interface cards (NICs) enabled commercial-off-the-shelf (COTS) servers to serve as a Grandmaster, Boundary Clock, or Slave Clock, a capability that was previously exclusive to dedicated hardware devices. Additionally, we have tested an O-RAN topology with six nodes from different vendors, incorporating two independent timing paths to match a merge between LLS-C2 and LLS-C3 setups.

Delay Asymmetry Detection/Measurement

Assisted Partial Timing Support Delay Asymmetry

Asymmetric delay on links carrying PTP messages is critical, and one of the most difficult issues for network time synchronization. At this year's event, in order to showcase the participating devices' abilities to detect the applied asymmetry, we have tested two different Delay Asymmetry scenarios.

In the first scenario, with one Grandmaster, two Boundary Clocks, and one Slave Clock, the Grandmaster and Boundary Clock-1 were referenced to GNSS (via a splitter) through an antenna on the roof of EANTC's lab. PTP profile G.8275.2 was used across the whole chain. When the Grandmaster and Boundary clock-1 were locked to GNSS, Boundary Clock-2 and the Slave Clock were using Boundary Clock-1 as the source for timing. When GNSS to Boundary Clock-1 was disconnected, Boundary Clock 1 reverted to using PTP from the Grandmaster as its timing source, with the Slave Clock 1PPS absolute time error being measured across this transition. We restarted the measurements and introduced the delay asymmetry using the Calnex Paragon-X device, and waited the Boundary Clock-2 to detect the Asymmetry.

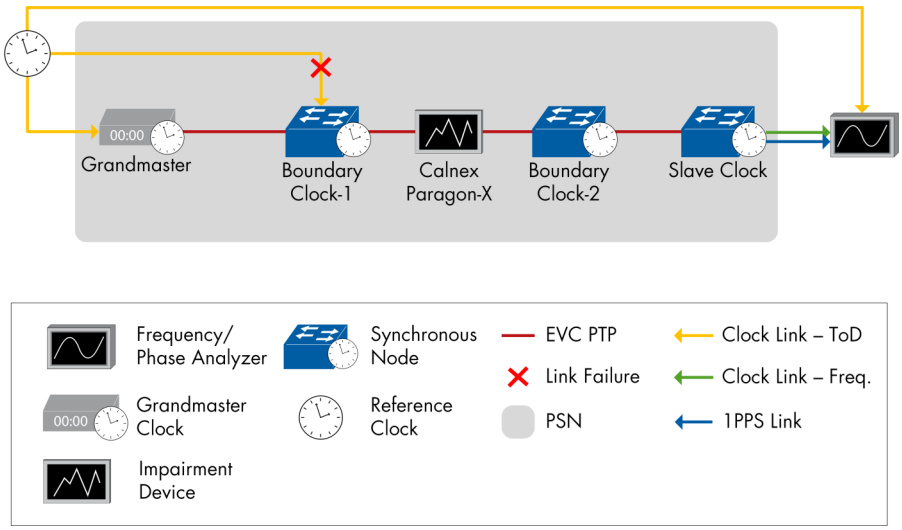

Figure 87: Assisted Partial Timing Support Delay Asymmetry

Grandmaster	Boundary Clock 1	Boundary Clock 2	Slave Clock	Time Error Analyzer
Microchip	Ericsson 6673	Cisco NCS	Ericsson 6675	Calnex Sentry
Microchip	Ericsson 6673	-	Juniper	Calnex Sentry
Microchip	Ericsson 6673	-	Intel E810-	Calnex Sentry

Table 24: Assisted Partial Timing Support Delay Asymmetry

During two testing combinations, the devices which have played the role of Boundary Clock-2 and had to use the 8275.2 PTP profile, had connectivity problems, and confusion with the configuration for this profile, so we had to exclude them from the test, and continue with one Boundary Clock, and one Slave Clock.

Time to lock and stabilize for devices using the G.8275.2 profile is typically longer than those using G.8275.1, which limited actual measurement duration.

Delay Asymmetry Measurement

The second asymmetry-related test performed was to detect and compensate the asymmetry either automatically, when supported, or manually.

The test cases are designed to cover the different methods vendors have of handling compensation in real-world implementations.

The test topology consisted of:

- Grandmaster-A (GM-A) connected to the GNSS as reference, used as main reference for the topology.
- Grandmaster-B (GM-B) connected to GNSS, used as backup.
- Boundary Clock (BC), connected to both GMs and configured using local priorities to select GM A when it is (or both GMs are) locked to GNSS

We employed three distinct methods for conducting the test, whereby the BC was linked to both GMs and set to lock on GM-A in all three methods.

We initiated the measurement process by detaching GM-A from the GNSS, which led to the BC locking onto GM-B. Nevertheless, we introduced the asymmetry subsequent to the GM switchover.

First Approach—Manual Delay Compensation: The BC followed the asymmetry and them the vendors engineers compensated the delay manually through the CLI commands. The following combinations have passed the test with this approach with 500 nanoseconds as one-way delay.

Second Approach—Automatic Delay Compensation: The BC followed the asymmetry and them the vendors engineers compensated the delay manually through the CLI commands. The following combinations have passed the test with this approach with 500 nanoseconds as one-way delay.

Third Approach—Manual Delay Compensation with GNSS Reference: The Boundary Clock, in this approach needed the manual compensation to overcome the introduced one-way delay, and needed the GNSS reference to detect the asymmetry. The one-way delay was 500 nanoseconds for this combination.

No interoperability issues were observed during this test, the only problem we faced was the lacking of Boundary Clocks which can have 2 slave ports with different PTP profiles as planned originally.

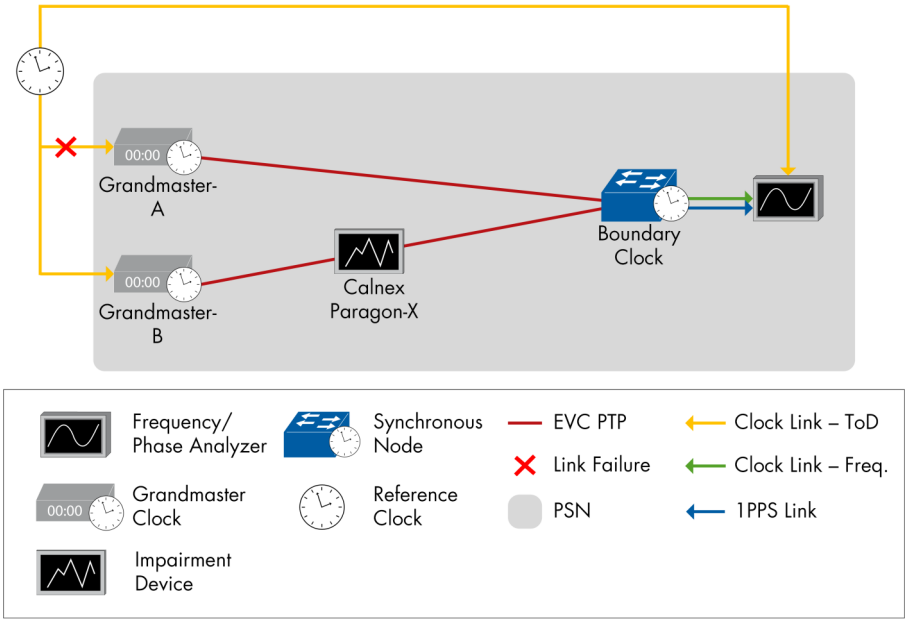

Figure 88: Delay Asymmetry Measurement

Grandmaster A	Grandmaster B	Boundary Clock	Time Error Analyzer
Ciena 5166	Microchip TimeProvider 4100	Juniper ACX-7100-48L	Calnex Sentry
Juniper ACX-7024	Huawei NetEngine 8000 M4	Cisco NCS 540X-16Z4G8Q2C	Calnex Sentry

Table 25: Delay Asymmetry Measurement—First Approach

Grandmaster A	Grandmaster B	Boundary Clock	Time Error Analyzer
Ciena 5166	Microchip TimeProvider 4100	Ericsson 6673	Calnex Sentry

Table 26: Delay Asymmetry Measurement—Second Approach

Grandmaster A	Grandmaster B	Boundary Clock	Time Error Analyzer
Ericsson 6675	Intel E810-XXVDA4T	Microchip TimeProvider 4100	Calnex Sentry

Table 27: Delay Asymmetry Measurement—Third Approach

Boundary Clock Class C/D Conformance Test

Class C/D Boundary Clocks have been specifically designed to fulfill the stringent demands for time synchronization in modern networks, thereby facilitating top-notch applications like ultra-reliable low-latency communication, which are integral to 5G mobile networks.

Over the years at the EANTC MPLS SDN Interoperability event, we have had the opportunity to witness the advancements in Boundary Clocks. Two years ago, it was unusual for a device to meet Class D specifications, whereas this year almost all Boundary Clocks tested passed the Class D conformance test.

This test is not one of interoperability as it tests the time error performance of only a single device, but it was used to qualify devices before their participation in the class D chain tests

The test was done by using the Calnex Paragon-neo to emulate the Grandmaster and the Slave Clock, with the device under test connected directly as the Boundary Clock. As per the requirements of G.8273.2, we measured the low pass filtered two-way time error with an applied limit of 5ns.

The boundary clocks enabled both PTP and SyncE on the link towards slave clock, as they configured the PTP 8275.1 hybrid profile. Additionally, we have performed the conformance test for Boundary Class C, complying with the latest ITU-T G.8273.2 Clause 7.1.4 by measuring the relative constant time error between two ports of the boundary clock.

Based on observations from previous years that a device's time error performance may vary across its different ports speeds, some devices were tested at various line rates.

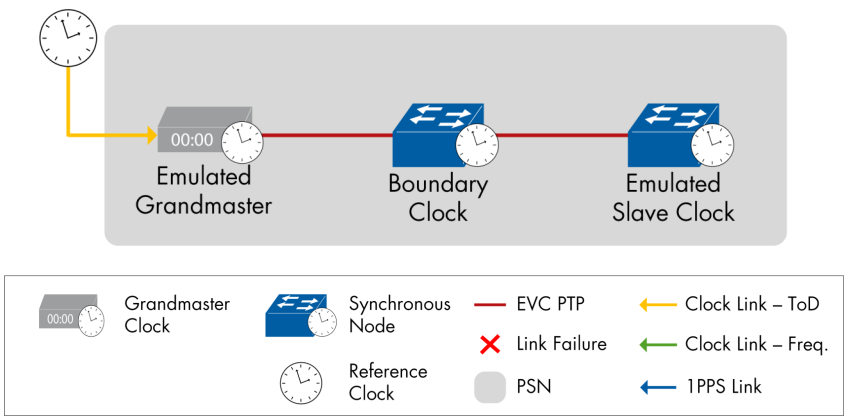

Figure 89: Boundary Clock Class D Conformance

The following devices have passed the test with the ports speeds:

Emulated GM	BC	Emulated SC	Port Speed
Calnex Paragon-neo	Ciena 5166	Calnex Paragon-neo	10 GbE
	Cisco NCS 540X-16Z4G8Q2C		10 GbE
	Ericsson 6673		10 GbE
	Huawei NE8000 M4		10 GbE
	Intel E810-XXVDA4T		10 GbE
	Intel E810-CQDA2T breakout		10 GbE breakout from 100 GbE
	Juniper ACX-7024		10 GbE
	Juniper ACX-7100-32C		10 GbE
	Juniper ACX-7100-48L		10 GbE
	Microchip TimeProvider 4100		1 GbE
	Microchip TimeProvider 4100		10 GbE
	Ericsson 6673		25 GbE
	Intel E810-XXVDA4T		25 GbE
	Ciena 5166		100 GbE
	Cisco NCS 540X-16Z4G8Q2C		100 GbE
	Ericsson 6673		100 GbE
	Intel E810-CQDA2T		100 GbE
	Juniper ACX-7024		100 GbE
	Juniper ACX-7100-32C		100 GbE
	Juniper ACX-7100-48L		100 GbE
	Ciena 5166		400 GbE
	Ciena 5166		400 GbE
	Juniper ACX-7100-32C		400 GbE
	Juniper ACX-7100-48L		400 GbE

Table 28: Boundary Clock Class D Conformance

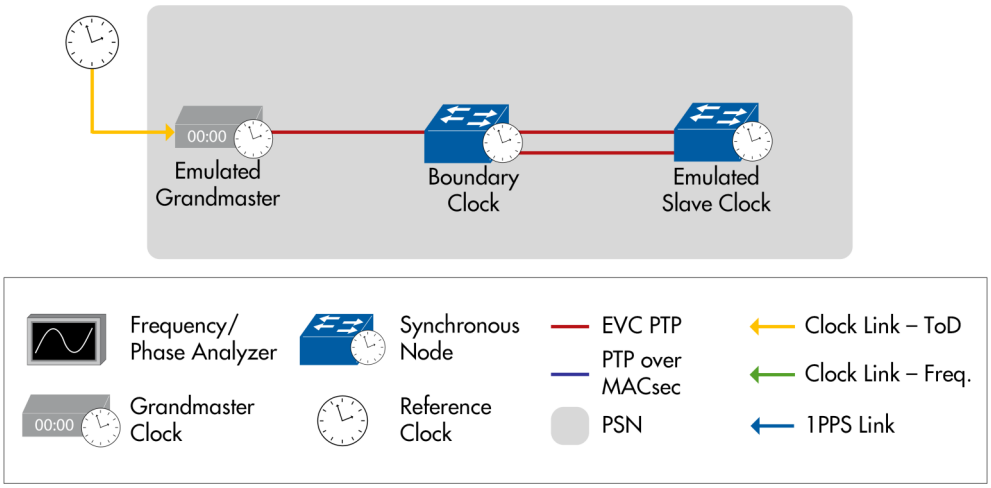

Figure 90: Boundary Clock Class C Relative Time Error

The following devices passed the class C with the ports speeds:

Emulated GM	BC	Emulated SC	Port Speed
Calnex Paragon-neo	Ciena 5166	Calnex Paragon-neo	10 GbE
	Ciena 5166		100 GbE
	Cisco NCS 540X-16Z4G8Q2C		10 GbE
	Cisco NCS 540X-16Z4G8Q2C		100 GbE
	Ericsson 6673		10 GbE
	Ericsson 6675		100 GbE
	Huawei NetEngine 8000 M4		10 GbE
	Huawei NetEngine 8000 M4		100 GbE
	Juniper ACX7100-32C		10 GbE
	Microchip TimeProvider 4100		10 GbE

Table 29: Boundary Clock Class D Conformance

High-Precision Clocking Source Failover

Testing time synchronization in a well-controlled lab environment typically yields favorable outcomes, but also fails to represent real-world conditions. This is the motivation behind this test case, which measures the time error produced by the Boundary Clock in a realistic topology with redundant Grandmasters. These Grandmasters may encounter GNSS connectivity interruptions which cause changes in the PTP source used by the Boundary Clock. This test case measures the time error of the Boundary Clock during the switchover between Grandmaster references (i.e. during the BMCA event), and also when the Boundary clock is in holdover due to both Grandmasters having lost GNSS connectivity. The used test topology consisted of:

- Grandmaster-A (GM-A) connected to the GNSS as reference, used as main reference for the topology.
- Grandmaster-B (GM-B) connected to GNSS, used as backup.
- Boundary Clock (BC), connected to both GMs and configured using local priorities to select GM A when it is (or both GMs are) locked to GNSS

At the test start, both GM-A and GM-B were locked to their GNSS inputs, and the Boundary Clock was locked via PTP to GM-A. The first measurement phase ran for 1000 seconds (to allow calculation of the Constant Time Error, cTE) The GNSS input to GM-A was then disconnected, causing the BC to select GM-B as its reference. The next step required also disconnecting GNSS from GM-B, forcing the BC to re-lock onto GM-A based on the configured local priorities. GNSS was reconnected to GM-B and then to GM-A. 1PPS and two -way-Time Error outputs from the Boundary Clock were measured at each of these steps.

Passing this test requires that the measured time error at the Boundary Clock output meets G.8271 accuracy level 6 or better, i.e. ≤ 260 ns.

No interoperability issues were seen during this test case, but one Grandmaster was observed to transmit clock Class 6 after disconnection from GNSS (rather than clock Class 7), which is non-compliant with the requirements of the relevant ITU-T recommendation G.8275.1, clause 6.4.

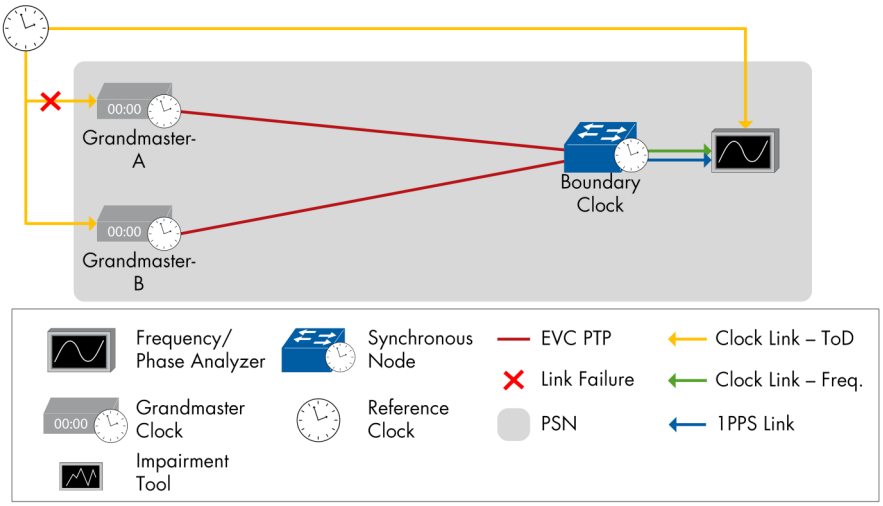

Figure 91: High-Precision Clocking Source Failover

Grandmaster A	Grandmaster B	Boundary Clock
Microchip TimeProvider4100	Ericsson 6673	Cisco NCS 540X-16Z4G8Q2C
Microchip TimeProvider4100	Ericsson 6673	Juniper ACX-7100-32C
Microchip TimeProvider4100	Ericsson 6673	Ericsson 6673
Microchip TimeProvider4100	Ericsson 6673	Microchip T imeProvider4100
Microchip TimeProvider4100	Ericsson 6673	Ciena 5166
Microchip TimeProvider4100	Ericsson 6673	Huawei NetEngine 8000 M4
Ericsson 6675	Microchip TimeProvider4100	Intel E810-XXVDA4T

Table 30: High-Precision Clocking Source Failover

Time Synchronization Source Failover

This test was a part of proposed resiliency tests of time synchronization, with adding an additional boundary clock at the end of the chain, which makes the topology more real-life scenario. The used test topology consisted of:

- Grandmaster A (GM A) connected to the GNSS as reference, used as main reference for the topology.
- Grandmaster B (GM B) connected to GNSS, used as backup.
- Boundary Clock-1 (BC-1), connected to both GMs and configured to prefer the GM A as long as it has the GNSS antenna, using the local priorities of the links.
- Boundary Clock-2 (BC-2), connected to BC-1 and to the Calnex Paragon-neo providing both PTP and SyncE measurements

The test started when both GM A, and GM B are locked on the GNSS reference. The Boundary Clock-1 was locked with both PTP and SyncE on the GM A. The first phase of measurement was started for 1000 seconds, to be able to calculate the Constant Time Error, then the GNSS connection of the GM A was disconnected, causing the BCs to choose the GM B as source. The next step started with disconnecting the GM B from the GNSS causing the BCs to re-lock on the GM A as both GMs have no GNSS and the local priorities were set on the BC to do so.

Then we reconnected GM B, then GM A while we are measuring the 1PPS, 2 Way-Time Error from the output of the Boundary Clock.

This test aimed to keep the time error from the output of the Boundary Clocks within the G.8271 accuracy level 6, all the following combinations passed the test successfully:

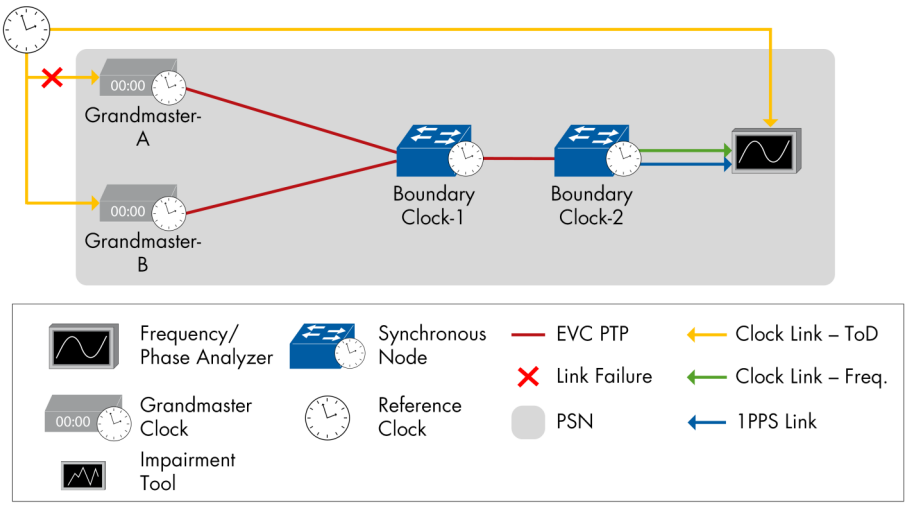

Figure 92: Time Synchronization Source Failover

Grandmaster A	Grandmaster B	Boundary Clock 1	Boundary Clock 2
Microchip TimeProvider 4100	Ciena 5166	Huawei NE8000 M4	Intel E810-XXVDA4T
Microchip TimeProvider 4100	Ciena 5166	Juniper ACX7100	Ericsson 6673
Microchip TimeProvider 4100	Ciena 5166	Ericsson	Juniper ACX-7100-48L
Microchip TimeProvider 4100	Cisco NCS 540X-16Z4G8Q2C	Ciena 5166	Microchip TimeProvider 4100
Ericsson 6673	Cisco NCS 540X-16Z4G8Q2C	Juniper ACX-7100-48L	-

Table 31: Time Synchronization Source Failover

Chain Ring of Class D Boundary Clocks

This test measured the accumulated time error of a chain of Class D Boundary Clocks in a ring topology, resembling a typical real-world service provider implementation. All participating devices had passed the G.8273.2 Class D Boundary Clock conformance test, indicating performance met the maximum absolute time error, low-pass filtered, max|TEL|, limit of 5ns.

The test topology consisted of:

- Grandmaster A (GM A) locked to GNSS generating clock Class 6 and ESMC QL-PRC. Used as the primary reference for the topology.
- Grandmaster B (GM B) locked to GNSS, generating clock Class 6 and ESMC QL-PRC. Used as backup reference.
- Boundary Clocks (BCs): Seven BCs formed a ring: all BCs were locked on to GM A using their local priorities configuration.
- Calnex Paragon-neo emulated a Slave Clock and measured the PTP time error.
- Calnex Sentry measured the 1PPS absolute time error.

Initially, all BCs were synchronized to PTP and SyncE from GM A, and a baseline measurement performed. Subsequently, GNSS was disconnected from GM A, causing all BCs to switch PTP and SyncE reference to GM B.

GNSS was then reconnected to GM A and a measurement performed while all BCs re-established synchronization with it and stabilized.

All BCs were configured with a one minute "Wait to Restore" (WtR) period, meaning each would wait one minute before reacting to a change in input quality i.e. before switching reference on receipt of clock Class 6 when GM A was reconnected to GNSS.

The following devices passed the test successfully:

GM-A: Microchip TimeProvider 4100

GM-B: Huawei NE8000 M4

BCs: Cisco NCS 540X-16Z4G8Q2C, Juniper ACX-7100-32C, Intel E810-XXVDA4T, Ericsson 6673, Ciena 5166, Juniper ACX-7100-48L, Microchip TimeProvider 4100

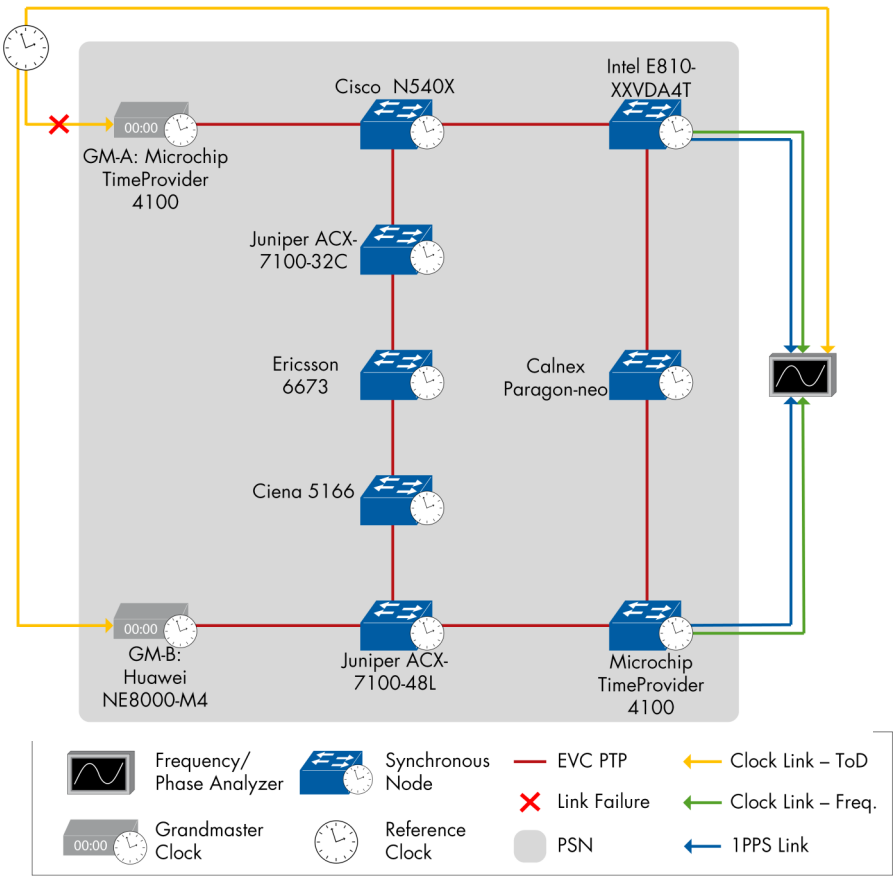

Figure 93: Chain of Class D Boundary Clocks

Phase/Time Holdover with Enhanced Sync-E Support

Enhanced Synchronous Ethernet (eSyncE) provides physical layer support to PTP-aware devices in full timing support networks, enhancing performance to enable the stringent synchronization requirements of modern telecommunication networks.

This test verified the ability of a chain of boundary clocks, configured to use eSyncE, to maintain acceptable values of time error during loss of the Grandmaster GNSS reference. eSyncE ESMC messages were also captured and analyzed to verify that the eSyncE TLV was being processed as required by each device in the chain.

The following devices passed the test successfully:

Grandmaster: Microchip TimeProvider 4100

Boundary Clocks: Cisco NCS 540X, Juniper ACX-7100 -32C, Intel E810-XXVDA4T, Huawei NE8000 M4, Ericsson 6673, Ciena 5166

O-RAN Fronthaul Network Time Synchronization

The Open Radio Access Network, commonly referred to as O-RAN, has become one of the most significant trends in the telecommunications industry. With its immense potential, professionals across the networking world are eagerly exploring, testing, and implementing O-RAN solutions.

One of the most crucial components of an O-RAN architecture is its fronthaul network, which plays a vital role in the overall system. Ensuring accurate and reliable time synchronization in this area is of paramount importance.

This has prompted us to investigate and conduct various scenarios of time synchronization within the fronthaul network to verify its performance and reliability.

For this combination, we tried to emulate the O-RAN Fronthaul LLS-C2 (Option-A), with one difference, which is the absence of the Distributed Unit. We connected a Grandmaster, a Boundary Clock, and two different timing paths starting from the Boundary Clock.

Each timing path had one Hub-Site Router (HSR) and one Cell-Site Router (CSR), Both CSR routers were connected to the Calnex Paragon-neo analyzer in order to measure the relative PTP time error, and the 1PPS absolute error. The test was passed all the measurement requirements stated by the O-RAN Alliance in the document O-RAN.WG9.XTRP-TST-v02.00 for FR2.

It is important to state that the O-DU would have increased the time error budget, but with the great results we have achieved in this test, the O-DU time error budget will not affect the test results.

The devices participated in the test are shown in Table 33. We performed this test to emulate LLS-C3 scenario, as per O-RAN.WG9.XTRP-SYN-v03.00 document, where the GM in Midhaul.

The topology consisted of:

- Grandmaster (GM): was placed in the Midhaul and connected to the Hub-Site Router.
- Hub-Site Router (HSR)
- Cell-Site Router (CSR)
- Emulated Open Ran Central Unit (Emulated O-CU): Connected to HSR
- Emulated Open Ran Distributed Unit (Emulated O-DU): Connected to CSR
- Emulated Open Ran Radio Unit (Emulated O-RU): Connected to CSR

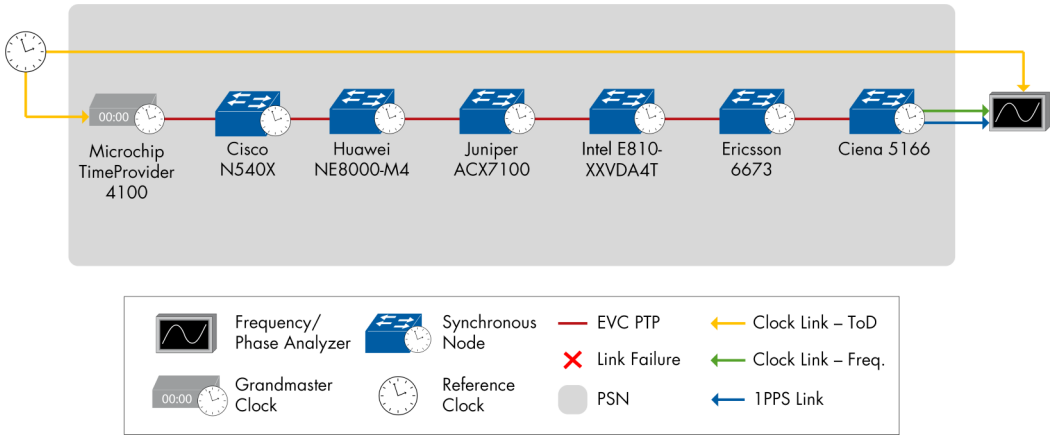

Figure 94: Phase/Time Holdover with Enhanced Sync-E Support

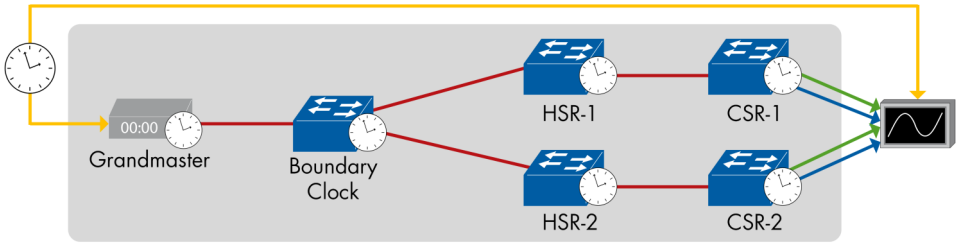

Figure 95: O-RAN Fronthaul LLS-C2 Topology

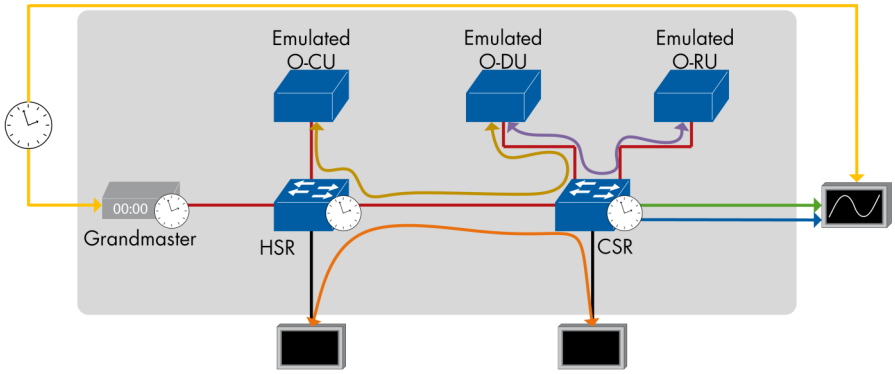

Figure 96: Timing Solution by C3 Configuration with GM from Midhaul

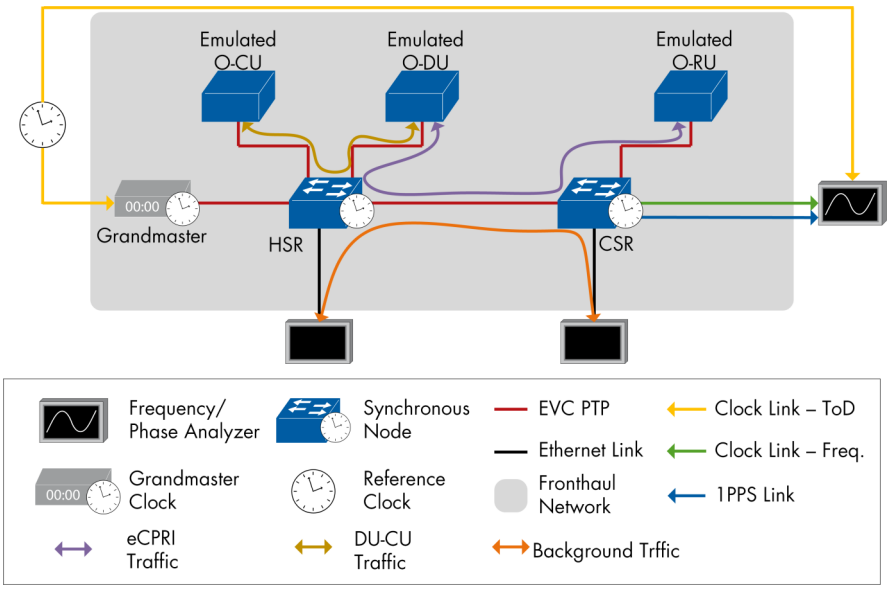

Figure 97: Timing Solution by C3 Configuration with GM from Fronthaul

GM	BC	HSR-1	CSR-1	HSR-2	CSR-2
Microchip	Intel	Juniper	Ericsson	Cisco NCS	Ciena

Table 32: O-RAN Fronthaul LLS-C2 Topology

Grandmaster	HSR	CSR	Emulated O-CU	Emulated O-DU	Emulated O-RU
Microchip TimeProvider 4100	Juniper ACX7100-32C	Juniper ACX7024	Keysight IxNetwork	Keysight IxNetwork	Keysight IxNetwork

Table 33: O-RAN Fronthaul Network Time Synchronization

The last setup we have tested for O-RAN Fronthaul time synchronization was the LLC-C3 configuration with the GM from Fronthaul.

- Grandmaster (GM): was placed in the Fronthaul and connected to the Hub-Site Router.
- Hub-Site Router (HSR)
- Cell-Site Router (CSR)
- O-CU: Connected to HSR
- O-DU: Connected to HSR
- O-RU: Connected to CSR

The Time Error measurements were done on the output of the CSR node. Keysight IxNetwork was used to simulate Midhaul traffic between O-CU and O-DU and ORAN Fronthaul eCPRI traffic between O-DU and O-RU. O-DU and O-RU were configured to simulate ORAN WG4 CU-plane eCPRI traffic for FDD use case with 100 MHz carrier bandwidth in both downlink and uplink direction, 30 KHz sub-carrier-spacing and BFP9 IQ compression. Along with these streams, a background traffic stream is also sent between CSR and HSR to emulate regular traffic in the network.

Class of service was configured on HSR as well as on CSR and the following traffic pattern was applied in the test network:

- eCPRI (O-RU – O-DU) traffic: goes to Low Latency queue with 1.7Gbps traffic load.
- PTP packets: goes to Network Control queue by default.
- O-DU – O-CU traffic: goes to a queue with medium priority with 1.2Gbps traffic load.
- Bidirectional background traffic between HSR and CSR: goes to Best Effort queue with lowest priority with 9Gbps traffic load.

All the time error measurements passed the testing requirements from O-RAN WG9 documents. No loss was reported in O-RAN Fronthaul traffic and also latency variations were less than a few nanoseconds even with the heavy presence of background traffic.

The relative time error was not measured, as only one timing path was tested in the topology. For the both LLS-C3 setups, the following devices participated successfully.

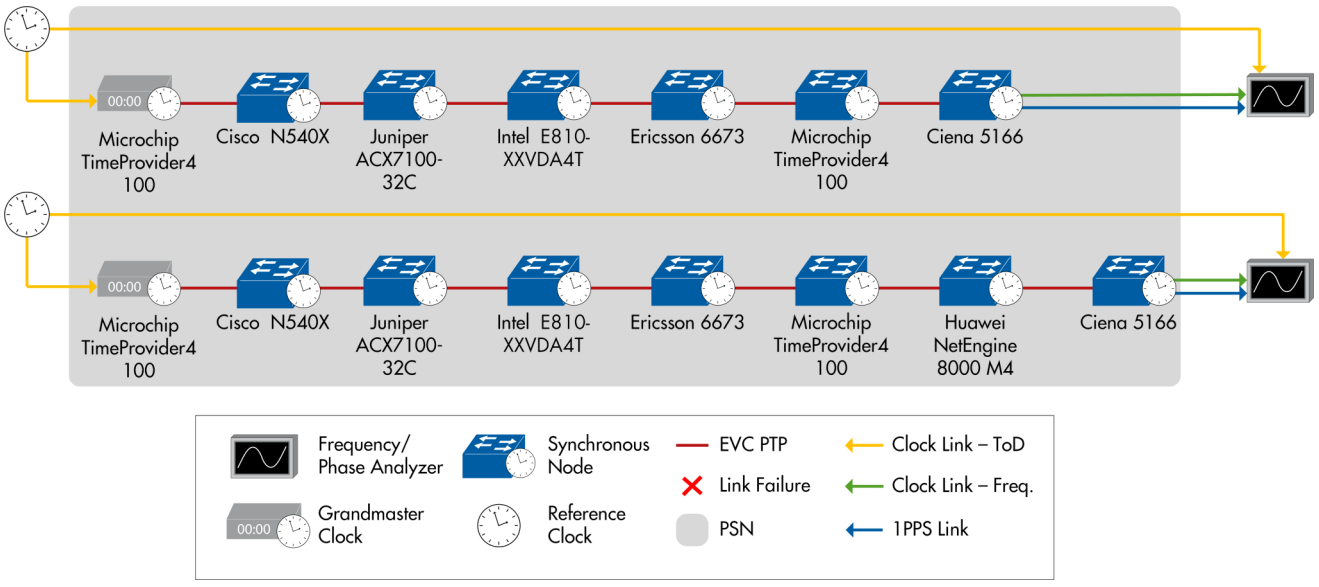

Figure 98: Calculating Time Error Limits for Boundary Clocks

GM	BC 1	BC 2	BC 3	BC 4	BC 5	BC 6	BC 7
Microchip TimeProvider 4100	Cisco NCS 540X-16Z4G8Q2C	Juniper ACX7100-32C	Intel E810-XXVDA4T	Ericsson 6673	Microchip TimeProvider 4100	Ciena 5166	-
Microchip TimeProvider 4100	Cisco NCS 540X-16Z4G8Q2C	Juniper ACX7100-32C	Intel E810-XXVDA4T	Ericsson 6673	Microchip TimeProvider 4100	Huawei NetEngine 8000 M4	Ciena 5166

Table 34: Calculating Time Error Limits for Boundary Clocks

Calculating Time Error Limits for Boundary Clocks

In real-life implementations, the time synchronization chains consist usually of multiple Boundary Clocks, which makes it crucial to measure the time error limits at the end of the chain. The Appendix V of the document G.8273.2 Performance estimation for chain of Boundary Clocks, specifies details for calculating limits for chains of Boundary Clocks.

We performed this test by creating a chain of multiple boundary clocks, connected to a grandmaster.

We measured the constant time, Dynamic time error – low pass filtered, Dynamic time error – high pass filtered, and maximum absolute time error.

For the Constant Time Error limit, as per the ITU-T recommendation, we used the accumulative value depending on the number of boundary clocks in the chain.

For other values we used the recommended formula √ (N x 2) where N is the number of the boundary clock in the chain.

PTP over MACsec

Security of the networks is a fundamental and crucial aspect, in the modern world where the cyber threats are forming a large chunk of the industry. PTP security is left a bit behind, but slowly trying to keep up with the current level of threats.

Encapsulating PTP packets with Layer 2 or Layer 3 encryption is one way of protecting the time synchronization network from being compromised, but the special nature of the PTP and the Hardware Time Stamping, this task is very difficult to implement.

These complications prevented the industry from having a standard interoperable solution between multiple vendors yet.

This did not stop us from testing PTP over MACsec between a boundary clock, and a slave clock from the same vendor—in this test Juniper—using the 8275.1 PTP hybrid profile (SyncE and PTP) and comparing the generated absolute 1PPS time error, when MACsec enabled and disabled.

We used Microchip TimeProvider 4100 as a grandmaster for this setup, and Calnex Sentry for measurement.

In all test steps the output of the slave clock preserved an absolute 1PPS time error less than 5ns.

The following devices passed the test successfully:

GM	BC 1	SC
Microchip TimeProvider 4100	Juniper ACX7100-32C	Juniper ACX7100-32C

Table 35: PTP over MACsec

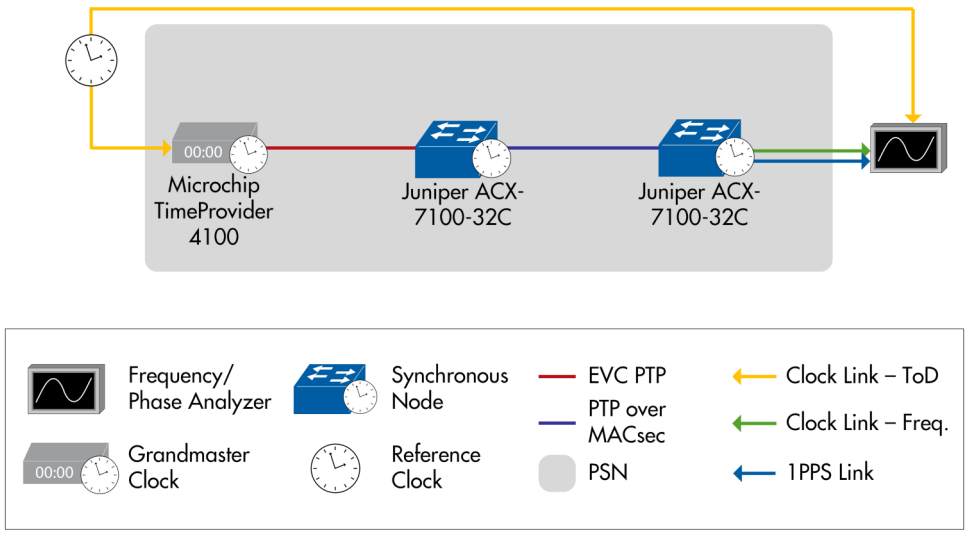

Figure 99: PTP over MACsec

Conclusion

In conclusion, advancements in networking technologies such as EVPN, SRv6, SR-MPLS, BGP Classful Transport Planes, OSPF segment routing, SDN, and time synchronization have significantly improved the efficiency, flexibility, and scalability of modern networks. The EANTC MPLS SDN Interoperability Test event showcased the successful implementation and interoperability of these technologies in multi-vendor environments, covering various services and use cases.

The tests focused on addressing the increasing demands of data centers, 5G networks, and multi-domain service provider environments. Notably, this year's event featured the first implementation of uSID in SRv6, Flex-Algo and FAPM in OSPF segment routing, and the use of virtualized devices for time synchronization. By continually pushing the boundaries of networking technology, these innovations promise to support the growing needs of our increasingly interconnected world.

EANTC AG, Salzufer 14, 10587 Berlin, Germany
info@eantc.de, www.eantc.de
[v1.2 20230419]