EANTC

MPLS SDN & NETCONF
Multi-Vendor Interoperability Test

TEST REPORT 2022

MPLS SD&AI NET WORLD22

Table of Contents

Editor's Note

It is great to be back in Paris this year! Finally we are able to present the results of our annual multi-vendor interoperability event again, in a real on-site showcase with live demonstrations. During the previous two years, we were only able to conduct staging test events in our lab in Berlin in a hybrid on-site/remote mode without public presentations.

The good news is that the industry has not slowed down innovations and product development during the pandemic. Since 2020 and including this year, we have witnessed strong vendor participation in all technology areas that we covered—from Ethernet VPNs and Segment Routing, to MPLS, VXLAN, and SRv6; through traffic engineering and FlexAlgo; to clock synchronization and network management.

The telecom transport, management, and synchronization technologies included in this interoperability test and demonstration are all maturing. Each of the building blocks has a rich family of standards, increasingly reliable implementations, and accepted use case scenarios. Which aspects are different this year, and what are the driving forces behind the scenes that require technology evolution?

The first one is simple: **Service diversity**. Customers are demanding flexible service support; legacy services are being migrated to SDN; and new use cases specifically in cloud access and mobile networks are evolving.

Secondly, **network automation** is becoming a necessity. Managing complex, advanced services is simply no longer possible without automation on the SDN control and management level. It is too error-prone and time-consuming to provision and manage a modern software-defined network manually. Additionally, it becomes more difficult to hire adequate support staff, adding to the reasons for automated provisioning and operations.

Finally, **5G Standalone** and following 3GPP releases require transport network support of an entirely different use case scenario compared to previous mobile network generations. So far, operators have not been required to deploy traffic engineering yet because 5G Non-Standalone (NSA) used today is only a best-effort data throughput booster. All mobile management processes in 5G NSA are still provided by legacy LTE. Serious 5G Standalone deployments at scale with high-density cell site deployments in urban areas and multiple slices are yet to come. These will require automated, traffic engineered networks from aggregation to edge.

Fortunately, the transport network manufacturers seems to be well prepared for these requirements. It is worrying to me, though, that only a relatively small number of vendors implement the full range of advanced hardware and control plane software in support of complex use case scenarios. Even fewer vendors have demonstrated readiness to support multi-vendor network management in our tests so far. I hope the ecosystem will grow further.

Automated end-to-end multi-vendor network provisioning and maintenance (towards fully autonomous networks) will likely be a key aspect of the future network operator business case. There is more work to be done and EANTC will focus on such aspects in our interoperability events in the next years.

This vendor-neutral white paper contains lots and lots of technical details. Its goal is to educate about the state of the art of transport technologies and to enable independent reproduction of our results. I hope this report will provide the expected insights and will help to accelerate innovative network deployments!

Introduction

The MPLS & SDN World Congress in Paris has always been an anchor point in our industry: Insightful presentations and tutorials, excellent networking opportunities, consistent high quality of the whole event carefully curated by the Upperside organizers.

More than 15 years ago, we have started a series of multi-vendor interoperability test events at the annual Upperside congress, which has evolved from the early times of MPLS to now, always focusing on innovative technologies. (For coverage of older technologies and implementations, just browse our archive at eantc.de).

This time, EANTC started preparing for the interoperability event in autumn 2021; we proposed a range of test areas and seed test cases based. Together with interested vendors we developed a detailed test plan with more than 110 test cases. Then, vendors marked their intentions to support for each of the test cases. With this information, we created a multi-vendor test matrix by January. This matrix was amended with detailed configurations to be prepared by vendors. For two weeks in February, we carried out an intense two-week hot staging event with 60+ engineers from the participating vendors in our Berlin lab. We collected more than 635 results data sets which are compiled in this report.

The tests covered the following technical areas:

- Ethernet VPN (EVPN) multi-vendor interoperability has evolved in advanced use case scenarios. Seven vendors successfully tested integrated routing and bridging (IRB), multi-homing, and inter-subnet multicast, just to name a few. These are relevant for carriers to implement a wide range of flexible Ethernet transport services.
- Path computation, traffic engineering, and traffic policies were the main focus areas of the Segment Routing (SR-MPLS) tests. Nine vendors participated in this test area.
- FlexAlgo as an end-to-end policy definition tool has been included in the scope as well again. An increased number of four router vendors plus two test equipment vendors supported it already. Together, FlexAlgo and SR-Traffic Engineering are gaining practical relevance for 5G Standalone slicing services.
- Segment Routing over IPv6 (SRv6) has slowly but steadily growing vendor support. The tests covered pretty much the same scope as in past years, but were expanded to five router vendors plus two test tool vendors this time.
- Our network management tests of SDN controllers and routers are getting more advanced. We used the Path Computation Element Protocol (PCEP), topology discovery by BGP Link State Protocol (BGP-LS), Egress Peer Engineering (EPE) and other technologies. Meanwhile, test cases are complex end-to-end user stories that each take a considerable effort to design and implement.
- Clock synchronization testing focused on Class C/D boundary clocks and multiple grandmaster clock failover scenarios. It is important for operators that slave clocks can always be synchronized without any single point of failure.
- For the first time, we combined our efforts with another project. EANTC is a partner in the "i14y Lab", together with Deutsche Telekom, Telefónica, Vodafone, Nokia, Capgemini and SMEs. This research project is co-funded by the German Ministry of Economic Affairs. We focus testing of disaggregated 5G Standalone Radio Access Network (RAN) services, applications, and platforms. The project contributed a real end-to-end O-RAN solution which we used to verify Open RAN fronthaul requirements as implemented by fronthaul routers participating in our test event.

Individual test scenarios and results are outlined in the remainder of this document.

Interoperability Test Results

This white paper documents all test results, whether positive or negative. We believe that truthful reporting is the key to innovation. Successful test combinations are reported individually including vendor and device names. Failed test combinations are not mentioned in diagrams; they are referenced anonymously to describe the state of the industry. Our experience shows that participating vendors quickly 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 to test and learn.

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.

Participating Vendors and Devices

Vendor	Device	
Arista	7050X3	7280R
Arrcus	UfiSpace S9600-72XC	
Calnex	Paragon-neo Paragon-t	Paragon-X SNE
Ciena	5164	8112
Cisco	Network Services Orchestrator (NSO)	
Huawei	ATN 910D-A iMaster NCE-IP	NetEngine 8000 M14 NetEngine 8000 X8 NetEngine A821
Juniper	ACX5448-M ACX710 ACX7100-32C ACX7100-48L MX10008 MX204 MX240-MPC10E Paragon Insights Paragon Pathfinder	PTX10001-36MR PTX10004 QFX10002-72Q QFX5110-48S QFX5120-32C QFX5120-48Y QFX5120-48YM QFX5130-32CD
Keysight	IxNetwork	
Microchip	TimeProvider 4100	
Nokia	7750 SR-1	Network Services Platform (NSP)
Ribbon	NPT-1300	
Spirent	SPT-N4U	
ZTE	ZENIC ONE	ZXR10 M6000-8S Plus

Table 1: Participating Vendors and Devices

Orchestrator & Controller/PCE

- Cisco NSO
- Huawei iMaster NCE-IP
- Huawei NCE Path Computing Element
- Nokia NSP
- Juniper Paragon Pathfinder
- ZTE ZENIC ONE

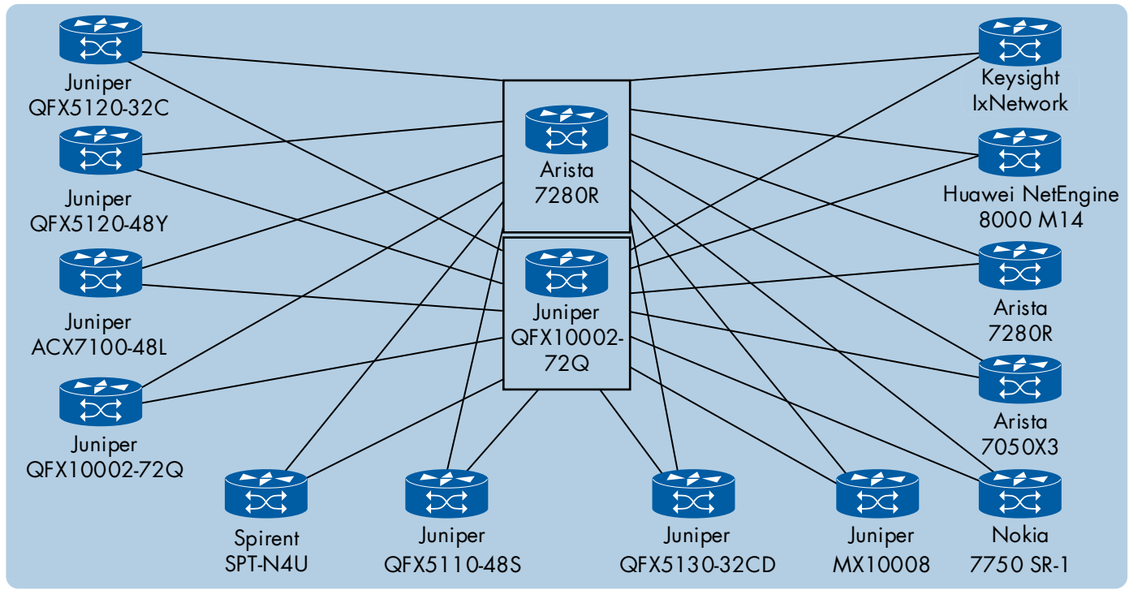

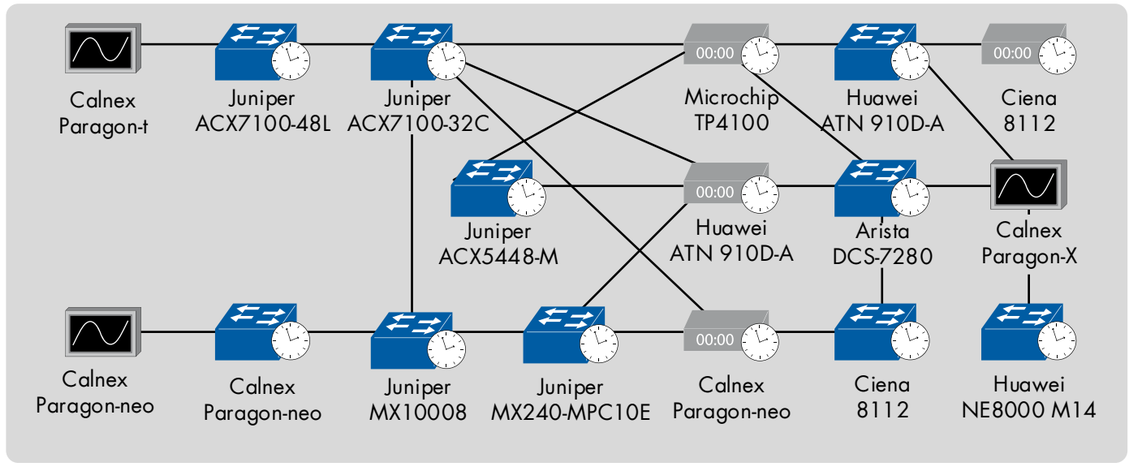

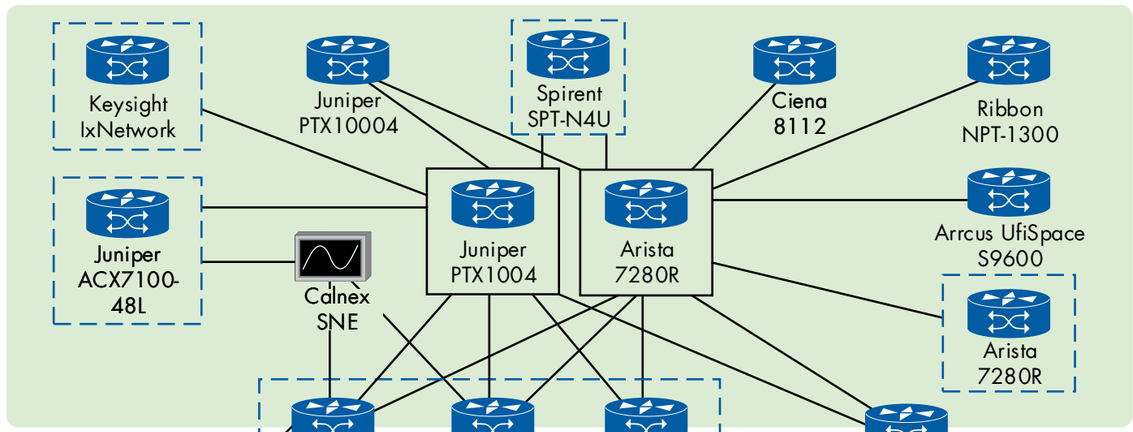

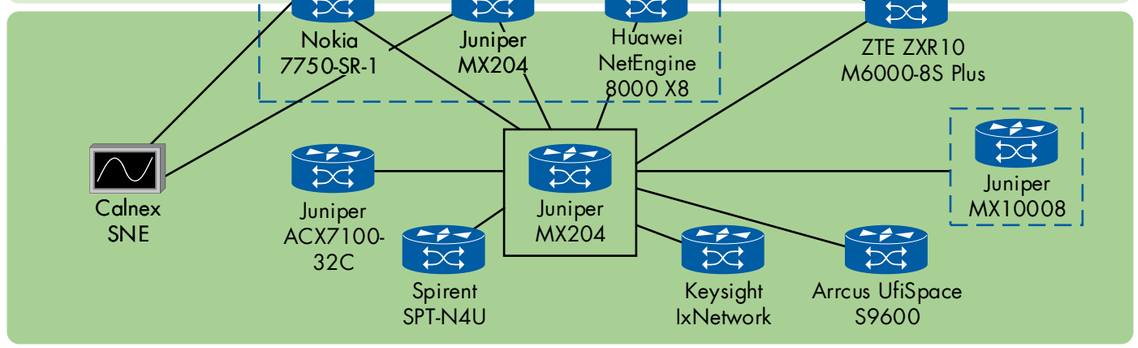

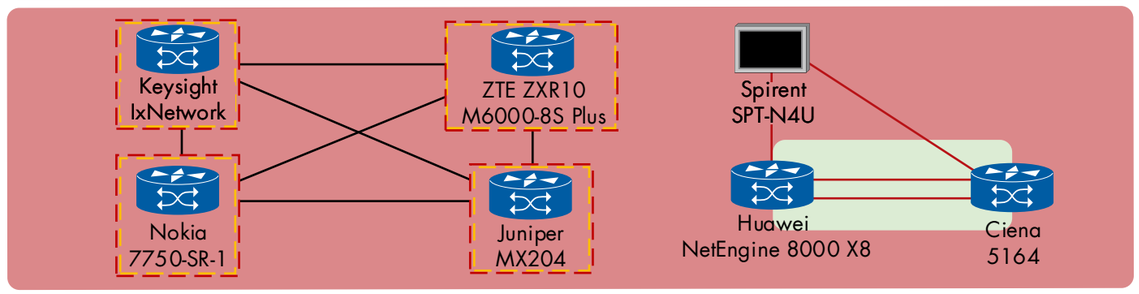

- Core Router
- Grandmaster
- PCE/ Management server
- Flex Algo
- VXLAN
- SR-MPLS
- Emulator
- Boundary Slave Clock
- Route Reflector
- PCEP
- Clocking
- SRv6
- SDN Edge Router
- IP/MPLS
- FlexE Link
- Physical Link

MPLS SDN Interoperability Test

Segment Routing

This year tests covered major aspects of Segment Routing (SR).

The interworking between different domains (SR-MPLS, SRv6 and VXLAN) had great share of focus. We used advanced scenarios with multiple gateways at the same time in aim to reflect real situations and confirm wide network interoperability support. Besides we covered many traffic steering methods in SR-MPLS, great use case for Prefix summarization in SRv6 and also the important tests of resiliency that include TI-LFA (topology-independent loop-free alternate) and S-BFD (Seamless BFD).

SR-MPLS VPN Services and Traffic Steering

SR-MPLS enhances packet forwarding behavior based on application requirements by using source routing and the available extensions of the BGP and IGP protocols. In this section we first verified the SR-MPLS creation and transport for end-to-end VPN services over ISIS/OSPF. Then we explored multiple SR-TE (Segment Routing Traffic Engineering) traffic steering modes (Local on demand, per traffic and per flow steering).

L3VPN over SR-MPLS

We used L3VPN services as a basic test of interoperability over SR MPLS. All participant nodes established ISIS/OSPF sessions with each other over a mesh topology. The routing tables included the loopback addresses and the respective SID of the involved PEs.

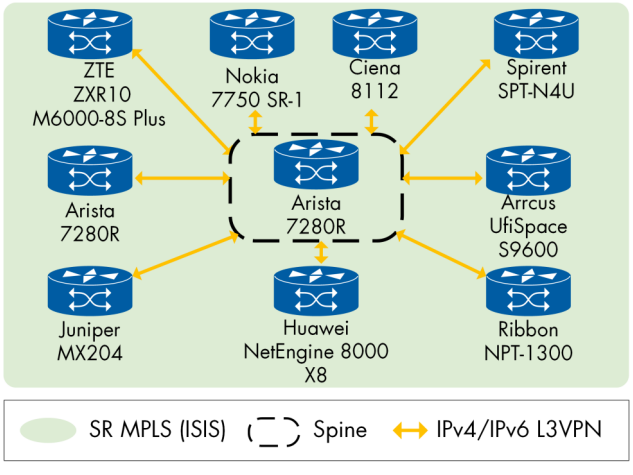

Figure 1: L3VPN over SR-MPLS (ISIS)

Arista 7280R, Arrcus UfiSpace S9600-72XC, Ciena 8112, Huawei NetEngine 8000 X8, Juniper MX204, Nokia 7750 SR-1, Ribbon NPT-1300, Spirent SPT-N4U, and ZTE ZXR10 M6000-8S Plus successfully participated as PEs in the test.

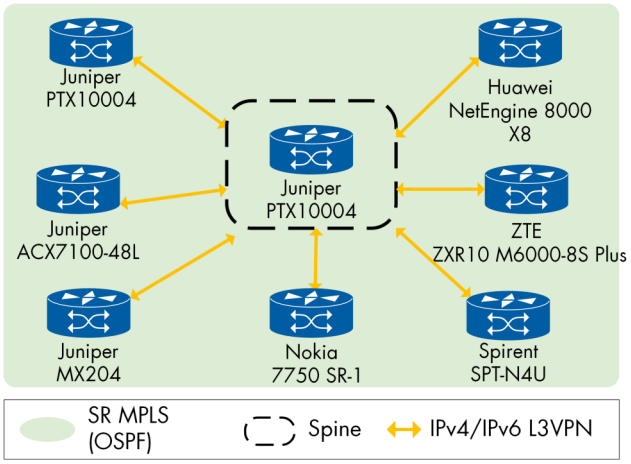

Figure 2: L3VPN over SR-MPLS (OSPF)

Huawei NetEngine 8000 X8, Juniper PTX10004, Juniper ACX7100-48L, Juniper MX204, Nokia 7750 SR-1, Spirent SPT-N4U, and ZTE ZXR10 M6000-8S Plus successfully participated as PEs.

SR-MPLS Dynamic Path Computation on Headend

In next-generation networks, network-performance criteria (e.g., latency) are becoming as critical to data-path selection as other metrics.

As RFC 8570 describes, the extensions to IS-IS traffic engineering provide performance metrics in the underlay network. Theses distributed information can then be used to make path selection decisions based on network performance. We verified path selection for IPv4/IPv6 traffic in IGP TE extension network. TWAMP-based performance measurement provides such performance value to be distributed via ISIS TE extensions, which is the preferred method supported by the test.

The testbed consisted of four nodes all supporting delay measurement. The DUT computed an SR-TE policy based on delay to reach the egress point.

Traffic was generated between the two PEs and it used the path with the least delay. Then using an impairment device we increased the delay on one link and observed the DUT automatically recalculated the SR-TE path to chose the one with better delay and the traffic being re-routed accordingly.

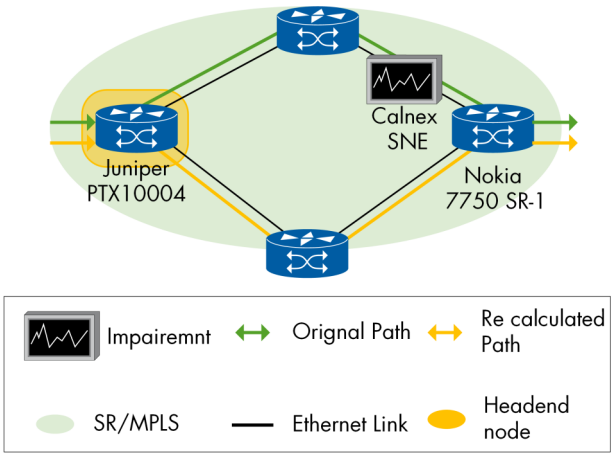

Figure 3: SR Dynamic Path Computation on Headend

The following devices participated successfully as:

DUT: Juniper PTX10004
PE: Nokia 7750 SR-1
Impairment device: Calnex SNE
Traffic Generator: Spirent SPT-N4U

Local On-demand SR MPLS-TE Policy Instantiation

In Segment Routing networks a large amount of configuration is required for setting up the Segment Routing label-switched paths (LSPs). The SR-TE on-demand LSP simplifies provisioning for networks and reduce the amount of configuration in such deployments.

SR provides VPN traffic steering to TE paths, and automatic instantiation of such paths according to the SR policy. The calculation of the path is carried out by the headend in an internal network, which has TE-information exchanged by IGP-TE. At the headend, the instantiation is triggered after the BGP route is received from remote PE.

We verified instantiation of SR-TE policies at the headend, upon receipt of BGP routes associated with the SR-TE policies. Vendors created admin groups to influence path computation and configured SR-TE on-demand policies that match the incoming routes to LSP templates for LSP instantiation. Participating nodes disabled routing-instances or services, forcing a withdrawal of BGP service routes, and it was verified that no SR-TE LSP was instantiated.

Later, we enabled the BGP peering, SR-TE LSPs were created automatically and mapped accordingly to the BGP colors.

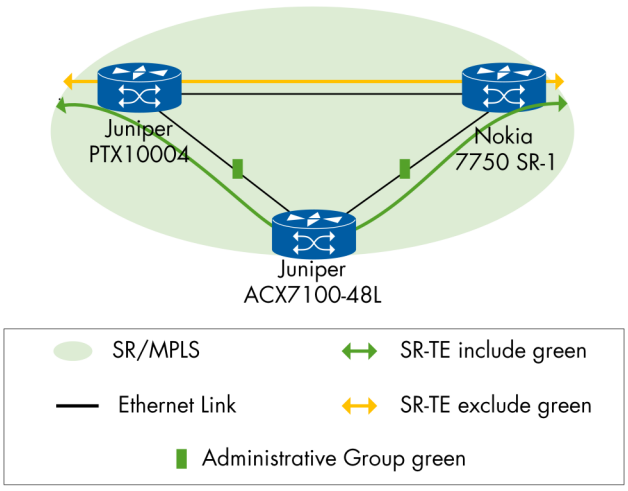

Figure 4: Local on-demand SR MPLS-TE Policy Instantiation

Juniper PTX10004 and Nokia 7750 SR-1 successfully participated as headend node, Juniper ACX7100-48L functioned as transit node.

SR-TE—Per-Destination Traffic Steering

Per-destination automated steering, it automatically steers service route onto SR policy based on color and next-hop address.

We verified per-destination traffic steering over a locally configured SR policy. An SR Policy was configured on the head end with color and an end-point of the policy and next-hop address with explicit path. Traffic was generated between the Egress and Ingress nodes and the flow was steered to the destination according to the policy.

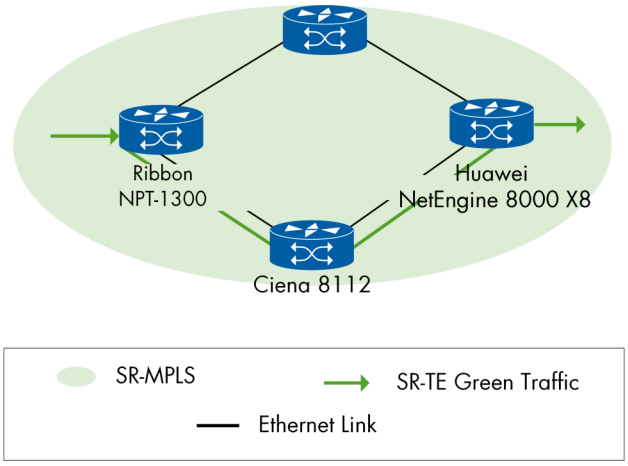

Figure 5: SR-TE—Per-Destination Traffic Steering

In this test, Ribbon NPT-1300 acted as ingress node, Ciena 8112 as transit node and Huawei NetEngine 8000 X8 as egress node.

SR-TE—Per-Flow Traffic Steering

In order to satisfy Service Level Agreements (SLAs), TE is used to assign traffic flows to network paths.

The SR-TE Per-flow policy is one of the mechanisms that allow the steering of traffic on an SR policy based on the attributes of the packets like DSCP value or source address etc. So it basically requires a match criteria which can identify specific flow rather than all traffic.

Two SR-TE policies were configured on the nodes. Traffic with two different DSCPs values was generated between the Ingress PE and the egress PE.

The traffic will be mapped to LSPs according to the BGP color community assigned to the DSCP values.

Traffic was received correctly on the right path with no packet loss.

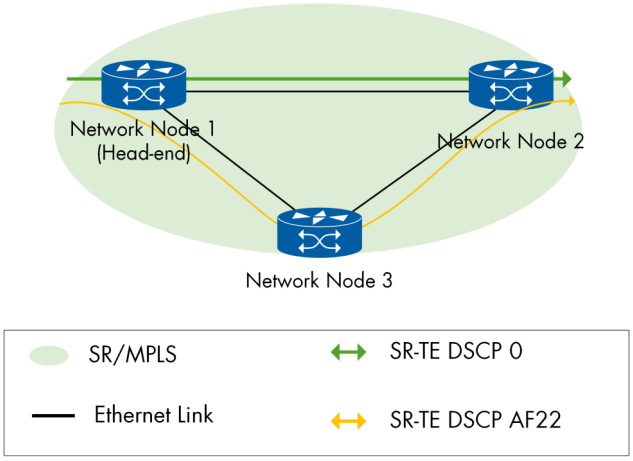

Figure 6: SR-TE—Per-Flow Traffic Steering

SR-MPLS BGP-LU with Prefix SID Redistribution

When the BGP LU is used, service providers who divide their MPLS networks into multiple regions with different IGP instances running within those regions benefit from increased network scale and faster convergence times.

BGP-LU provides connectivity between regions by advertising PE loopbacks and label bindings to the Regional Border Routers (RBR). ABRs then advertise the loopbacks and label bindings to remote PEs in other regions.

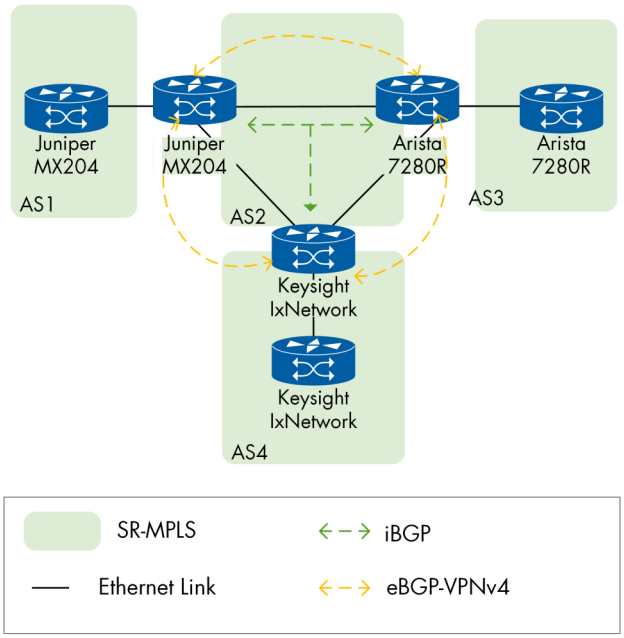

Figure 7: SR-MPLS BGP LU Prefix SID Redistribution

Arista 7280R, Juniper MX204, and Keysight IxNetwork successfully participated as customer edge (CE) and area border router (ABR).

We prepared a topology with four ASs that include three ABRs and three CEs. All the CEs maintained VPNv4 peering with the remote nodes, while the ABRs established BGP LU peering between each other through iBGP.

Network Node 1 (Headend)	Network Node 2	Network Node 3
Juniper MX204	Arista 7280R	Huawei NetEngine 8000 X8
Arista 7280R	Juniper MX204	Huawei NetEngine 8000 X8
Huawei NetEngine 8000 X8	Juniper MX204	Arista 7280R
ZTE ZXR10 M6000-8S Plus	Arista 7280R	Juniper MX204

Table 2: SR-TE - Per-Flow Traffic Steering

Each autonomous system boundary router (ASBR) is redistributing from ISIS into BGP-LU, and from BGP-LU into ISIS. BGP-LU is carrying prefix-SID and SRGB information, thus allowing remote ASBR and PE to use globally significant MPLS labels. Through the FIB of the CEs we confirmed that the LU tunnel is being received as SR tunnel between each remote CE and they can see their VPN routes as well thanks to the prefix redistribution on the respective ABR.

400GE ZR

400G pluggable modules represent an architectural change in high-bandwidth data center interconnects because they can be plugged directly into switches and routers offering the same density for both coherent DWDM and client optics in the same chassis. This architectural change helps network operators support their growing bandwidth demands in a more cost-efficient manner.

We had SR-MPLS control plane with L3VPN service between two locally connected optics.

200Gb/s bidirectional traffic was monitored across the 400ZR link with no issues.

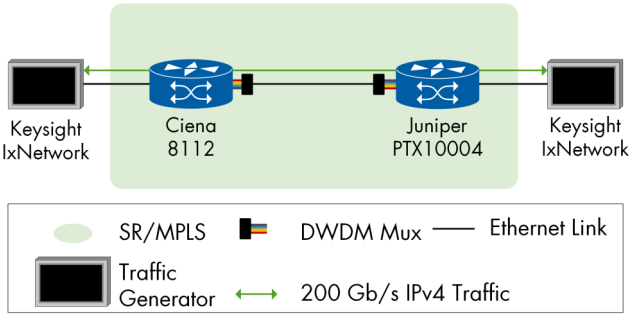

Figure 8: 400GE ZR Interoperability

Ciena 8112 and Juniper PTX10004 were tested as PEs and Keysight IxNetwork functioned as measurement equipment.

Segment Routing Anycast

Anycast-SID in Segment Routing plays a vital role in achieving node resiliency, traffic load-sharing or even creating separate network planes for different types

Anycast-SID is a Node Prefix-SID that is advertised by more than one node (typically two). The set of nodes advertising the same anycast-SID form a group called an anycast set. Using an Anycast-SID in the SID list of an SR policy path provides improved traffic load-sharing and resiliency.

The network architecture was made of four devices, with two Anycast-SIDs. We confirmed that the anycast set can advertise the Anycast-SID. Through learning this SID, an SR policy is configured on the headend node to steer the traffic and select and include the Anycast SID into the segment list. Then traffic can be load-balanced to reach the remote end using the Anycast set as next-hops.

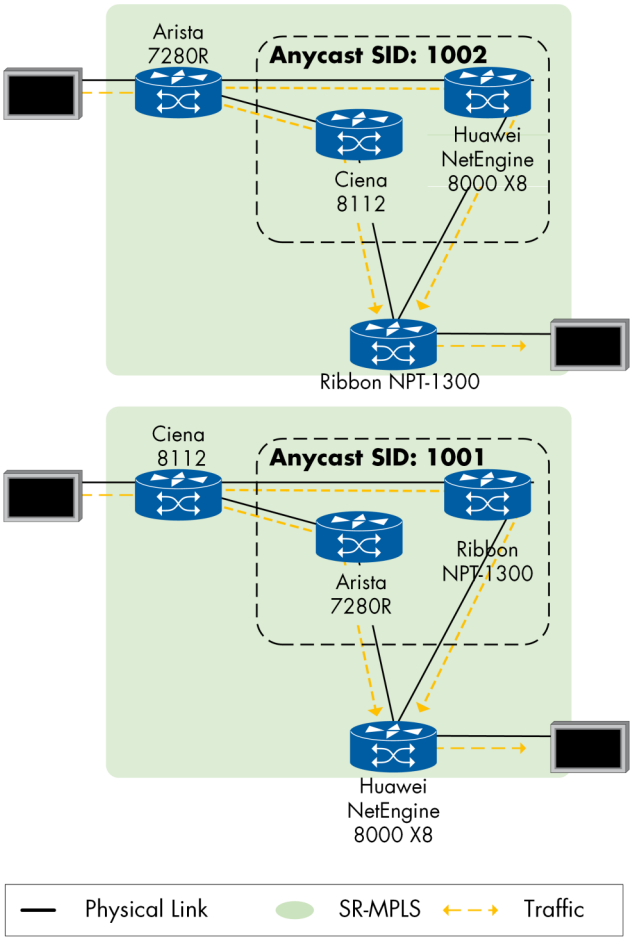

Figure 9: SR MPLS Anycast SID

Arista 7280R, Ciena 8112, Huawei NetEngine 8000 X8, and Ribbon NPT-1300 successfully participated in this test.

Segment Routing LSP Ping/Traceroute

Ping and traceroute operations can be used to check the connectivity of label distribution protocol (LDP) label switched path (LSPs) that carry IPv4 or IPv6 packets also can locate the fault point on the path.

We verified ping and trace route for SR-MPLS.

One test pair showed malformed TLV types, failing to send either an echo request or an echo response. We excluded this pair from the test.

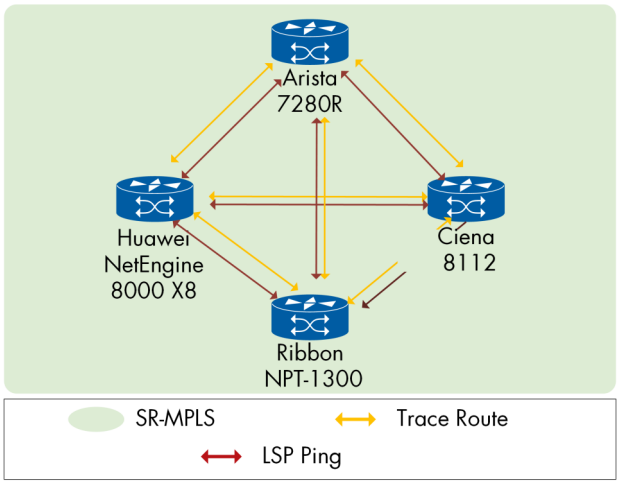

Figure 10: Segment Routing LSP Ping/Traceroute

The participating devices were Arista 7280R, Ciena 8112, Huawei NetEngine 8000 X8, and Ribbon NPT-1300.

SRv6 VPN Services

SRv6 simplifies the network even further by relying on the native IPv6 header and header extension to provide the same services and flexibility as SR-MPLS directly over the IPv6 data plane.

The following tests verified functionality of various VPN services over SRv6.

L3VPN over SRv6

As defined in "SRv6 BGP based Overlay services", draft-ietf-bess-srv6-services, we used BGP as control plane and SRv6 as data plane to build up L3VPN services between the PEs.

BGP advertises the reachability of prefixes of a specific service from an egress PE node to ingress PE. BGP messages exchanged among PEs deliver SRv6 service SIDs, which BGP makes use of to interconnect PE devices to shape VPN sessions.

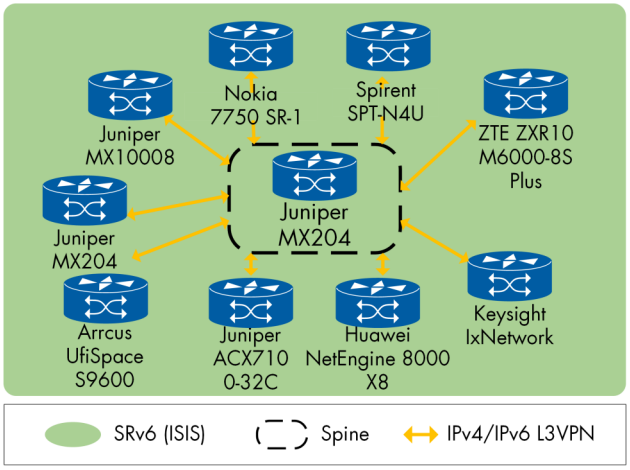

Figure 11: L3VPN over SRv6 (ISIS)

PEs in this tests were Arrcus UfiSpace S9600-72XC, Huawei NetEngine 8000 X8, Juniper ACX7100-32C, Juniper MX10008, Juniper MX204, Keysight IxNetwork, Nokia 7750 SR-1, Spirent SPT-N4U, and ZTE ZXR10 M6000-8S Plus.

EVPN VPWS over SRv6

This test verified the interoperability of VPWS over the SRv6 network with multi-homing scenario.

All-Active Multi-Homing enables an operator to connect a CE device to two or more provider edge (PE) devices to provide load balancing and redundant connectivity. With All-Active Multi-Homing, all the PEs can forward traffic to and from the multi-homed device.

We verified that the CE node which was connected to two PEs through ethernet links and all the multi-homed PEs forwarded traffic to/from that Ethernet segment for a given VLAN.

The traffic flow was load-balanced to both PE1 and PE2 and received with no traffic loss.

As Link Aggregation Control Protocol (LACP) was configured on the multihomed CE, we emulated a failure on one of the links and observed traffic continuing through the second PE with no packet loss.

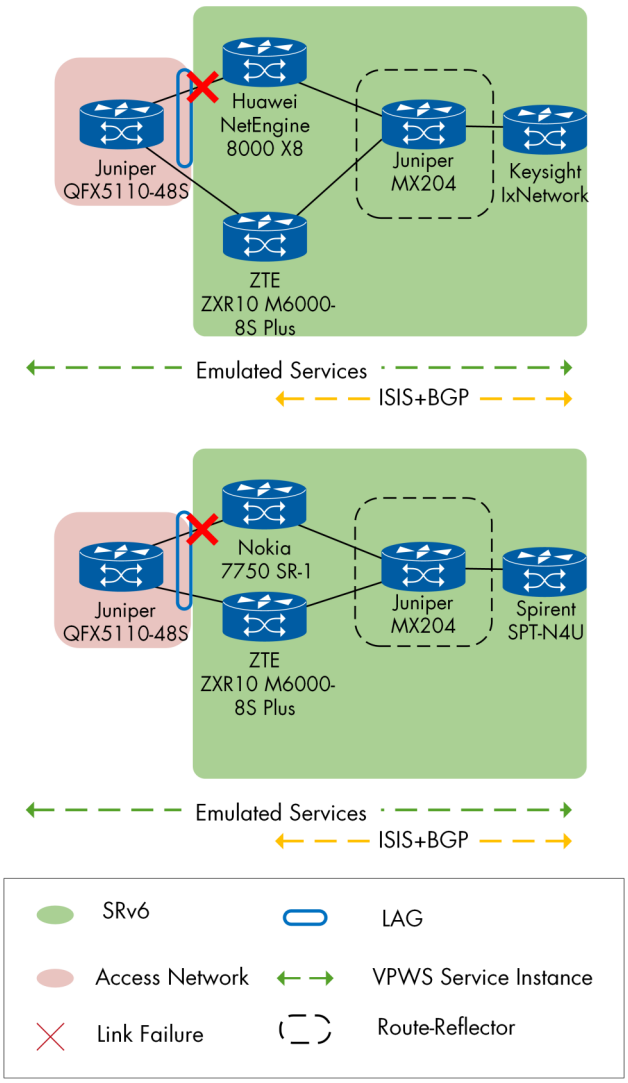

Figure 12: EVPN VPWS Multi-homing Service over SRv6

For EVPN VPWS the participant devices were Huawei NetEngine 8000 X8, Nokia 7750 SR-1, and ZTE ZXR10 M6000-8S Plus as multi-homed PEs and Keysight IxNetwork and Spirent SPT-N4U functioned as single-homed PEs. Juniper QFX5110-48S acted as CE.

BGP IPv4/IPv6 Global Routing Table over SRv6

BGP is used to advertise the reachability of prefixes of a particular service from an egress PE to ingress PE nodes. The BGP messages exchanged between PE devices carry SRv6 service SIDs.

We verified the advertising IPv4 and IPv6 prefixes over BGP peerings using IPv6 transport with SRv6 information in the Prefix-SID attribute for forwarding over an SRv6 data plane. We observed advertisements of VPN route prefixes with END.DT4, End.DT6 and END.DT46 behaviors (identifier of the endpoint with decapsulation and specific IPv6, IPv4, or IP table lookup) via IPv6 BGP over the SRv6 network. All traffic was received for the advertised routes as expected.

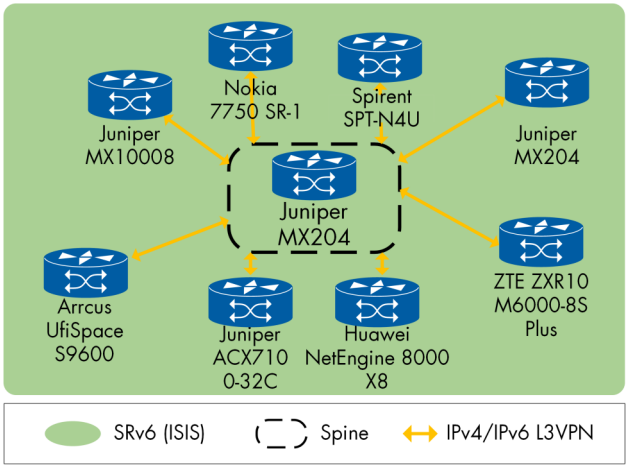

Figure 13: BGP IPv4/IPv6 Global Routing Table over SRv6 (END.DT4,END.DT6)

Arrcus UfiSpace S9600-72XC, Huawei NetEngine 8000 X8, Juniper ACX7100-32C, Juniper MX10008, Juniper MX204, Nokia 7750 SR-1, Spirent SPT-N4U, ZTE ZXR10 M6000-8S Plus participated in the test.

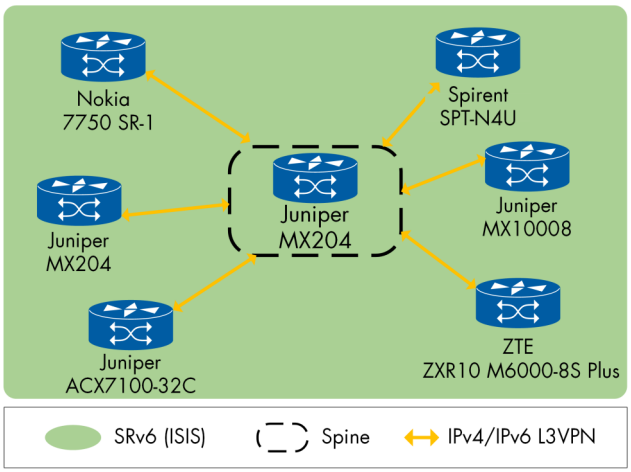

Figure 14: BGP IPv4/IPv6 Global Routing Table over SRv6 (END.DT46)

The participating devices were Huawei NetEngine 8000 X8, Juniper ACX7100-32C, Juniper MX10008, Juniper MX204, Nokia 7750 SR-1, Spirent SPT-N4U, and ZTE ZXR10 M6000-8S Plus.

One device had an issue with installing the route into the RIB/FIB when the SID structure TLV contained non-zero values. But later that was fixed by adjusting the code to handle that case and all routes were installed.

SRv6 Locator Summarization

In MPLS networks, the lack of prefix summarization led to very complex inter-AS options and hierarchical BGP labels. Thanks to SRv6 prefix summarization, SRv6 gets rid of all of these complexities and achieves massive-scale reachability.

From point of view of small access devices, the service providers prefer to ask for a small amount of route summarization rather than thousands of routes. So we chose to perform the summarization per-Algo including base Algo 0.

The devices deployed two Flex Algos FA0 and FA133. We confirmed that the ABRs are advertising the prefixes summary as SRv6 locators (required by the standards) as well IPV6 Reachability (so all kinds of devices whether they support SRv6 or not can route using the prefixes) as expected.

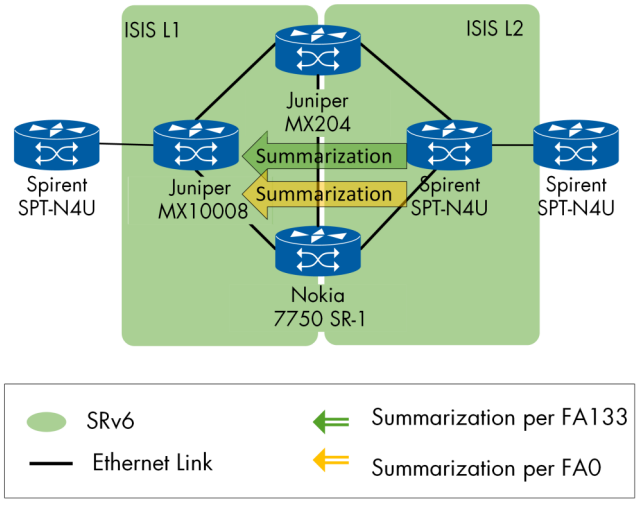

Figure 15: SRv6 Locator Summarization

Juniper MX204 and Nokia 7750 SR-1 successfully participated as ABRs, Juniper MX10008 and Spirent SPT-N4U were tested as PEs.

Domains Interworking

SRv6 and MPLS Service Interworking

The SRv6/MPLS L3 Service Interworking Gateway supports both transport and service termination at the border node. For all prefixes in the VRF set for re-origination, the gateway creates both SRv6 VPN SIDs and MPLS VPN labels. By popping the MPLS VPN label, checking up the target prefix, and pushing the appropriate SRv6 encapsulation, the gateway facilitates traffic forwarding from the MPLS domain to the SRv6 domain. The gateway removes the outer IPv6 header, looks for the target prefix, and pushes the VPN and next-hop MPLS labels from the SRv6 domain to the MPLS domain.

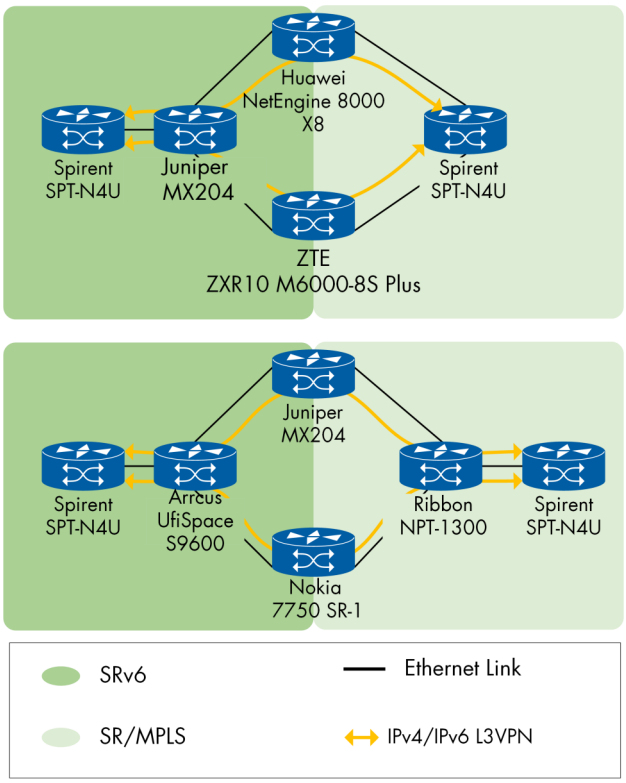

Figure 16: SRv6 and MPLS Service Interworking

These devices participated successfully as:

Interworking Gateways: Huawei NetEngine 8000 X8, Juniper MX204, Nokia 7750 SR-1, and ZTE ZXR10 M6000-8S Plus

PE: Arrcus UfiSpace S9600-72XC, Juniper MX204, Ribbon NPT-1300, Spirent SPT-N4U

This functionality allows SRv6 L3VPN domains to communicate with MPLS L3VPN domains that already exist. And also makes it possible to transition from MPLS L3VPN to SRv6 L3VPN. We deployed a testbed consisting of two gateways between the SR-MPLS/SRv6 domains. The interworking devices needed to Import service routes received from one domain (MPLS or SRv6), then re-advertised the exported service routes to the other domain (next-hop-self) and stitch the service on the data plane.

With Nokia and Juniper combination as interworking gateways a domain id for each service instance was also configured and advertised in the d-path to avoid loops.

IPVPN over SRv6/EVPN RT5 over SR-MPLS Interworking

One of the encountered situations while integrating different technologies, is having different services of advertising prefixes between the access and the core.

In this test we verified the interworking between IPVPN over SRv6 and EVPN L3VPN over SR-MPLS. This interworking relies on the ability of the gateways between the domains to receive IPVPN/EVPN prefixes and then readvertised them to the other side.

Pure IP traffic was injected into the PEs and observed on the gateways BGP IPVPN routes and EVPN RT5 routes sent and received from the PEs.

Traffic was received with no packet loss. Also a D-Path attribute was used in this scenario to help prevent control plane loops. The interworking PE flags route as a loop, and does not re-advertise it to the neighbors since its D-PATH contains the gateway's local domain.

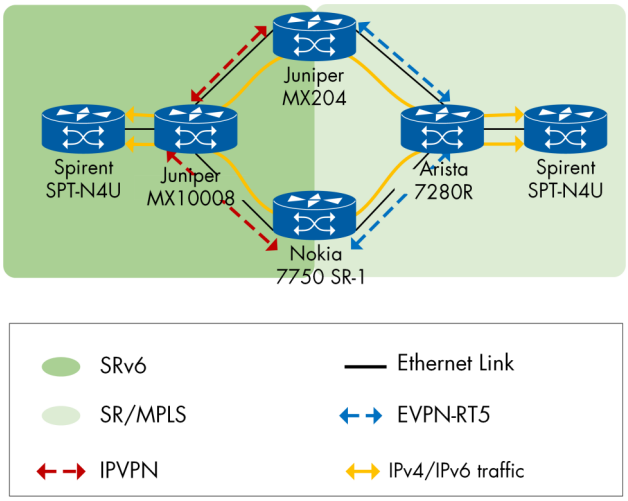

Figure 17: IPVPN over SRv6/EVPN RT5 over SR-MPLS Interworking

Juniper MX204 and Nokia 7750 SR-1 successfully participated in the test as GW, Arista 7280R, and Juniper MX10008 as PEs, and Spirent SPT-N4U as CE.

EVPN over VXLAN and EVPN over SRv6 Interworking

Operators encounter the interoperability of VXLAN when introducing SRv6 transport for end-to-end network services.

A gateway can stitch the EVPN at the border of both networks with each other. Based on the support of control plane, generating EVPN prefix (with VNI and next-hop) from VXLAN, as well as EVPN prefix with SID from SRv6. The gateway shall also support encapsulation of EVPN packets into VXLAN packets and translate them to SRv6 packets (vice versa).

We confirmed that the gateway is carrying out the routes from one domain to the other and observed no packet loss in the end to end traffic.

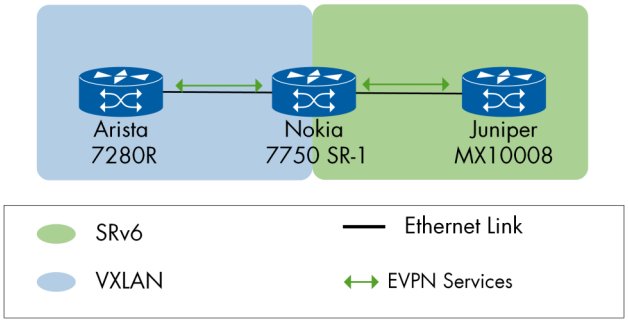

Figure 18: EVPN over VXLAN and EVPN over SRv6 Interworking

Nokia 7750 SR-1 successfully participated as Gateway router and Arista 7280R and Juniper MX10008 were tested as PEs in this test.

TI-LFA over SR-MPLS

In these tests, we verified the TI-LFA local SRLG (Shared Risk Link Group) and measured the convergence time after a link failure. We also ran tests to confirm the remote micro-loop avoidance feature. Micro Loop happens when different nodes in the network have different convergence times, and when loop duration is longer than their TTL, it may cause traffic loss.

TI-LFA Local SRLG and Link Protection

We built a topology consisting of four nodes to test link and SRLG TI-LFA over the SR-MPLS network. The participated vendors configured the network nodes with an L3VPN service.

Prior to the link failure, the ingress PE (PLR) forwarded the traffic to the directly connected egress PE. To simulate the link failure, we asked the vendor of egress PE to disconnect the link between egress and ingress nodes (the protected link), simultaneously the traffic was still flowing from the traffic generator toward the ingress PE.

We observed in three of the cases the out-of-service time between 7 ms and 26 ms. In one of the combinations, the out-of-service time was 106 ms above the expected results.

For local SRLG, the PLR nodes used a port to repair the link fault, regardless of the cost, because it shares the same SRLG of the failed port. The failover time was between 13 ms and 32 ms for the two combinations we performed.

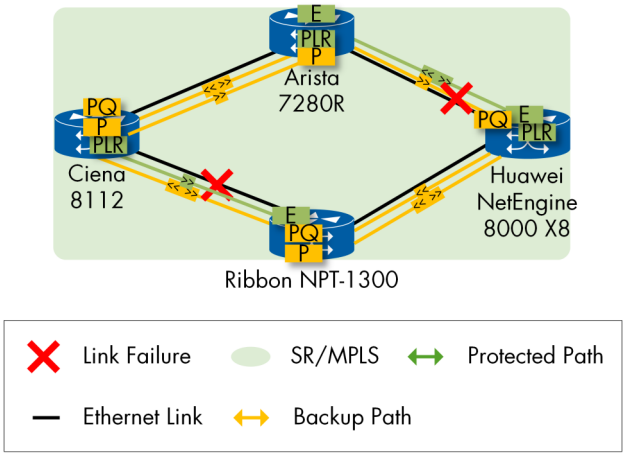

Figure 19: TI-LFA over SR MPLS Link Protection

Arista 7280R, Ciena 8112, Ribbon NPT-1300, and Huawei NetEngine 8000 X8 participated in multi-vendor combinations in this test case as egress nodes, Arista 7280R, Ciena 8112, Huawei NetEngine 8000 X8, and Ribbon NPT-1300 as P nodes, Arista 7280R, Ciena 8112, and Huawei NetEngine 8000 X8 as PLR nodes and Arista 7280R and Huawei NetEngine 8000 X8 as PQ nodes.

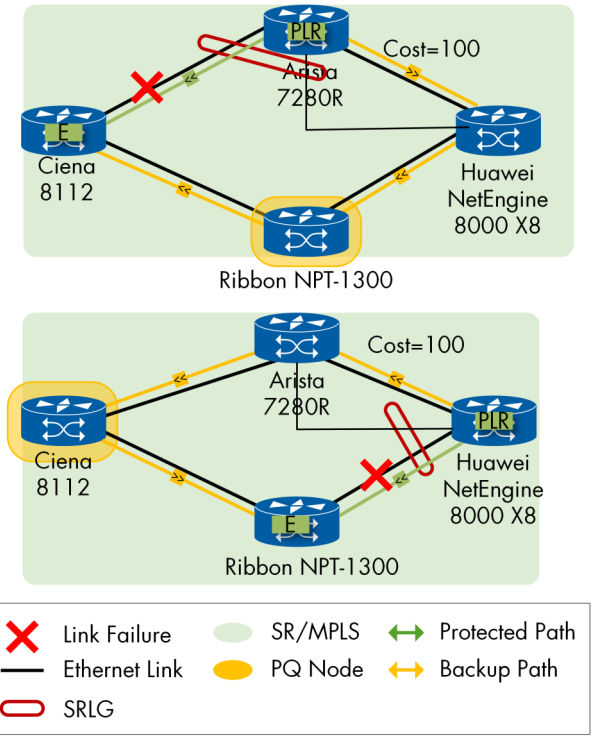

Figure 20: TI-LFA over SR MPLS with Local SRLG Protection

Arista 7280R and Huawei NetEngine 8000 X8 successfully participated as PLR nodes, Ciena 8112 and Ribbon NPT-1300 functioned as egress nodes in the test.

TI-LFA with Remote Micro Loop Avoidance

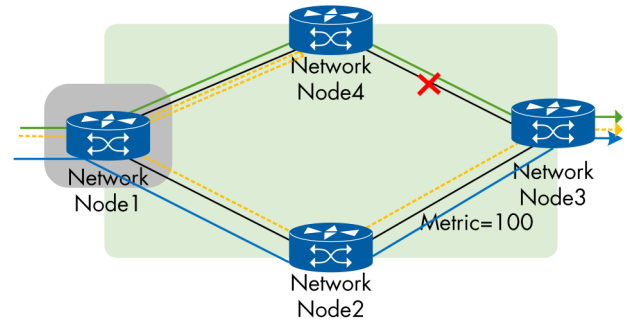

Figure 21: TI-LFA with Micro Loop Avoidance—Link Failure

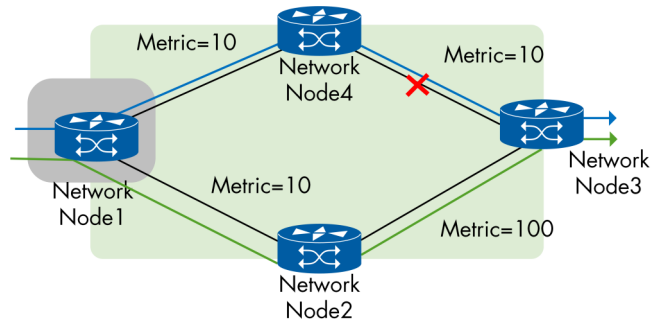

Figure 22: TI-LFA with Micro Loop Avoidance—Link Restoration

IETF draft-bashandy-rtgwg-segment-routing-uloop states that forwarding loops happen during the convergence of the IGP, as a result of transient inconsistency among forwarding states of the nodes of the network. Micro loop avoidance feature will instantiate a TI-LFA computed tunnel which enforces an explicit path without creating any state along the desired path. This repair tunnel is created on Node1 when a network failure or restoration results in a change of next-hop for a given prefix, and remains programmed in Node1 for a time determined by a configured FIB delay Timer C. Using four nodes topology, we configured the FIB (forwarding information base) download delay timer to a high value—30 seconds—so that we could determine that the micro-loop repair tunnel was programmed in network node 1 after both the failure and restoration of the link between Node4 and Node3. Before failing the link the NetworkNode1 (DUT) was pushing the Destination SID.

Then link was disconnected and DUT pushed [AdjSID (NetworkNode2->destination)]. As a result, traffic follows the desired path, regardless of the forwarding state for the destination at NetworkNode2. The same behavior was expected with the link restoration and we confirmed that as well.

The following devices participated successfully as DUT1: Juniper MX204, Juniper PTX10004, and Nokia 7750 SR-1.

TI-LFA over SRv6

The next step is to verify the TI-LFA with the SRv6 data plane. Square topology was implemented and all vendors took turns in participating as PLR nodes. Out-of-service times were between 2 ms and 27 ms. One combination had a failure time of 95 ms which was not included in the report.

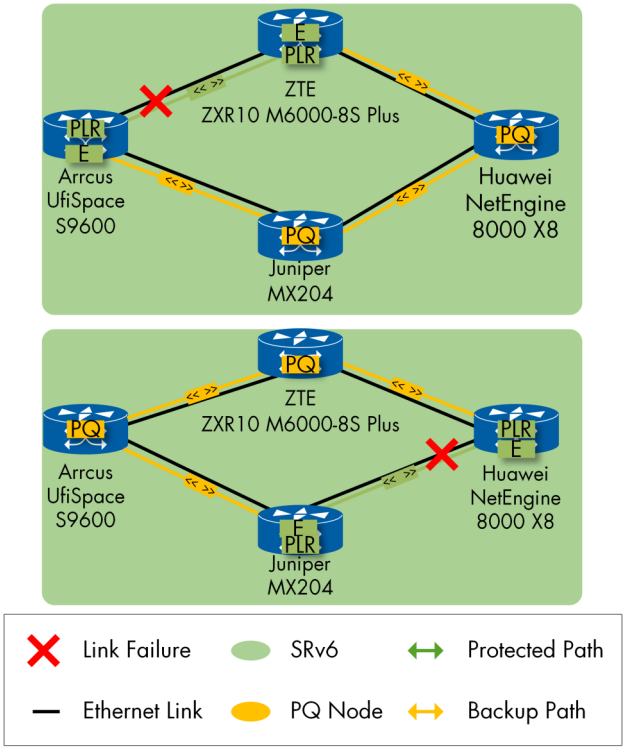

Figure 23: TI-LFA over SRv6

The following devices successfully participated in the test as: Arrcus UfiSpace S9600-72XC, Huawei NetEngine 8000 X8, Juniper MX204, ZTE ZXR10 M6000-8S Plus, and Spirent SPT-N4U as measurement device.

Seamless BFD over SR-MPLS

Seamless BFD (S-BFD) is a simplified mechanism for using BFD with a large proportion of negotiation aspects eliminated, thus providing benefits such as quick provisioning, as well as improved control and flexibility for network nodes initiating path monitoring.

In this test, we verified that the SR policy can be used to steer traffic into the SR-TE tunnel, and the Seamless BFD can detect the link failure and trigger SR-TE hot standby protection.

We built a triangle topology with egress PE, ingress PE, and one P router. We asked each pair of PEs to configure two SR-TE policies; one was the primary path and the other one was the backup.

The initiator interval was configured 50 ms making the acceptable out of service time between 100-170 ms.

We started generating traffic between Initiator and Reflector through the P node (the longer and primary path). And to display the S-BFD role in network convergence we emulated a tear-down session by shutting down a remote port and we observed the traffic switching to the backup SR MPLS TE backup.

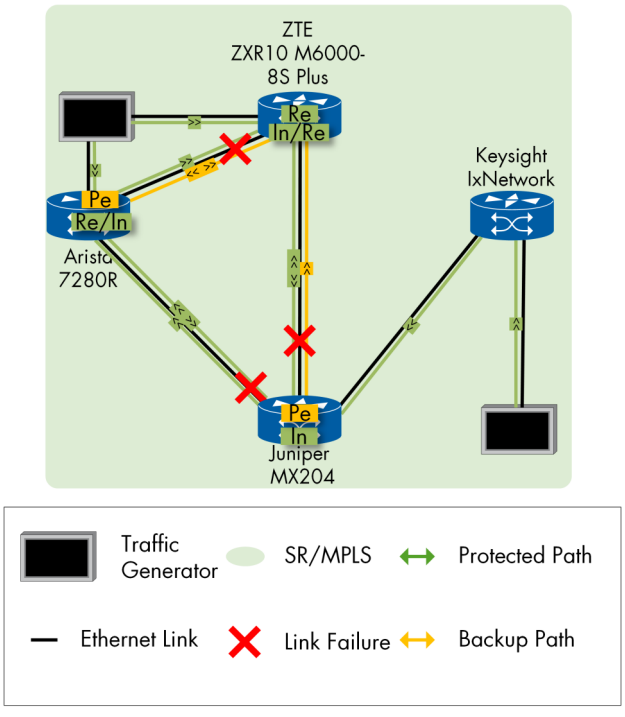

Figure 24: S-BFD over SR-MPLS

Arista 7280R, Juniper MX204, Keysight IxNetwork, and ZTE ZXR10 M6000-8S Plus participated as initiators and reflectors in this test.

After restoring the port we examined the S-BFD session and primary SR-TE path and both were up and traffic was flowing through the primary path. Out-of-service times were relatively high at first but after changing the node SIDs for building the LSP path and using instead adjacency SID, the times were between 145 ms and 156 ms.

One run had 221 ms out of service time later the vendor explained that they were still using the Node SID.

Seamless BFD over SRv6

We verified S-BFD over SRv6 in the control plane only.

Both nodes deployed SRv6 over ISIS. An SRv6 policy was configured to be used to steer traffic into the SR-TE tunnel when the Seamless BFD can detect the link failure and trigger SR-TE .

One node was configured as a responder while the other as a reflector and we observed the sessions are up between both of them.

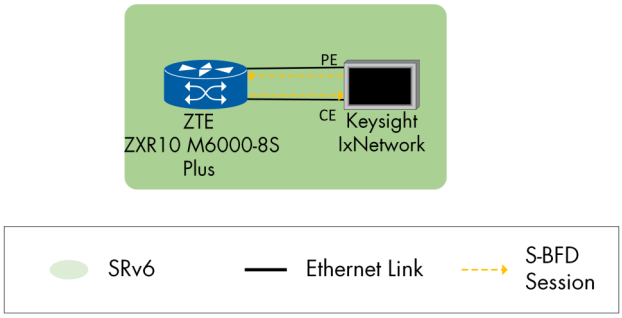

Figure 25: S-BFD over SRv6

Keysight IxNetwork and ZTE ZXR10 M6000-8S Plus successfully participated in the test.

EVPN

The event once again witnessed what was already a significant trend: Data centers and cloud environments play a vital and expanding role in this era of digital transformation with 5G. This test area aimed to verify interoperability of existing data center networking solutions to reflect experience.

We observed Ethernet VPN (EVPN) support in Carrier Ethernet services, providing E-Line services. The unified control solutions represented by BGP EVPN were self-learning through the network device control plane protocol: unicast integrated with Layer 2 and Layer 3 services, and multicast services. A variety of capabilities under EVPN categories benefit from the control layer, so the data plane becomes much simpler: Integrated Routing and Bridging (IRB) function which enables EVPN extension between subnets, optimized inter-subnet multicast (OISM), IGMP proxy, and MAC mobility. We observed multi-homing and single homing setups for these network services. Today, end-to-end cloud service delivery depends heavily on edge cloud processing and cloud interconnect services. We verified interconnection technologies of data centers: EVPN and IPVPN interworking, EVPN VXLAN-VXLAN network, seamless EVPN and VPLS.

E-Line Service

E-Line, as a point-to-point service model, is realized in an enhanced architecture in data centers. IETF working group BESS built VPWS on top of EVPN, which provides a powerful VPWS framework for data center designs. Benefits of VPWS of EVPN are single-active or all-active multihoming capabilities and support for Inter-autonomous system (AS) options associated with BGP-signaled VPNs. In addition, higher layer services can be encapsulated by the existing transport layer, EVPN SR with MPLS data plane, allowing label-based encapsulation of a label-switched packet network.

We created all E-Line services over SR-MPLS, including a mix of multi-homing and single-homing PEs. The test steps required the observation of three points as follows, firstly, network status that met the conditions for establishing E-Line services, like established BGP EVPN sessions and successful EVPN-VPWS signaling. Secondly, the DF election took place for Single-Active Multi-homing or Traffic Load balancing for All-active Multi-homing ES's. Finally, each PE's Ethernet A-D per-ES and per-EVI routes originating from remote PEs arrived at the routing table. As expected, none of the E-Line services showed traffic loss.

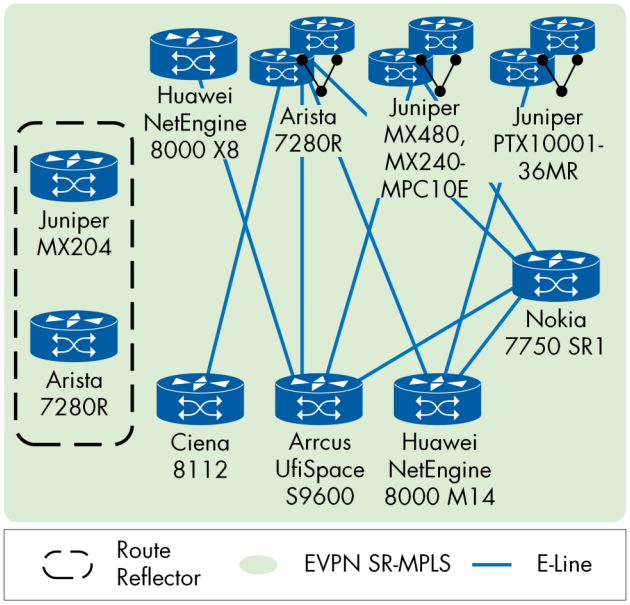

Figure 26: E-Line

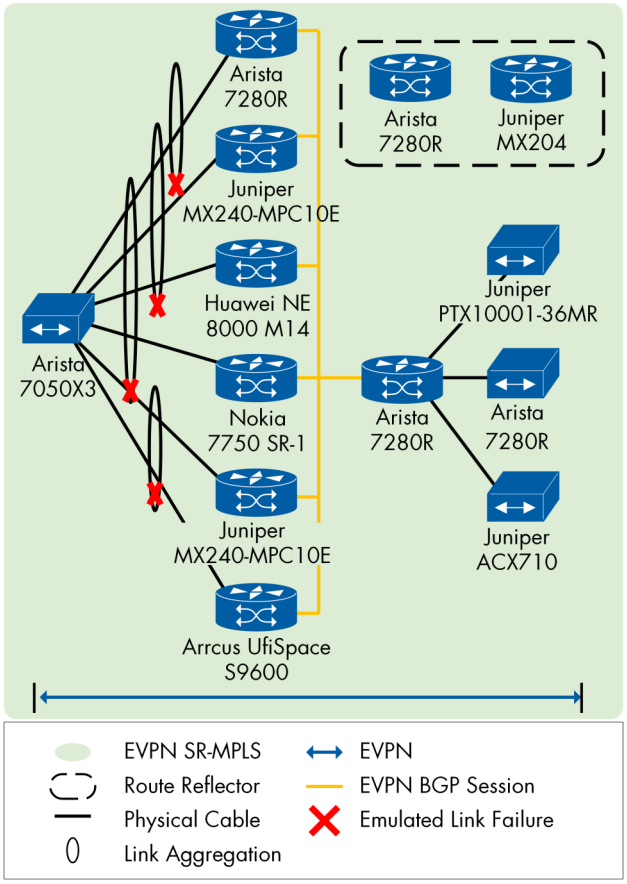

Figure 27: E-Line Service Multi-homing

Arista 7050X3, Arista 7280R, Juniper ACX710, and Juniper PTX10001-36MR participated in multi-vendor combinations in this test case as CEs. Arista 7280R was tested as single-homed PE. Arista 7280R, Arrcus UfiSpace S9600-72XC, Juniper MX240-MPC10E, Nokia 7750 SR-1, Huawei NE 8000 M14 participated as multihomed PEs in a multi-vendor environment.

Devices with below pairings have also tested Type3 ESI. Arista 7050X3, Juniper PTX10001-36MR participated as CEs. Arista 7280R as single homed PE, and Arrcus UfiSpace S9600-72XC, and Juniper MX240-MPC10E as multi-homed PEs in a multi-vendor environment.

Flexible Cross-Connect Service

This test verifies the network ability to enable the flexible cross-service in an E-Line scenario. This technique is beneficial for EVPN VPWS to maximize the ability of the tunnel to carry the number of AC sites. Multiple ACs across multiple Ethernet Segments are multiplexed into a single EVPN VPWS service tunnel, which is represented by a VPWS service ID. The multiplex reduces the number of EVPN service labels associated with the EVPN-VPWS service tunnel, thereby reducing EVPN BGP signaling from the system.

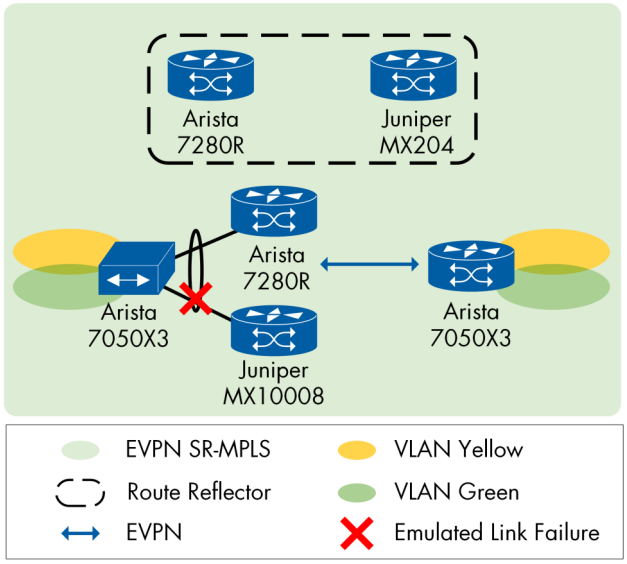

Figure 28: Flexible Cross-Connect Service

Arista 7050X3 successfully participated as CE in the test. Arista 7280R and Juniper MX10008 participated in multi-vendor combinations in this test case as multihomed PEs. Arista 7050X3 functioned as single-homed PE.

Integrated Routing and Bridging (IRB)

EVPN simplifies the architecture of integrated layer 2 layer 3 services, which allow hosts to communicate with each other within or across subnets in the EVPN. The integrated routing and bridging (IRB) provides a gateway between switched and routed networks. We verified symmetric and asymmetric IRB functionalities, using EVPN VLAN-based and VLAN-aware services.

Symmetric IRB

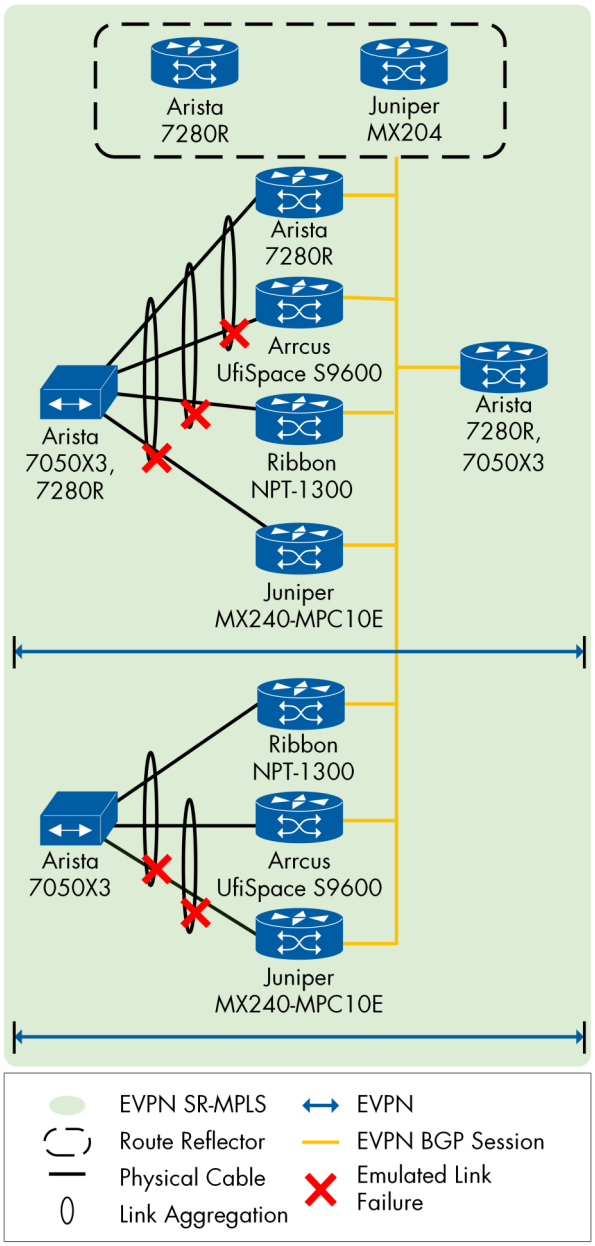

Figure 29: Symmetric IRB with VLAN-Based SR-MPLS

The symmetric IRB mode is analogous to a Layer 3 routing interface between different switches, where each tenant is assigned to a unique logical connection for IP-VRF. Two BGP route types are significant:

RT-2: The MAC and IP Route Type 2 is advertised with both Bridge-Domain/EVI label and IP VRF label with their respective route-targets

RT-5: IP prefix Route, is an alternative solution. A pure type-5 route operates without an overlay next hop or a type-2 route for recursive route resolution.

Arista 7280R, Arrcus UfiSpace S9600-72XC, Huawei Net-Engine 8000 M14, Juniper MX240-MPC10E, Juniper QFX5120-32C, Ribbon NPT-1300, Spirent SPT-N4U participated in multi-vendor combinations in this test case as multi-homed PEs. Arista 7050X3 and Arista 7280R were tested as single-homed CEs and PEs.

The identifier is the VXLAN network identifier (VNI) in VXLAN data plane and needs to be the same on all peers participating the same tenant's symmetric IRB. In MPLS data plane, this identifier is the MPLS Label2 associated with the IP-VRF. The test steps required the observation of following points: established BGP sessions between peers. When the link failure occurs, each PE receives Route Type 2 update.

The features under test included: VLAN-based or VLAN-aware verification. Route Type-2 and Route Type-5 tables.

None of the services showed any packet loss. There was no packet loss after link failure and recovery.

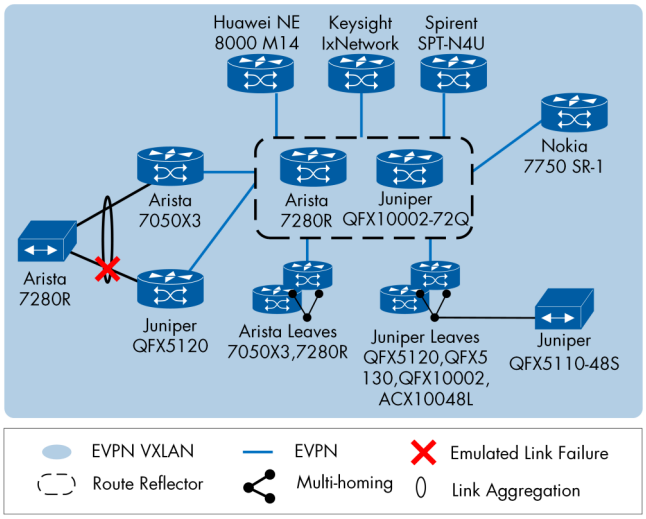

Figure 30: Symmetric IRB with VLAN-based VXLAN

These devices successfully participated in the test:

Multi-homed PEs: Arista 7280R, Arista 7050X3, Juniper ACX7100-48L, Juniper QFX5120-48Y, Juniper QFX5120-32C, Juniper QFX5130-32CD, Juniper QFX10002-72Q

Multi-homed PEs in a multi-vendor environment: Arista 7050X3 and Juniper QFX5120

Single-homed PEs: Huawei NetEngine 8000 M14, Keysight IxNetwork, Spirent SPT-N4U, Nokia 7750 SR-1

CE: Arista 7280R and Juniper QFX5110-48S

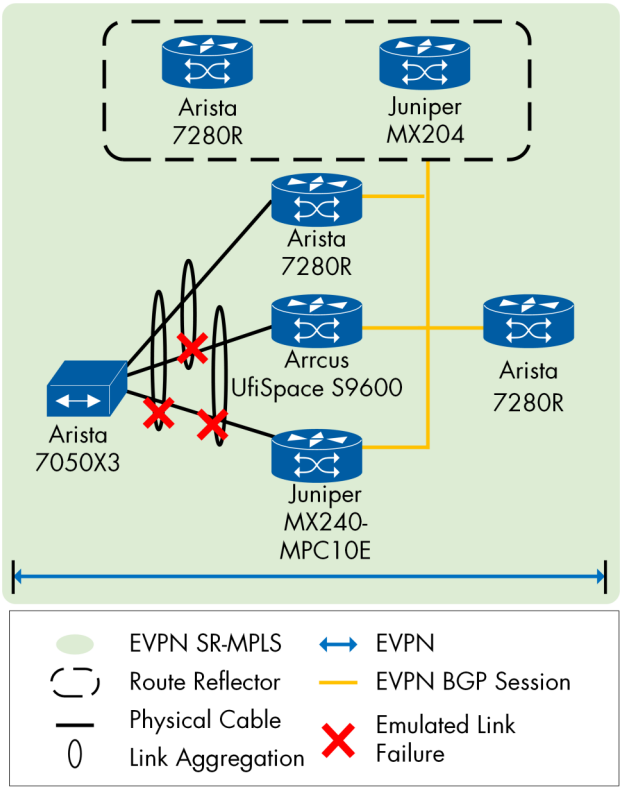

Figure 31: Symmetric IRB with VLAN-Aware bundle SR-MPLS

Arista 7280R, Arrcus UfiSpace S9600-72XC, and Juniper MX240-MPC10E participated in multi-vendor combinations in this test case as multi-homed PEs. Arista 7280R was tested as single-homed PE and Arista 7050X3 as CE.

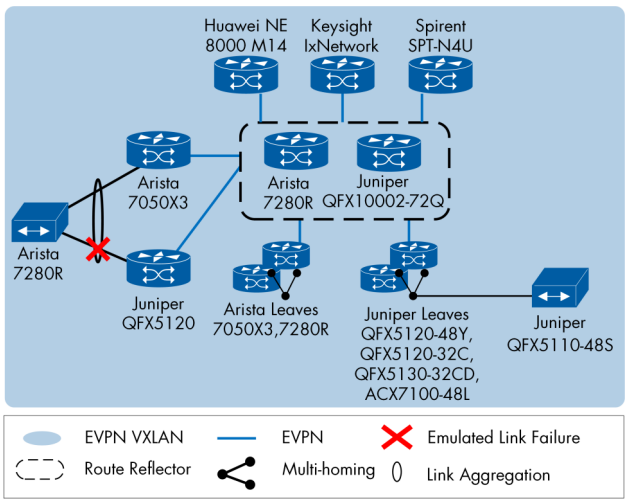

Figure 32: Symmetric IRB with VLAN-Aware bundle VXLAN

Arista 7280R, Arista 7050X3, Juniper QFX5110-48S, Juniper QFX5120-48Y, Juniper QFX5120-32C, Juniper QFX5130-32CD, and Juniper ACX7100-48L participated in multi-vendor combinations in this test case as multi-homed PEs. Arista 7050X3 and Juniper QFX5120 were tested as multi-homed PEs in a multi-vendor environment. Huawei NetEngine 8000 M14, Keysight IxNetwork, and Spirent SPT-N4U as single-homed PEs. Arista 7280R and Juniper QFX5110-48S acted as CEs.

Asymmetric IRB

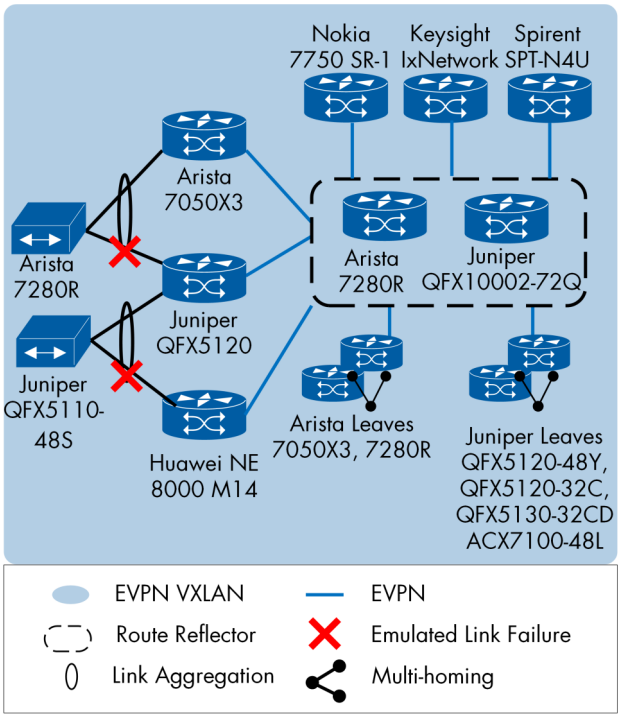

Figure 33: Asymmetric IRB with VLAN-based VXLAN

We verified asymmetric IRB functionality. In the asymmetric IRB semantic, both IP and MAC lookups are required at the ingress PE, whereas only MAC lookup is needed at the egress PE. The test steps required the observation of three points as follows: Established BGP EVPN sessions. ARP tables shall have remote MAC addresses. When the link failure occurs, each PE receives Route Type 2 update.

The features under test included: VLAN-based or VLAN-aware verification, Route Type-2 tables. None of the service showed packet loss. No packet loss was observed after link failure and recover.

Arista 7280R, Arista 7050X3, Juniper QFX5110-48S, Juniper QFX5120-48Y, Juniper QFX5120-32C, Juniper QFX5130-32CD, Juniper ACX7100-48L successfully participated as multi-homed PEs. Arista 7050X3, Juniper QFX5120, and Huawei NetEngine 8000 M14 as multi-homed PEs in a multi-vendor environment. Keysight IxNetwork, Nokia 7750 SR-1, and Spirent SPT-N4U as single-homed PEs. Arista 7280R and Juniper QFX5110-48S acted as CEs.

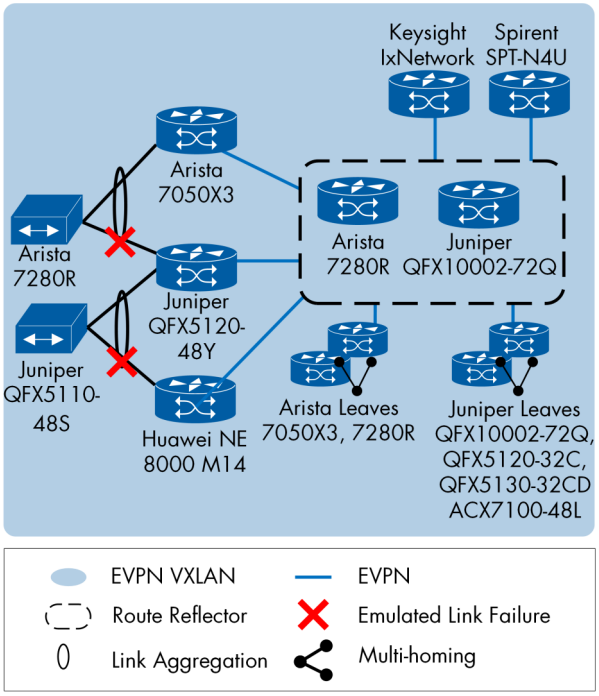

Figure 34: Asymmetric IRB with VLAN-aware bundle VXLAN

These devices successfully participated in the test:

Multi-homed PEs: Arista 7280R, Arista 7050X3, Juniper QFX5120-32C, Juniper QFX5130-32CD, Juniper ACX7100-48L, Juniper QFX10002-72Q

Multi-homed PEs in a multi-vendor environment: Arista 7050X3, Juniper QFX5120-48Y, Huawei NetEngine 8000 M14

Single-homed PEs: Keysight IxNetwork and Spirent SPT-N4U

CEs: Arista 7280R and Juniper QFX5110-48S

Proxy MAC-IP Advertisement

EVPN allows distribution of traffic to connected hosts over different PEs. Traffic follows between the active links based on the hashing algorithm known as All-Active multi-homing. The network will learn that not all the host routes (MAC-IP bindings) will be learned through the same PE. Thus different knowledge of host routes (MAC-IP bindings) appears. To solve this problem, the Proxy MAC-IP advertisement provides L3 ECMP of host routes (MAC-IP bindings) across the PEs sharing the ESI.

We sent test traffic to ensure that during the steady-state, all the PE's which belong to the same ESI learned a different set of MAC addresses and advertise different sets of (IP, MAC) EVPN route-type 2.

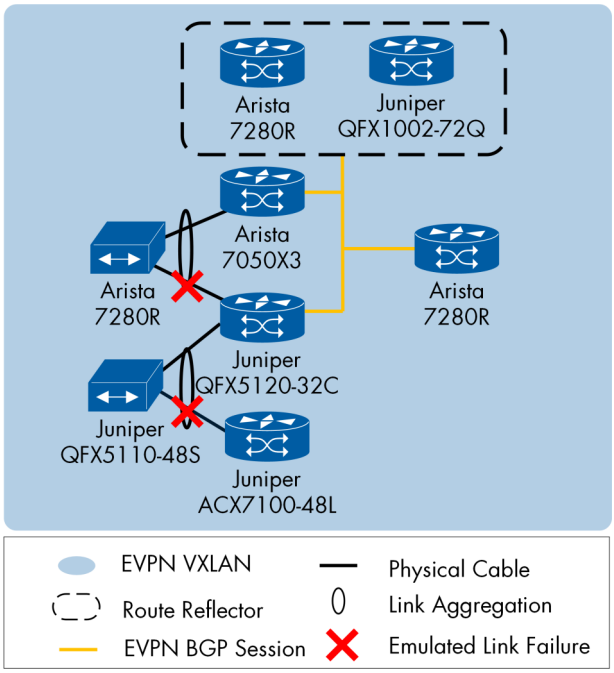

Figure 35: Proxy MAC-IP Advertisement VXLAN

Juniper QFX5120-32C and Juniper ACX7100-48L were tested as multi-homed PEs. Arista 7050X3 and Juniper QFX5120-32C successfully participated as multi-homed PEs in a multi-vendor environment. Arista 7280R acted as single-homed PE, Arista 7280R and Juniper QFX5110-48S as CEs.

In a test scenario of link failure emulated, we verified that proxy MAC-IP advertisement was enabled. In case of link or node failure, EVPN type 2 routes were withdrawn and these (IP, MAC) addresses were not be reachable, till the traffic is sent out through a different active link by the CE. Proxy (IP-MAC) allowed all the PE's in the same ESI to re-advertise the same EVPN route-type 2 even it was not learned locally, provided that the proxy bit is set.

As expected, traffic to flow in and out of the multi-homed CE during the transient time of link failure, indicating successful proxy MAC-IP advertisement.

Arista 7280R and Arrcus UfiSpace S9600-72XC participated as multi-homed PEs in a multi-vendor environment, Arista 7280R as single-homed PE and Arista 7050X3 as CE.

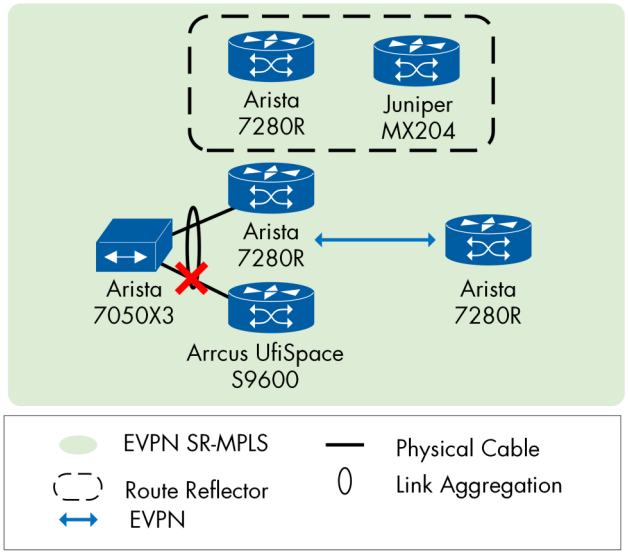

Figure 36: Proxy MAC-IP Advertisement SR-MPLS

Optimized Inter-Subnet Multicast (OISM)

EVPN optimized inter-subnet multicast (OISM) provides multicast VPN services, especially allowing optimization in a network that integrates Layer 2 and Layer 3; that is, EVPN can connect networks of different subnets or connect to the same subnet. Initially, the inter-domain multicast first appeared in the IRB draft (draft-ietf-bess-evpn-inter-subnet-forwarding), called BUM, but based on the unicast forwarding path. For an optimized multicast path, OISM (draft-ietf-bess-evpn-irb-mcast) defines optimized multicast path across subnets.

In an OISM without SBD (Supplementary Broadcast Domain) scenario, we tested OISM in which all BDs (Broadcast Domains) belong to the same tenant domain. We also verified OISM with SBD scenario, which is associated to separate tenant domain.

We used Ingress replicator mode in this multicast setup. The ingress PE (connected to the multicast source) duplicates multicast traffic based on interest in the multicast group received. The participating PEs established PMSI (P-Multicast Service Interface) tunnel with each other based on the RT-3 route (Inclusive multicast Ethernet Tag route). The PEs sent and learned RT6 SMET (Selective Multicast Ethernet Tag Route) in each domain for interested multicast groups. We sent multicast traffic from the emulated source, the ingress replicator forwarded the multicast traffic to all egress PEs. After receiving the multicast traffic, the PE re-encapsulated the traffic and forwarded it to the corresponding domain without any packet loss.

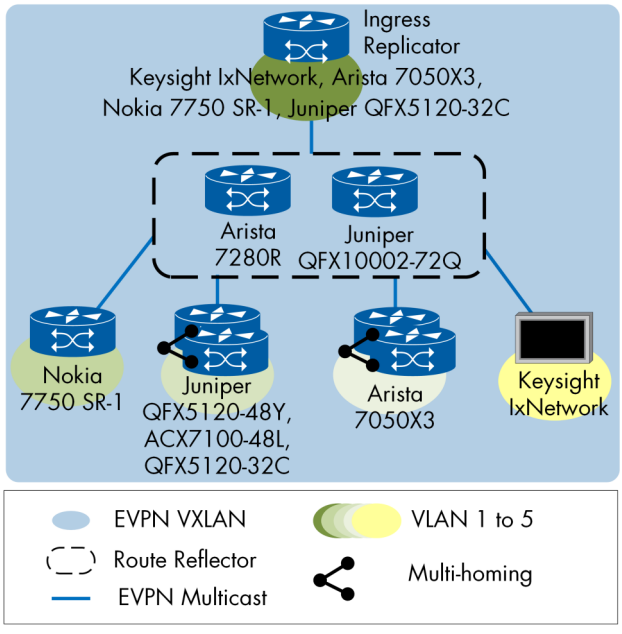

Figure 37: OISM With BD Everywhere

The PEs Arista 7050X3, Juniper QFX5120-32C, Juniper QFX5120-48Y, Juniper ACX7100-48L, Keysight IxNetwork, and Nokia 7750 SR-1 participated in the test.

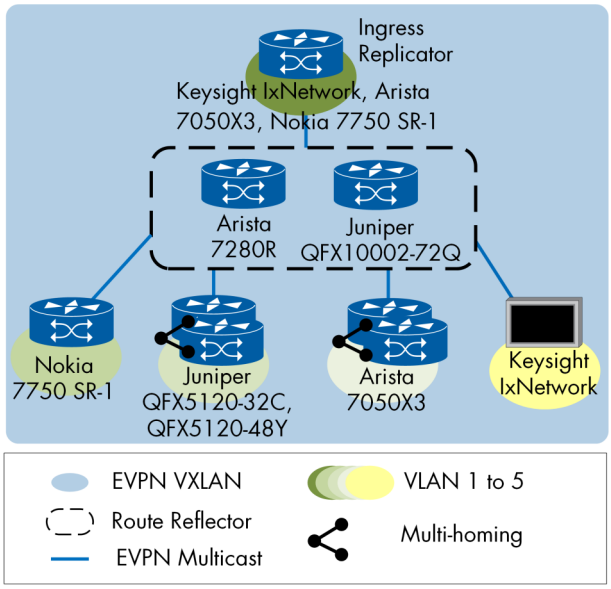

Figure 38: OISM with SBD

Arista 7050X3, Juniper QFX5120-48Y, ACX7100-48L, Juniper QFX5120-32C, Keysight IxNetwork, and Nokia 7750 SR-1 successfully participated in the OISM with SBD test.

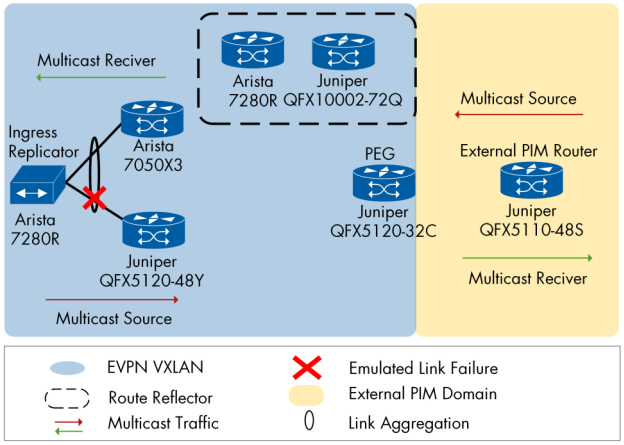

Figure 39: OISM with PEG

Arista 7050X3, Arista 7280R, Juniper QFX5120-32C, Juniper QFX5120-48Y, Juniper QFX5110-48S successfully participated as PEs in the test.

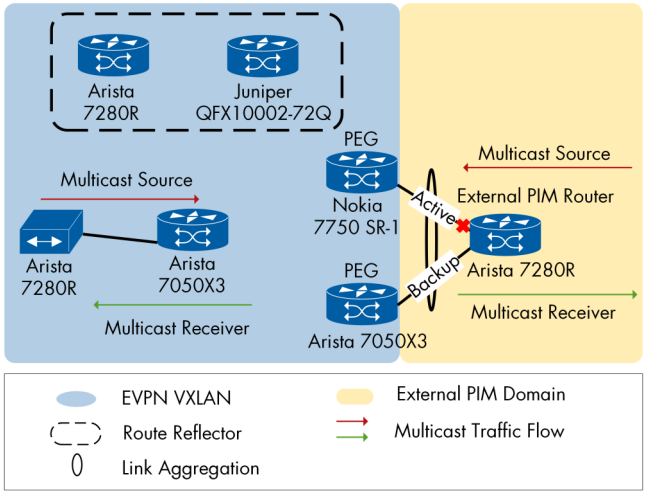

Figure 40: OISM with PEG Election

Arista 7280R and Arista 7050X3 functioned as PEs, Arista 7050X3 and Nokia 7750 SR-1 participated as multihomed PEs in a multi-vendor environment. Arista 7280R functioned as CE.

IGMP Proxy

EVPN forwards IGMP messages through EVPN routes to minimize the scope of related PEs and reduce the flood of IGMP messages (queries and reports).

The goal of the IGMP proxy mechanism is to reduce the flood of IGMP messages (both Queries and Reports) in EVPN instances among PE routers. Furthermore, if there is no physical/virtual multicast router attached to the EVPN network for a given (*,G) or (S,G), it is desired for the EVPN network to act as a distributed anycast multicast router for all the hosts attached to that subnet.

We observed that the PE device received selective Multicast Ethernet Label (Type 6: SMET) from EVPN service, indicating that participating leaf peers of the service successfully registered their interest in the selected multicast group. For leaf PEs that joined the group, depending on whether these were multi-homed, we also verified IGMP Proxy (Type 7: IGMP Join Sync Route and Type 8: IGMP Leave Sync Route) exchanged, ensuring that only the designated PE has been for multicast forwarding registered.

No packet loss was observed in the multicast groups. All leaf PEs received successfully the test traffic without any loss.

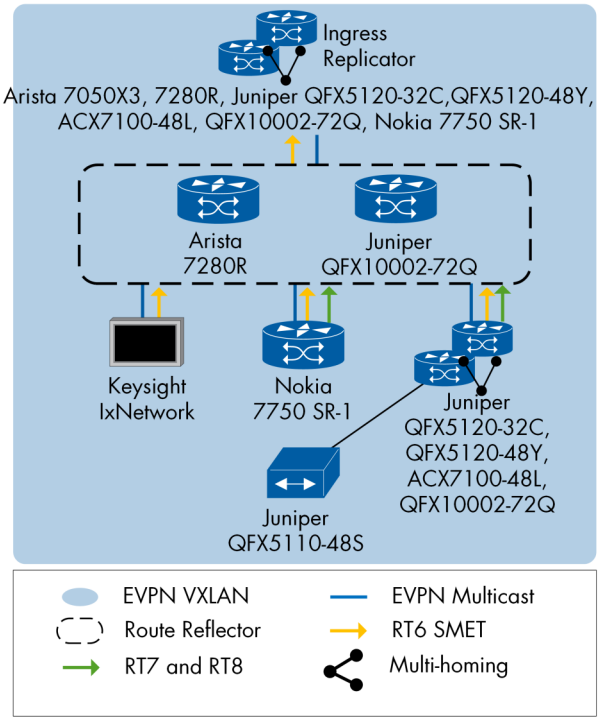

Figure 41: IGMP Proxy Layer 2

These devices successfully participated in the tests where both the receivers and source were on the same subnet. All vendors were involved in the tests as ingress replicator or as receivers.

Keysight IxNetwork participated as multicast source.

Arista 7050X3, Arista 7280R, Juniper QFX5120-32C, Juniper QFX5120-48Y, Juniper ACX7100-48L, Juniper QFX10002-72Q, Nokia 7750 SR-1, as All-active Multihoming, and Keysight IxNetwork were tested as ingress replicators.

Arista 7050X3, Arista 7280R, Juniper QFX5120-32C, Juniper QFX5120-48Y, Juniper ACX7100-48L, Juniper QFX10002-72Q as All-active Multihoming, and Keysight IxNetwork as receivers. Juniper QFX5110-48S functioned as CE.

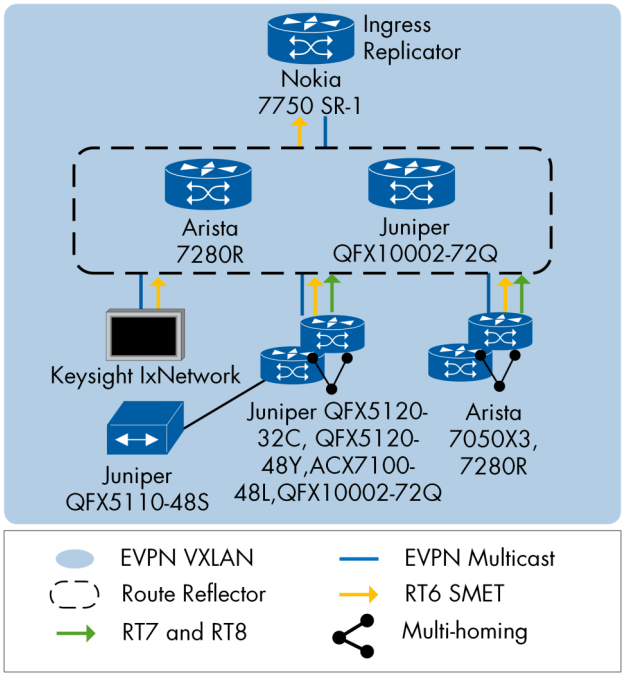

Figure 42: IGMP Proxy

Keysight IxNetwork acted as multicast source and Nokia 7750 SR-1 as ingress replicator.

Arista 7050X3, Arista 7280R, Juniper QFX5120-32C, Juniper QFX5120-48Y, Juniper ACX7100-48L, Juniper QFX10002-72Q as All-active Multihoming, and Keysight IxNetwork were involved as receivers. Juniper QFX5110-48S functioned as CE.

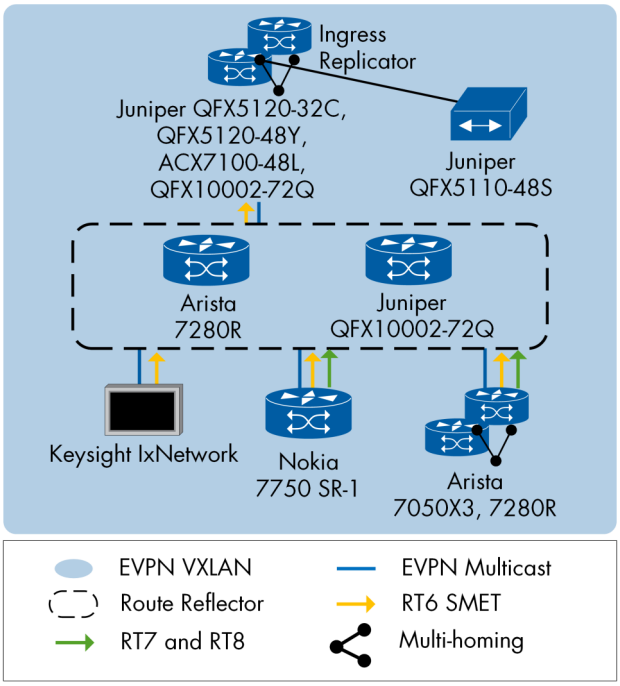

Figure 43: IGMP Proxy

Keysight IxNetwork acted as multicast source and Arista 7050X3 and Arista 7280R as ingress replicators.

Juniper QFX5120-32C, Juniper QFX5120-48Y, Juniper ACX7100-48L-1 and Juniper QFX10002-72Q as All-active Multihoming, and Keysight IxNetwork and Nokia 7750 SR-1 were involved as receivers. Juniper QFX5110-48S functioned as CE.

EVPN and IP-VPN Interworking

Multi-domain operators can interconnect different data center networks running EVPN for end-to-end service delivery. A WAN running MPLS-based IP-VPN works as a central network to transport EVPN with VXLAN encapsulation from data centers. We verified the Interconnection of Data Center Networks Through WAN.

We verified the D-PATH capability on multi-homed gateways for loop prevention. D-PATH is optional and transitive BGP path attribute as specified in draft "EVPN Interworking with IPVPN" (draft-ietf-bess-evpn-ipvpn-interworking). Similar to AS_PATH, D-PATH is composed of a sequence of Domain segments.

Arista 7280R, Arrcus UfiSpace S9600-72XC, and Nokia 7750 SR-1 participated as single-homed PEs in this test. Traffic flow from DC1 to DC2 was successful. D-Path validated on the resilient gateways into DC1 (Arista 7280R, Nokia 7750 SR-1).

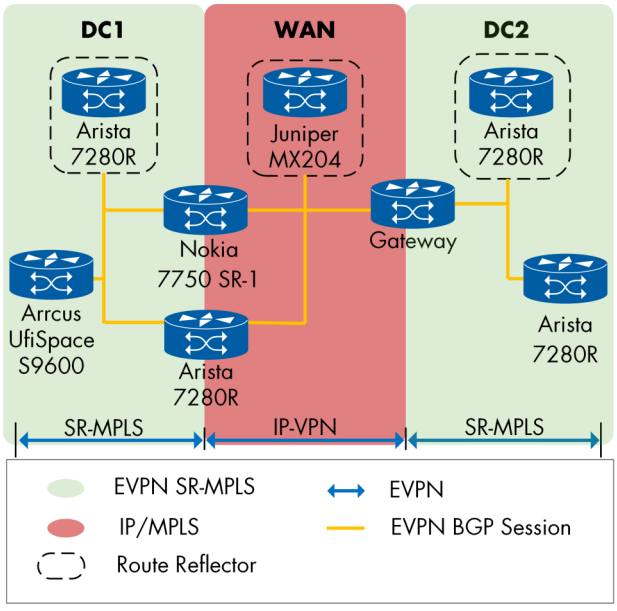

Figure 44: EVPN and IP-VPN Interworking SR-MPLS

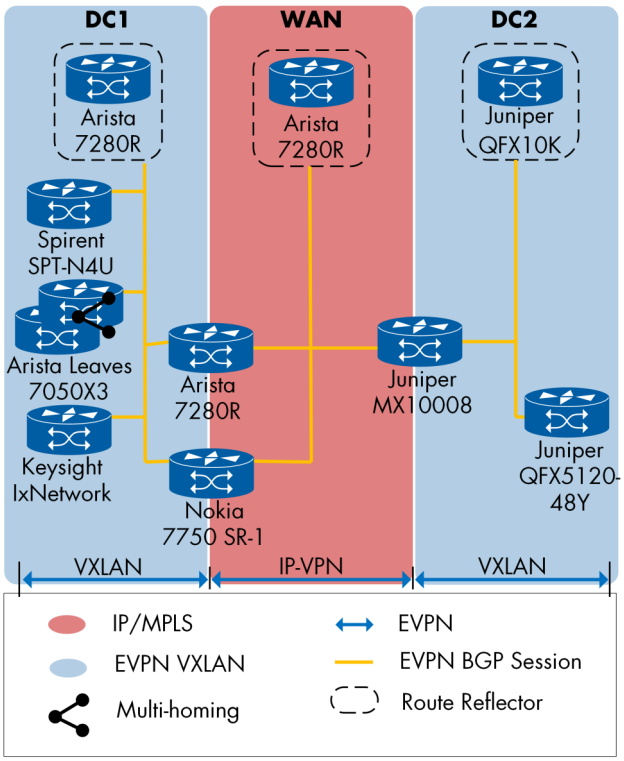

Figure 45: EVPN and IP-VPN Interworking VXLAN

Arista 7050X3 successfully participated as multi-homed PEs. Arista 7280R, Juniper MX10008, Juniper QFX5120-48Y, Keysight IxNetwork, Nokia 7750 SR-1, and Spirent SPT-N4U were tested as single-homed PEs. D-Path validated on the resilient gateways into DC1 (Arista 7280R, Nokia 7750 SR-1).

EVPN VXLAN and VXLAN Interworking

Multi-domain operators can interconnect different data center networks running EVPN for end-to-end service delivery. A WAN running VXLAN works as a central network to transport EVPN with VXLAN encapsulation from data centers. We verified the Interconnection of Data Center Networks Through WAN.

VXLAN is a widely supported data plane technology, which encapsulates a MAC frame in a UDP datagram for transport across an IP network. To be able to offer a regional or national EVPN network, service providers are seeking flexible approaches to extend the reach of EVPN beyond a single data center. One mechanism is the use of VXLAN in the Metro Area Network (MAN) to interconnect multiple EVPN domains.

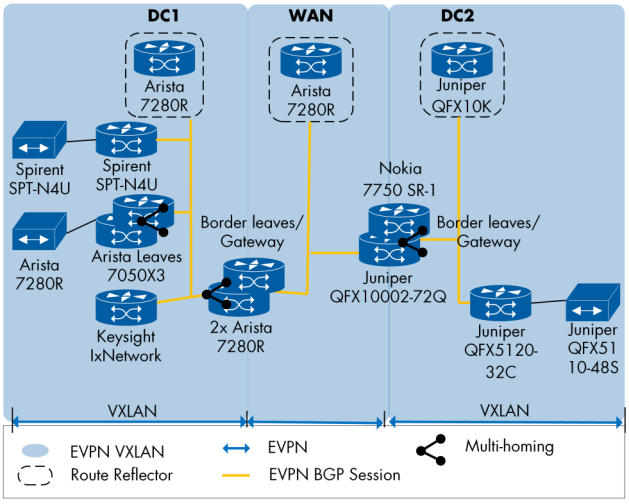

Figure 46: EVPN-VXLAN to EVPN-VXLAN Interworking

Arista 7050X3 and Arista 7280R successfully participated as multihomed PEs, Juniper QFX10002-72Q and Nokia 7750 SR-1 as multi-homed PEs in a multi-vendor environment.

Arista 7280R, Keysight IxNetwork, Juniper QFX5120-32C, and Spirent SPT-N4U functioned as single-homed PEs. CEs were Arista 7280R, Juniper QFX5110-48S, and Spirent SPT-N4U.

MAC Mobility

EVPN provides a mechanism to automatically track the location of the host and update the MAC address. This mechanism saves the effort of manual provisioning in provider's network. Where does a host come from the datacenter, EVPN learns it through the RT2 route. If the host moves, the unaged MAC address would lead to an inconsistency EVPN therefore also provides a sequencing mechanism to track to where a host moves, referred to MAC Mobility Extended Community as defined by RFC 7432. When a newly learned MAC address would be found in the MAC table which had been learned from a remote end, the sequence number of the MAC Mobility Extended Community shall increase by one and the value is carried out via the RT2. The EVPN learns from the highest sequence number the latest update of where the host is connected to, this view prevents race conditions which might exist with multiple rapid moves.

In this test we first connected an emulated host that has not been moved before to an EVPN segment, to confirm that within this initial state MAC/IP advertisement of the MAC address on the PE showed the sequence number 0. This information was required because we used it for comparison in the next step when we moved the host to a different EVPN segment by changing the traffic from previous PE to a new PE. Then the value increased by 1. This proved that a PE receiving a MAC/IP Advertisement route for a MAC address with a different Ethernet segment identifier and a higher sequence number than that which it had previously advertised from its MAC/IP Advertisement route. We sent test traffic and did not observe any frame loss, we also did not receive any flooded traffic.

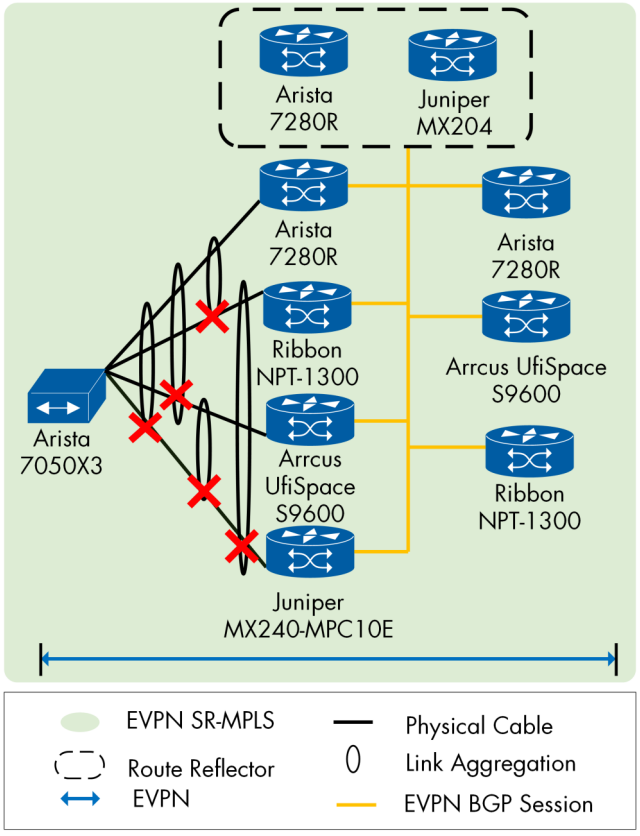

Figure 47: MAC Mobility SR-MPLS

These devices successfully participated in the test:

Multi-homed PEs in a multi-vendor environment: Arista 7280R, Arrcus UfiSpace S9600-72XC, Juniper MX240-MPC10E, Ribbon NPT-1300

Single-homed PEs: Arista 7280R, Arrcus UfiSpace S9600-72XC, Ribbon NPT-1300

CE: Arista 7050X3

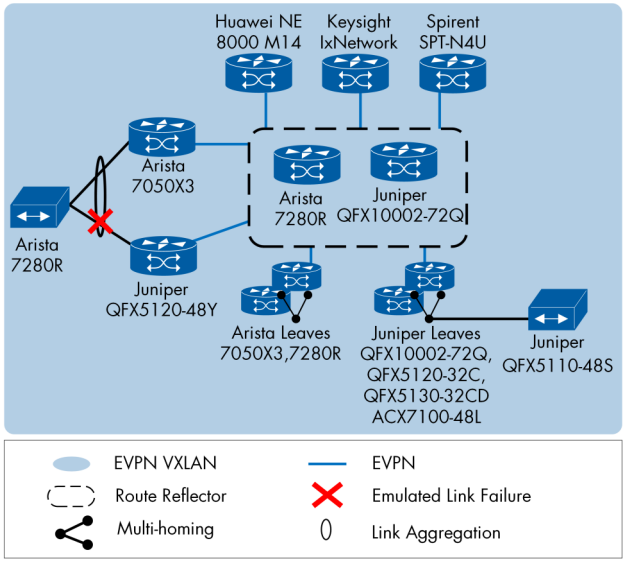

Figure 48: MAC Mobility VXLAN

These devices successfully participated in the test:

Multi-homed PEs: Arista 7280R, Arista 7050X3, Juniper ACX7100-48L, Juniper QFX5120-32C, Juniper QFX5130-32CD, Juniper QFX10002-72Q, Juniper QFX5110-48S

Multi-homed PEs in a multi-vendor environment: Arista 7050X3 and Juniper QFX5120-48Y

Single-homed PEs: Huawei NetEngine 8000 M14, Keysight IxNetwork, Spirent SPT-N4U

CE: Arista 7280R and Juniper QFX5110-48S

Seamless EVPN and VPLS

EVPN provides backward compatibilities to VPLS PEs as defined in RFC8560. VPLS is a widely deployed l2VPN technology. Service providers who are looking at adopting EVPN want to pass the success of existing solutions to the new solution. The solution must not require any changes to existing VPLS, not even a software upgrade. In order to support seamless integration with VPLS PEs, the RFC requires that VPLS PEs support VPLS A-D per [RFC6074], and it requires EVPN PEs to support both BGP EVPN routes per [RFC7432] and VPLS A-D per [RFC6074]. All the logic for seamless integration shall reside on the EVPN PEs. The EVPN PE establishes VPWS to VPLS PE.

We verified end-to-end EVPN between VPLS and EVPN PEs.

The EVPN PEs joined and fully discovered the VPLS PEs, then they established full-meshed VPWS devices with each other. We sent traffic through the VPLS. Once the services have been established, all traffic went through without any frame loss as expected.

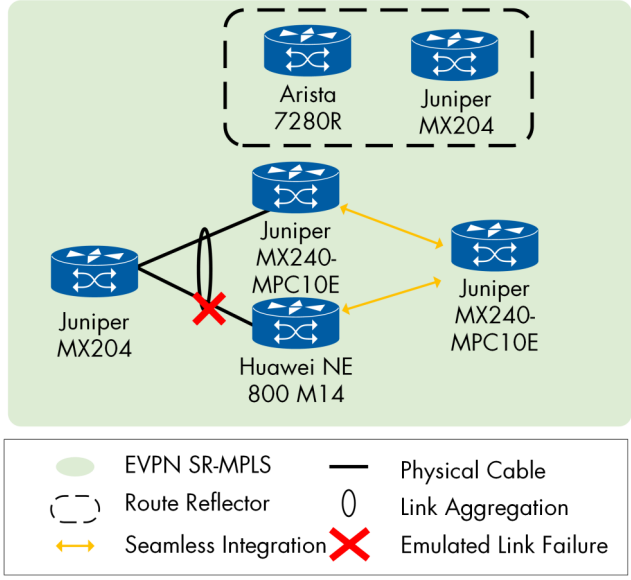

Figure 49: Seamless EVPN and VPLS

Huawei NetEngine 8000 M14, Juniper MX240-MPC10E acted as multi-homed PEs in a multi-vendor environment, and Juniper MX240-MPC10E as single-homed PE. Juniper MX204 functioned as CE.

PW Headend Multi-homed EVPN-VPWS Access to L3VPNs

We verified PW Headend termination with EVPN-VPWS access to L3VPN. We also verified redundant active-standby transport connectivity between multiple service PEs.

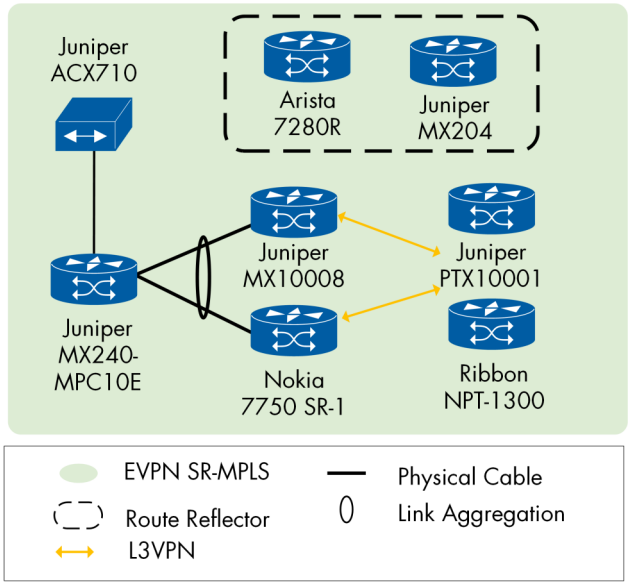

Figure 50: PW Headend Multi-homed EVPN-VPWS access to L3VPNs

Juniper 10008 and Nokia 7750 SR-1 successfully participated as multi-homed PEs in a multi-vendor environment in this test, and Juniper PTX10001 and Ribbon NPT-1300 as single-homed PEs. CE was Juniper ACX710.

Flexible Algorithm

Flex Algorithm (Flex Algo) is a great tool for designing route computations according to the network traffic engineering needs. This leverages the SR-TE in one hand and a provides a powerful key as slicing infrastructure for 5G networks.

The tests included creating paths with a subset of routers in the network as basic feature of deploying the Flex Algo. We also used measurements of delay per link to select the path that offered the least cumulative delay to a destination.

Prefix metric propagation between multiple IGP instances and creation of Flex Algo multi-planes with different policy constrains (TE Metrics, Affinity Links) was also tested and confirmed.

Flexible Algorithm Prefix Metric

The limitation of the existing functionality of the ISIS flexible algorithm is the inability to calculate the optimal path to the prefix of the remote area or remote IGP domain. Prefixes are advertised across ISIS regions or protocol domains, but existing prefix metrics do not reflect the constraints used in flexible algorithm paths. Flexible algorithms can calculate the optimal path to a prefix between areas or a redistributed prefix within an area, but the path does not represent the overall optimal path through multiple areas or IGP domains.

The Flexible Algorithm Prefix Metrics (FAPM) feature introduces flexible algorithm-specific prefix metrics into ISIS prefix advertisements. Prefix metrics provide a way to calculate the optimal end-to-end path across multiple ranges or domains optimized by flexible algorithms.

We created two ISIS levels in the network for SR-MPLS. By adding a Flex Algo over the underlay topology, we verified that Flex Algo 129 (based on delay metric) and Flex Algo 128 (based on IGP metric) include both levels in the network.

VPN prefixes were mapped to the different Flex Algos using the BGP color community.

Original traffic is following the lowest delay and due to inter-level FAPM, the delay metric is correctly propagated between levels. This assures correct routing decisions between levels.

Later we performed an increase of the delay on one of the links using an impairment device. By setting the M-flag the Flex Algorithm-specific prefix metric MUST be used for inter-area and external prefix calculation. So this change of delay will be leaked to the second level and the traffic is switched to the other route.

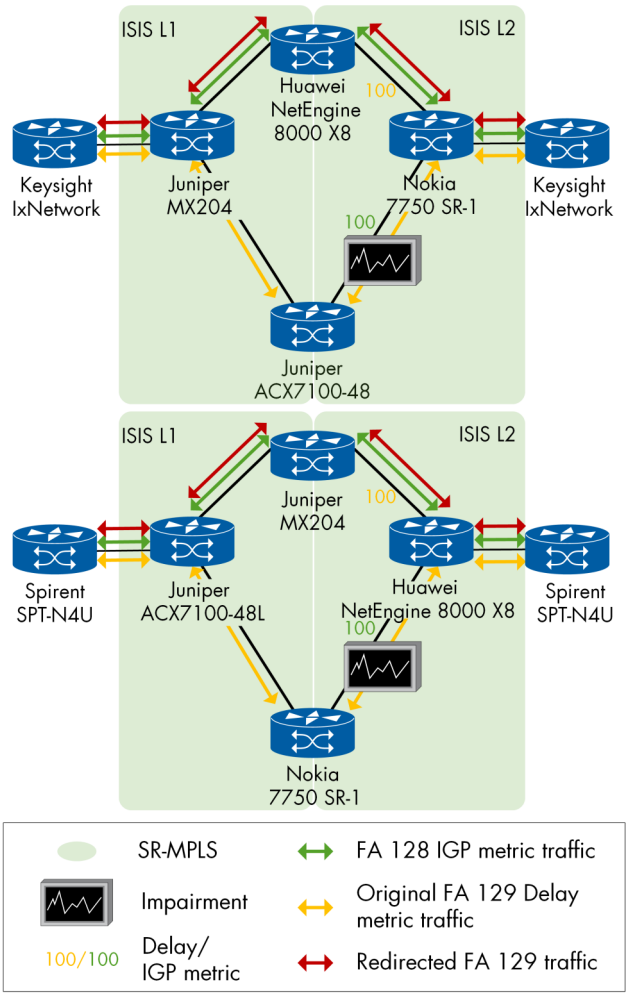

Figure 51: Flexible Algorithm Prefix Metric

Huawei NetEngine 8000 X8, Juniper ACX7100-48L, Juniper MX204, Keysight IxNetwork, Nokia 7750 SR-1, and Calnex SNE as Impairment device participated in the test.

SR-MPLS Flexible Algorithms

The goal of the test includes Flex Algo multi-planes and isolation, to allow the existing ISIS underlay network to unfold its full potential.

We created two Flex Algo definitions and expect that the ISIS underlay shall calculate a set of nodes and links of each Flex Algo, and based on the collected performance metric information to form two different Flex Algo planes.

We verified Flex Algo for SR-MPLS using the following definitions:

FA	Link Metric
FA 128	Low TE Metrics
FA 129	Low Delay

Table 3: SR-MPLS Flex Algo

We created L3VPNs in each of the Flex Algo definitions. In different parts of the network, we added traffic and expected each VPN to follow its own Flex Algo path.

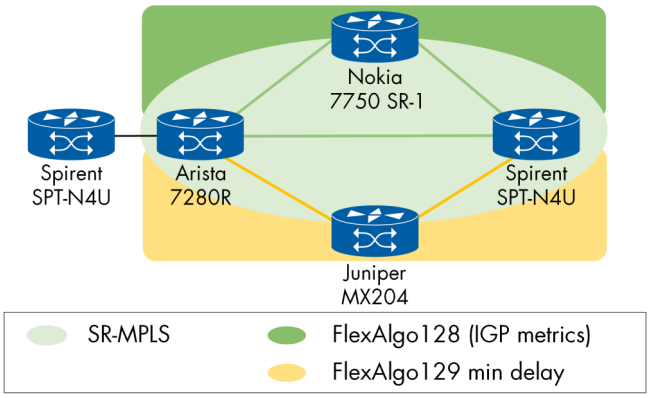

Figure 52: SR-MPLS FA

Arista 7280R, Juniper MX204, Nokia 7750 SR-1, and Spirent SPT-N4U successfully participated.

An extra step was added to the test to check the behavior of the network while a link failure on the green plane. In principle, the traffic will be dropped completely as there are no alternative nodes participating in the FA129 to the destination. But we have noticed that the DUT has fallen back to segment routing instead of dropping the traffic when no paths are available within the Algo.

SRv6 Flexible Algorithms

SRv6 introduces flexible algorithms to the IPv6 data plane. Each SRv6 locator is associated with an algorithm, representing a topologically-constrained forwarding path.

We verified Flex Algo for SRv6 using the following definitions:

FA	Link Metric
FA 131	Admin Groups (include all yellow)
FA 132	TE Metrics
FA 133	Delay metrics
FA 134	Admin Groups (include all green)

Table 4: SRv6 Flex Algo

FA131 and FA132 traffic were flowing between nodes including only yellow links or green links, and FA132 traffic avoided the links with the high TE metric. For FA133 the delay value came from the delay measurement value measured by the DUT via TWAMP-light. This method by default resulted in the shortest path with less delay to be included in the Flex Algo 133.

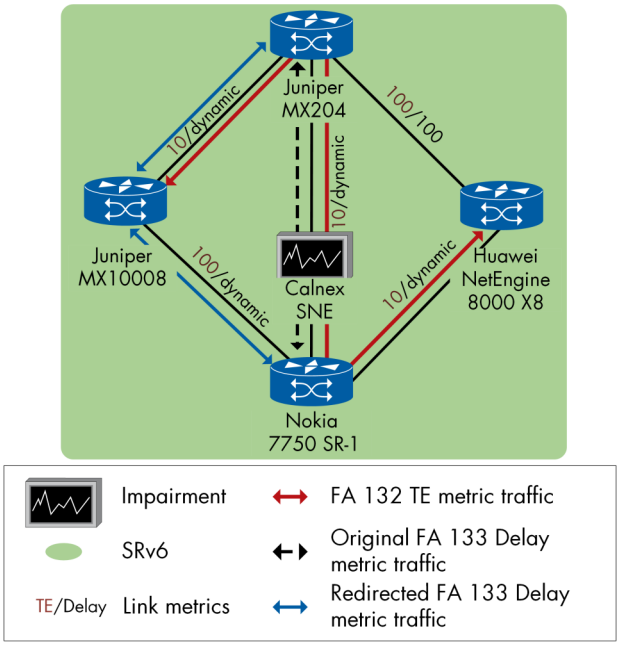

Figure 53: SRv6 FA with TE and Delay Metrics

Huawei NetEngine 8000 X8, Juniper MX10008, Juniper MX204, Nokia 7750 SR-1, and Calnex SNE as Impairment Device participated in the test.

In addition, the delay measurement feature added one more test step, to verify that when the delay value changed dynamically, the creation of Flex Algo 133 was based on the dynamic value. So, we increased with the impairment device 500 ms delay over the short path and observed as the Flex Algo133 included the links with smaller delay value.

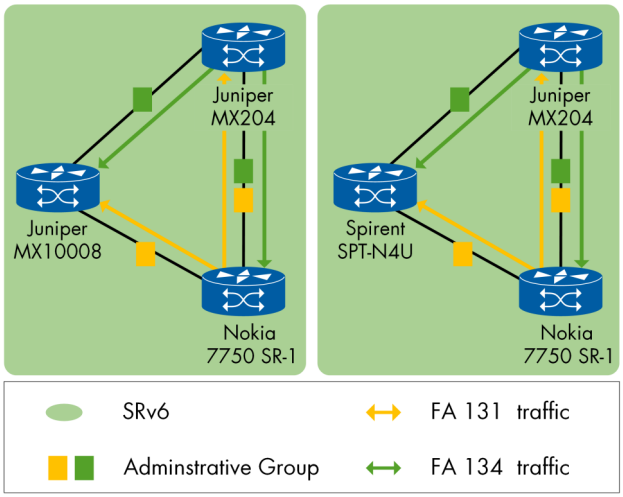

Figure 54: SRv6 FA with Administrative Groups

Juniper MX10008, Juniper MX204, Nokia 7750 SR-1, and Spirent SPT-N4U successfully participated in the test.

In one of the iterations of this test-case, using link delay as the metric-type for Flex Algo 133, one of the participating devices was not able to compute a new shortest path after the link delay was increased and consequently was not able to steer traffic over the shortest path based on end-to-end delay.

Also in the test case, using Administrative Groups as the metric-type for Flex Algo 131/134, one of the participating devices only supported the Extended Admin Groups without being compatible with Admin Groups as well. So the test was carried out only with Admin Groups supporting devices.

Software-Defined Networks

Software-Defined Networks (SDN) are getting mature more and more through the years. The technologies and tools under the SDN umbrella are more reliable and consistent, this was the basis of our plans for the SDN Interoperability Tests 2022. This year we were able to verify basic features such as Topology Discovery through BGP LS, and PCE Path Computation, and more advanced tests like managing SR policies via BGP SR-TE, and EPE.

As we have tested more than 35 successful combinations for the test cases only in SDN Area, we still did not see much progress regarding PCEP, PCEPv6, and some SRv6 features. On the contrary, we found improvement regarding the interoperability of the BGP SDN features.

Topology Discovery

Utilizing the Traffic engineering and Link State information to collect the topology information is a huge advantage to manage and administer modern networks. Using BGP-LS information and exporting them from the Path Computation Client (PCC) to the Path Computation Element (PCE) enables it to be able to create a topology of the network and compute the packets paths through the network elements.

Table 5 shows the combinations that passed the test successfully.

During the test we enabled the BGP-LS topology export from the PCC, then enabled the BGP-LS session on the PCE. When the PCE received the topology information through the BGP-LS session, the PCE was able to display the network topology correctly.

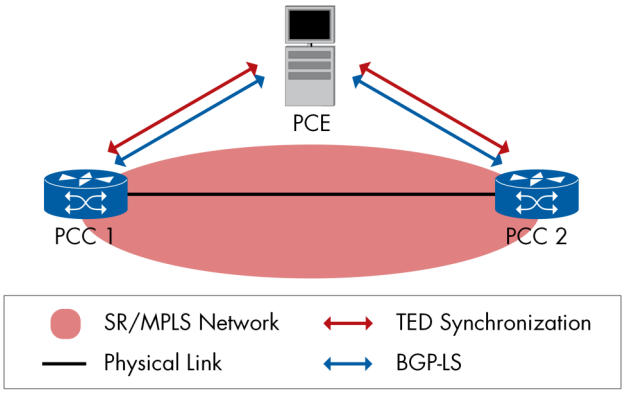

Figure 55: Test Topology—Topology Discovery

There were no fundamental interop issues during this test, although we observed that one vendor relied on the BGP-LS Identifier to recognize different topologies, which complies with RFC7752 as it was in 2016, but other vendors have already updated their software according to the draft RFC 7752bis in the year 2019 in which the BGP-LS Identifier TLV was deprecated.

This issue caused the mentioned vendor to interpret the multiple information coming from different network elements from the same topologies as multiple topologies.

PCE	PCC 1	PCC 2
Huawei NCE Path Computing Element	ZTE ZXR10 M6000-8S Plus	Juniper MX204
Huawei NCE Path Computing Element	Huawei ATN 910D-A	Nokia 7750 SR-1
Juniper Paragon Pathfinder	Nokia 7750 SR-1	ZTE ZXR10 M6000-8S Plus
Juniper Paragon Pathfinder	Nokia 7750 SR-1	Nokia 7750 SR-1
Keysight IxNetwork	Juniper MX204	Nokia 7750 SR-1
Keysight IxNetwork	Huawei ATN 910D-A	Nokia 7750 SR-1
Nokia NSP	Huawei ATN 910D-A	ZTE ZXR10 M6000-8S Plus
Nokia NSP	ZTE ZXR10 M6000-8S Plus	Juniper MX204
Nokia NSP	Juniper MX204	Huawei ATN 910D-A
ZTE ZENIC ONE	Juniper MX204	Huawei ATN 910D-A
ZTE ZENIC ONE	Nokia 7750 SR-1	Huawei ATN 910D-A
ZTE ZENIC ONE	Nokia 7750 SR-1	Juniper MX204

Table 5: Topology Discovery

PCE Path Computation

Network applications and their demands play a central role in the context of the SDN. The ability of the network to be agile and flexible in an environment where the network paths need to be changed on-demand is crucial.

This test verifies the ability of the PCE to trigger an LSP creation in response to Applications needs between different vendors.

We checked the LSP setup, state synchronization, update, and deletion of PCE-initiated LSPs under the stateful PCE model, without the need for local configuration on the PCC.

- The DUTs started the IGP adjacencies between them, and the connectivity was verified. For this test, the DUTs established IS-IS as IGP.
- We verified the Stateful PCEP session.
- We verified PCE path instantiation.
- LSP state synchronization was verified.
- For this test we did not create VPN services to generate traffic, we used the pings to confirm transport paths were installed.

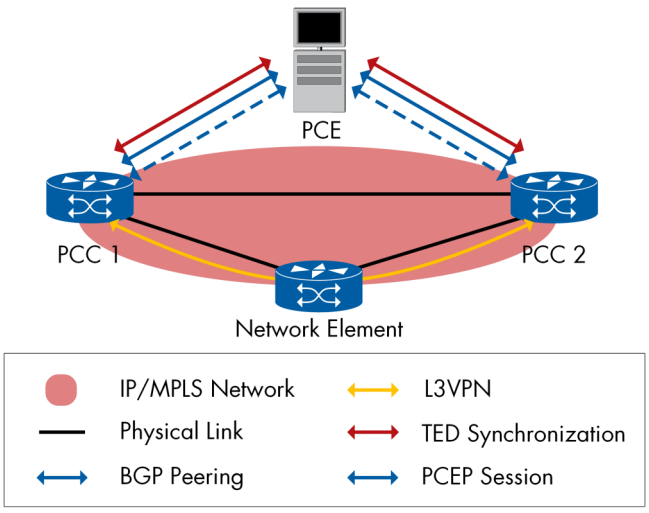

Figure 56: Test Topology—PCE Path Computation

Table 6 shows the combinations that passed the test successfully.

We faced two issues during this test. One combination was not successful because the PCC was not able to instantiate SR-TE LSPs via PCEP from PCE. The PCC signaled that the LSP was up but then immediately signaled it down again, causing PCEs to display the LSP is down on their CLI/GUI. Second issue was that some vendors did not support the PCEP either for PCE or PCC.

PCE Managing SR Policies via BGP SR-TE NLRI

Segment Routing allows a headend node to steer a packet flow along any path. Intermediate per-flow states are eliminated thanks to source routing.

In such a scenario the BGP possesses the capability to provide SR policy, and it can also give a candidate path and PCEP is not needed.

This test was performed using the SR-MPLS and SRv6. Although not all nodes from different vendors supported the SRv6 in this context, they participated as transport nodes.

The test was conducted as follows:

- BGP Sessions for Address Family SR-Policy were checked.
- VPN Routes for VRF configured on DUTs were checked.
- Verified BGP next-hop resolution and SR Policies for VRF Prefixes.
- Triggered creation of SR Policies on PCE and advertise it via BGP to PCCs.

The combinations that completed the test successfully with SR-MPLS are shown in Table 7. The combinations that passed the test successfully using SRv6 are shown in Table 8.

PCE	PCC 1	PCC 2	Network Element
Huawei NCE Path Computing Element	Juniper MX204	Nokia 7750 SR-1	ZTE ZXR10 M6000-8S Plus
Juniper Paragon Pathfinder	Nokia 7750 SR-1	Nokia 7750 SR-1	ZTE ZXR10 M6000-8S Plus
Nokia NSP	Juniper MX204	Juniper MX204	ZTE ZXR10 M6000-8S Plus

Table 6: PCE Path Computation

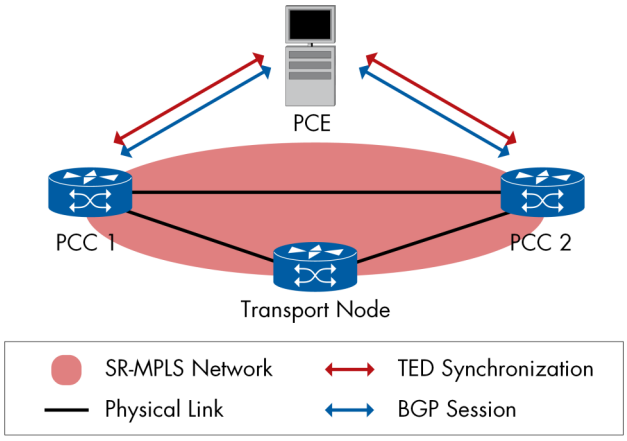

Figure 57: SR-MPLS Test Topology

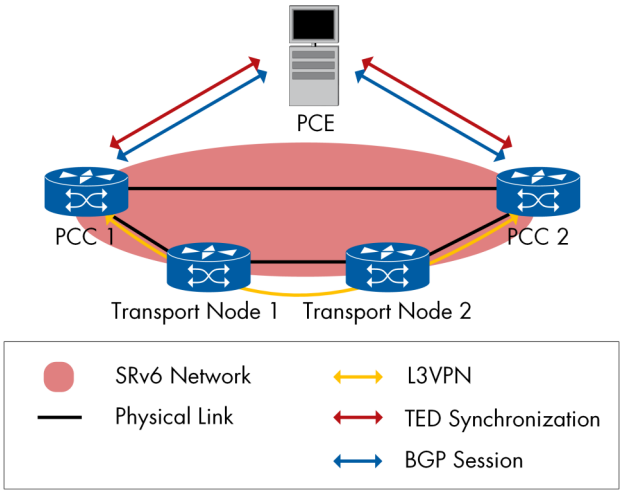

Figure 58: SRv6 Test Topology

PCE	PCC 1	PCC 2
Keysight IxNetwork	ZTE ZXR10 M6000-8S Plus	Huawei ATN 910D-A
Keysight IxNetwork	ZTE ZXR10 M6000-8S Plus	Nokia 7750 SR-1
Keysight IxNetwork	Juniper MX204	ZTE ZXR10 M6000-8S Plus
ZTE ZENIC ONE	Juniper MX204	Nokia 7750 SR-1
ZTE ZENIC ONE	Huawei ATN 910D-A	Nokia 7750 SR-1
Huawei NCE Path Computing Element	ZTE ZXR10 M6000-8S Plus	Juniper MX204
Huawei NCE Path Computing Element	Juniper MX204	Nokia 7750 SR-1
Huawei NCE Path Computing Element	Nokia 7750 SR-1	ZTE ZXR10 M6000-8S Plus
Nokia NSP	Huawei ATN 910D-A	ZTE ZXR10 M6000-8S Plus
Nokia NSP	ZTE ZXR10 M6000-8S Plus	Juniper MX204

Table 7: PCE Managing SR Policies via BGP SR-TE NLRI—SR-MPLS

PCE	PCC 1	PCC 2	Transport Node 1	Transport Node 2
ZTE ZENIC ONE	Huawei ATN 910D-A	ZTE ZXR10 M6000-8S Plus	Nokia 7750 SR-1	Juniper MX204
Huawei NCE Path Computing Element	ZTE ZXR10 M6000-8S Plus	Huawei ATN 910D-A	Nokia 7750 SR-1	Juniper MX204
Keysight IxNetwork	ZTE ZXR10 M6000-8S Plus	Huawei ATN 910D-A	Nokia 7750 SR-1	Juniper MX204

Table 8: PCE Managing SR Policies via BGP SR-TE NLRI—SRv6

Egress Peer Engineering with SDN Controller

The Segment Routing architecture can be directly applied to the MPLS data plane with no change on the forwarding plane. It requires a minor extension to the existing link-state routing protocols.

The SR-based BGP-EPE solution allows a centralized (Software Defined Network, SDN) controller to program any egress peer policy at ingress border routers or at hosts within the domain.

Thanks to the BGP-LS extension it is possible to export BGP peering node topology information (including its peers, interfaces and peering ASs) in a way that is exploitable in order to compute end-to-end SR-TE LSPs or SR-Policies where ingress and egress nodes are in different autonomous-systems.

In this test, we verified that EPE Segment Routing can be used to allocate MPLS labels for each engineered peer and use Label stack to steer traffic to a specific destination.

Table 9 shows the combinations that passed the test successfully.

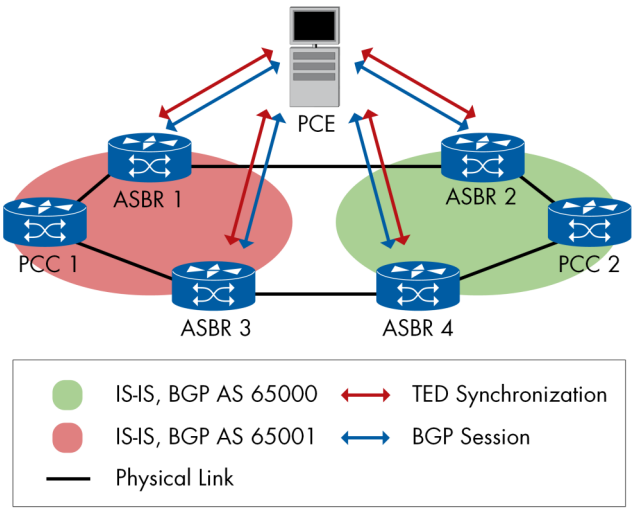

Figure 59: Test Topology—EPE

PCE	PCC 1	PCC 2	ASBR 1	ASBR 2	ASBR 3	ASBR 4
Huawei NCE Path Computing Element	Keysight IxNetwork	Keysight IxNetwork	Huawei ATN 910D-A	Juniper MX204	Nokia 7750 SR-1	ZTE ZXR10 M6000-8S Plus
Juniper Paragon Pathfinder	Keysight IxNetwork	Keysight IxNetwork	Huawei ATN 910D-A	Juniper MX204	Nokia 7750 SR-1	ZTE ZXR10 M6000-8S Plus
Nokia NSP	Keysight IxNetwork	Keysight IxNetwork	Huawei ATN 910D-A	Juniper MX204	Nokia 7750 SR-1	ZTE ZXR10 M6000-8S Plus
ZTE ZENICONE	Keysight IxNetwork	Keysight IxNetwork	Huawei ATN 910D-A	Juniper MX204	Nokia 7750 SR-1	ZTE ZXR10 M6000-8S Plus

Table 9: Egress Peer Engineering with SDN Controller

Clock Synchronization

Over the last decade and more it has become obvious that there is an exponential growth in the demand for fast, reliable and secure information both for the business community and that of the individual. This has driven, and continues to drive, the deployment of telecommunications networks with greater and greater capacity to deliver data to the consumer at a much higher throughput than ever before. This in turn has the driven the deployment of 5G networks with much higher performance requirements being placed on the network equipment. As a result of these requirements there is less and less tolerance for error within both the end to end network itself and that of individual network elements.

With this in mind the test scenarios designed and executed during the EANTC SDN Interoperability Event have provided accurate and measured data on the performance of multi-vendor network elements whilst performing either as part of a networks of inter-linked elements or in a standalone configuration. Using highly accurate measurement analysis equipment within the test set ups has allowed EANTC to effectively perform complex test scenarios which realistically reflect the type of conditions that could occur when such equipment is deployed.

Special attention has been placed on ensuring that the Class C and Class D network devices meet the performance requirements as defined in the current ITU-T standards. In addition the resilience of this equipment to the stresses caused due to loss of primary and secondary references has been closely measured. The deployment of Fronthaul based network architectures has also been addressed in testing the impact of the use of the Flex-E transport layer between devices for both frequency and timing distribution.

Phase/Time Partial Timing Support

Partial timing support within networks has been increasing over a number of years as the requirement to provide a cost effective Phase/Time solution that is more applicable to non-greenfield deployments.

This test was performed using the ITU-T G.8275.2 profile (PTP telecom profile for Phase/Time of day synchronization within partial timing networks) between the Grandmaster and the Boundary Clock. Since there could be no reliance using the traditional methods (e.g. SyncE), the Grandmaster clock was provided with a GPS input. In turn, both the slave and boundary clock were started from a free running state.

A Calnex Paragon-X measurement analyzer/tool was used to emulate a network load on input to the Boundary Clock using the application of the G.8261 test case 12 PDV profile.

For this test case, we had the involvement of three vendors namely Arista, Huawei, and Microchip.

In turn, the Calnex Paragon-t was used to accurately measure the Time of Day output from the Slave Clock against the defined limits.

The following devices successfully participated:

- BC: Arista DCS -7280
- GM: Huawei ATN910D-A
- SC: Microchip TimeProvider 4100
- Impairment tool: Calnex Paragon-X
- Phase Analyzer: Calnex Paragon-t

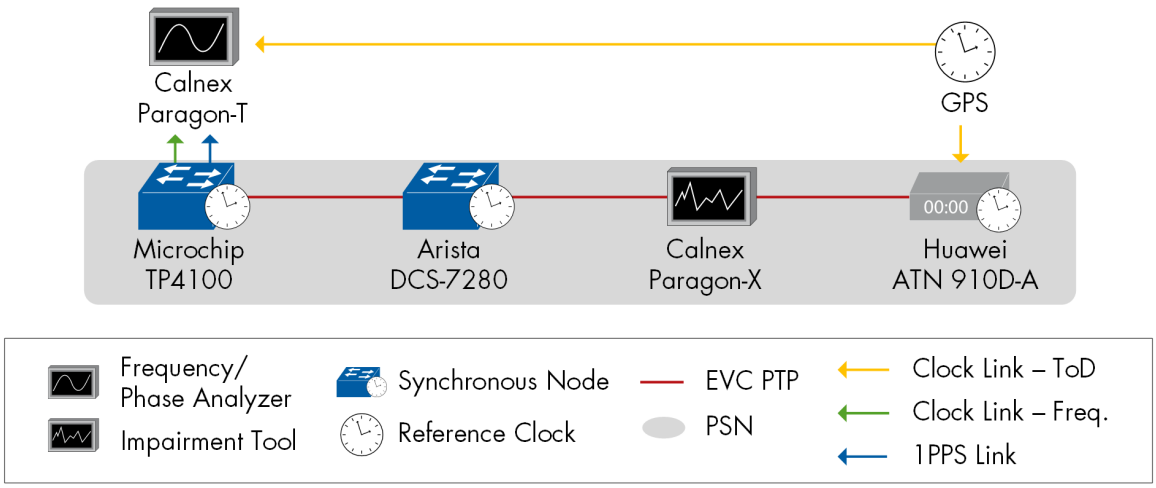

Figure 60: Phase/Time Partial Timing Support

High-Precision Clocking Source Failover

In any network, ensuring the resiliency of the time synchronization is paramount. To achieve this slave clocks and boundary clocks are provided with primary and secondary Grandmasters. This test exercised the boundary clock best master clock algorithm (BMCA) and confirmed switchover within required limits could be achieved using a lightweight topology comprising of a single boundary clock and 2 Grandmasters.

Both Grandmasters were locked to a GPS signal. The signal to the primary Grandmaster was then disabled within the Grandmaster to emulate loss of signal thus degrading the quality of the Grandmaster and causing the boundary clock to transition to the better source of time on its secondary Grandmaster. The resultant transient response was measured to ensure it fell within the expected performance limits. The G.8275.1 Telecom profile was used within the test with SyncE was enabled across the chain.

The depicted combinations all pass G.8271 level 6 accuracy.

- BC: Juniper MX240/MPC10E, Juniper ACX5448-M
- Primary GM: Microchip TimeProvider 4100, Calnex Paragon-neo (emulated Master)
- Secondary GM: Huawei ATN910D-A, Calnex Paragon-neo (emulated Master)
- Phase Analyzer Calnex Paragon-t
- Reference Clock: Microchip TimeProvider 4100

One vendor's Boundary Clock only actively exchanged PTP messages to the Grandmaster to which it was locked, and only generated a 1pps output when locked to a Grandmaster. It listened to other Grandmasters but would not send messages to the other Grandmasters until it lost lock to its primary. When it switched to another active master it stopped the 1pps output, so captures showed a period of circa 30s of dropped pulses, but this did not affect the Time Error produced.

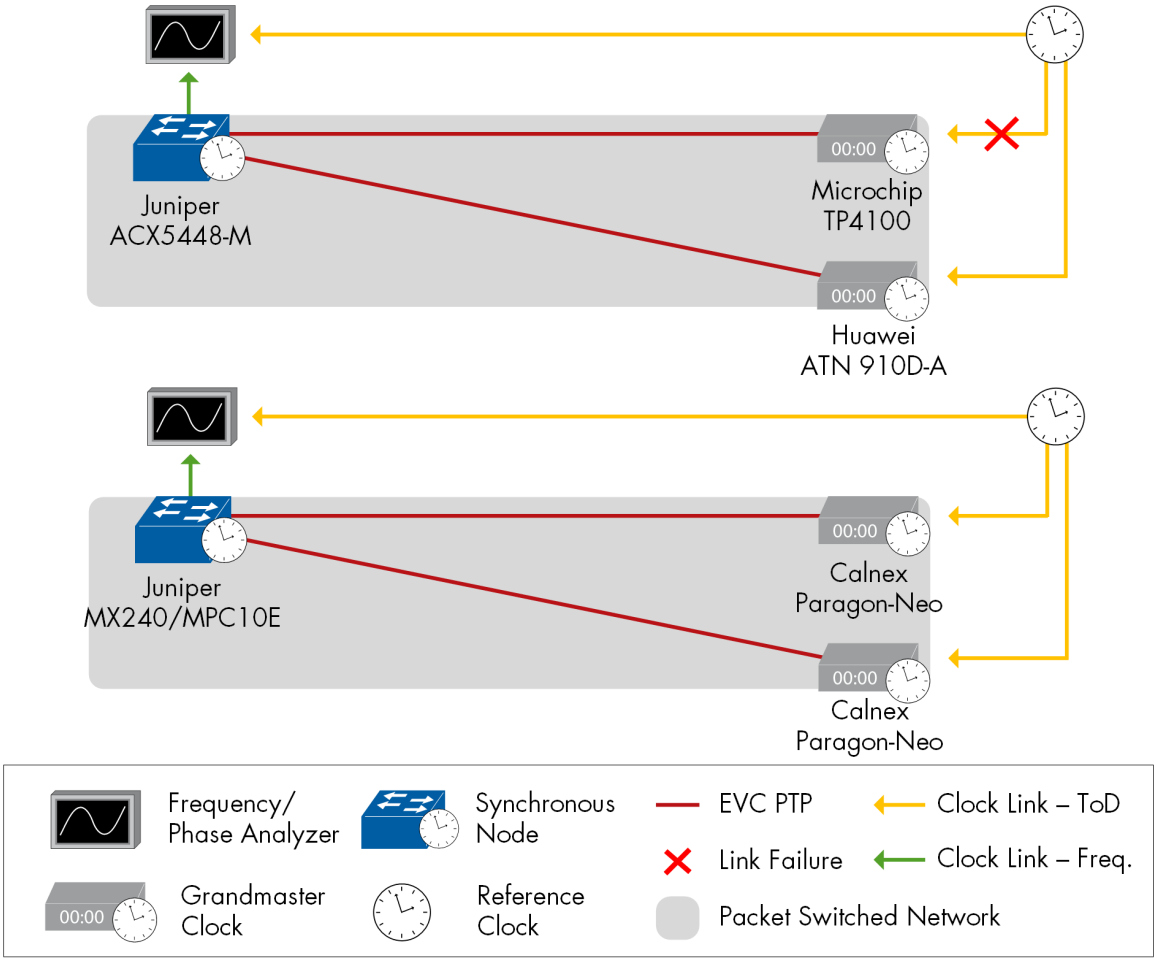

Figure 61: High-Precision Clocking Source Failover

Phase/Time Synchronization Source Failover

The IEEE 1588v2 standard mandates methods for ensuring reliable time delivery, by enabling the use of clock source redundancy with primary and secondary Grandmasters provided in a clocking system. Such a system allows for a chain of boundary clocks and slaves which all require accurate timing. Using the BMCA the boundary clock selects which of the two Grandmasters provides the best clock quality, selecting it as the primary clock. Using this selected timing path the boundary clock used PTP to deliver this best time to its associated slaves.

This test setup exercises the real life resiliency using two Grandmasters, a Boundary Clock and a Slave Clock. Initially, the Boundary clock was locked to the Grandmaster designated as the primary Grandmaster.

The Slave Clock was then allowed to lock to the Boundary Clock. We then caused a degradation in the clock quality on the primary Grandmaster by disconnecting its GPS interface.

For Paragon-neo as Grandmaster, we simulated loss of GPS by manually degrading the clock Class and ESMC-QL values to ensure both references switched.

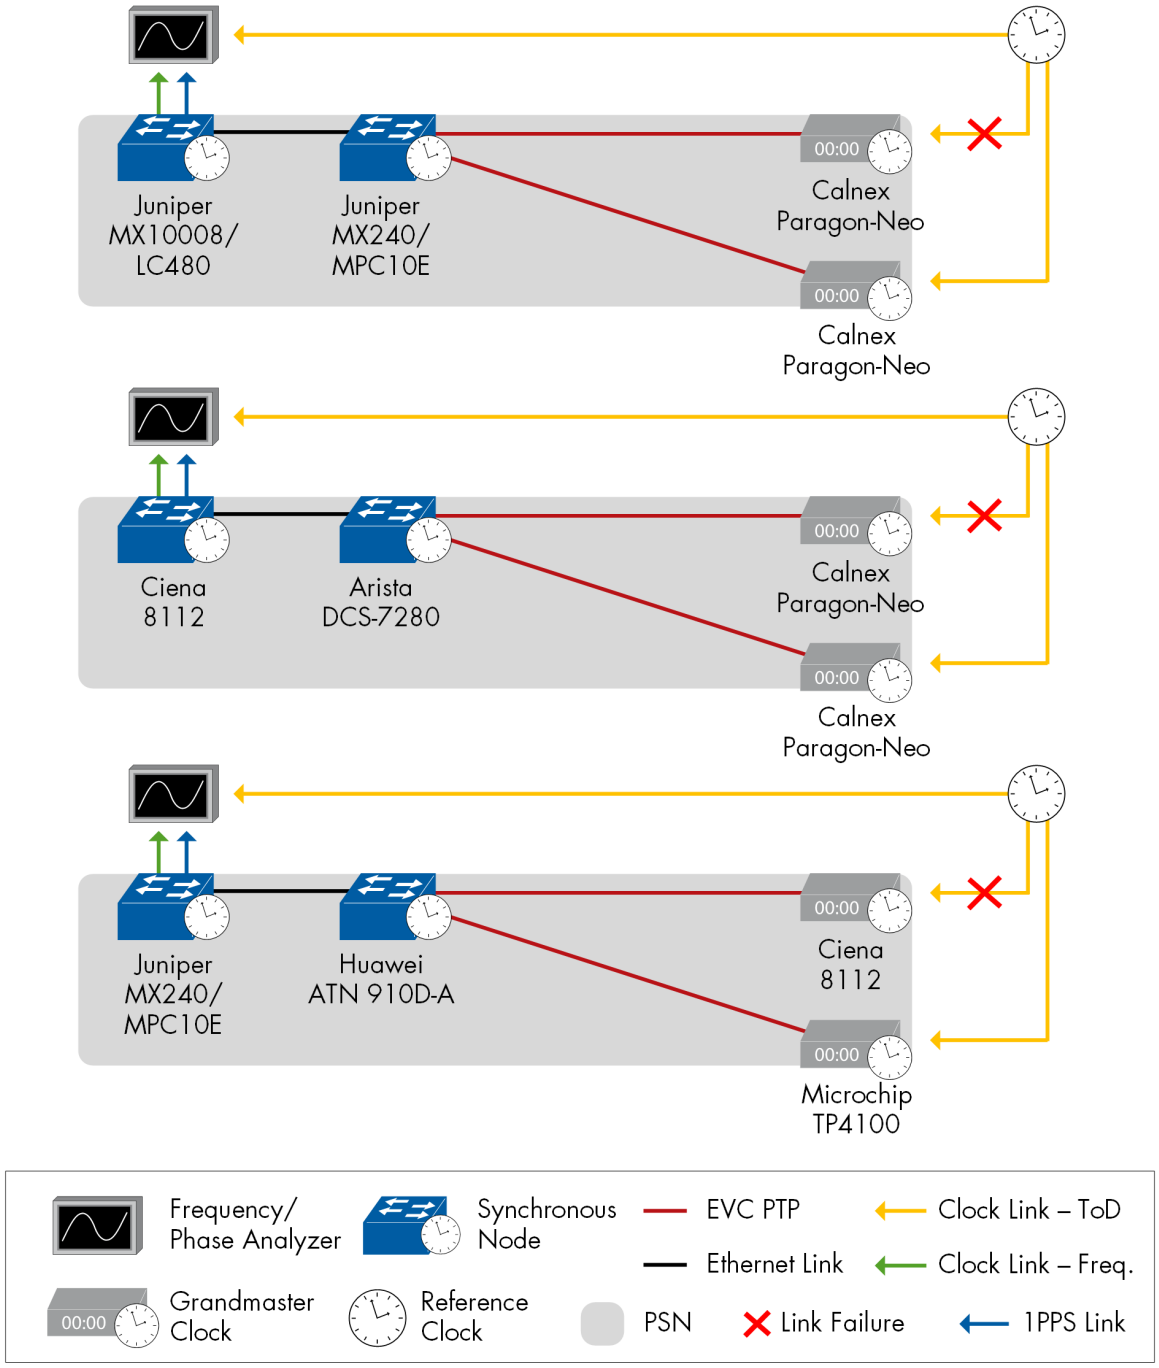

Figure 62: Phase/Time Synchronization Source Failover

We verified that clock quality had dropped below that of the secondary Grandmaster and that the boundary clock swapped its clock synchronization to the secondary Grandmaster. While doing so the slave clock's transient response was measured.

This test was performed using the ITU-T G.8275.1 profile between all of the network elements, while Sync-E was transmitted to ensure frequency lock.

The measurements made in this test were carefully calibrated to ensure accuracy including cable delays between the antenna and network equipment and between the slave output and the measurement equipment in the form of the Calnex Paragon-neo where it was acting as a Grandmaster or the Calnex Paragon-t when the Paragon-neo was not involved in the test. Testing was carried out at both 10GbE and 100GbE interface rates.

All combinations achieved G.8271 Level 6 accuracy.

- BC: Arista DCS-7280, Huawei ATN910D-A, Juniper MX240/MPC10E
- Primary GM: Ciena 8112, Calnex Paragon-neo (Master emulator)
- Secondary GM: Microchip TimeProvider 4100, Calnex Paragon-neo (Master emulator)
- SC: Ciena 8112, Juniper MX240/MPC10E, Juniper MX10008/LC480
- Phase Analyzer Calnex Paragon-neo
- Reference Clock: Microchip TimeProvider 4100

All steps passed except step 3 in one of the configurations, where limits and mask were slightly exceeded when one vendor's slave lost PTP lock to the Boundary Clock whilst the Boundary Clock was attempting to lock to the Grandmaster. Within 30s this had stabilized again.

Phase/Time Synchronization Degradation of Primary Source—Measuring the Effect of Source Failover to Secondary Source

The purpose of this test is to ensure that a Boundary Clock can maintain its phase/time synchronization when it loses its GPS connection and switches to use another GPS-led source of PTP.

In this test, the Boundary Clock was locked to GPS as its primary source and received PTP from another Boundary Clock connected to a GPS connected Grandmaster.

The test was performed using the G.8275.1 Telecom Profile, with the devices configured to use it with SyncE frequency reference in hybrid mode.

During the execution of the test, the performance when using its primary reference was recorded and measured against G.8271 Accuracy Level 4 limits. We then disconnected that primary source such that Boundary Clock switched to use its secondary source, that of the PTP flow from the other Boundary Clock.

Performance during this transition whilst acquiring lock and performance once locked were both measured and compared against the limits defined in the G.8271 ITU-T standard using either the Calnex Paragon-t measurement analyzer or the Calnex Paragon-neo Measurement Analyzer.

In addition to its role as Measurement Analyzer, the Paragon-neo was also used in test configurations as an emulated PTP Master.

Successful combinations in this test included the following equipment:

- BC: Arista DCS-7280, Juniper MX10008/LC40, Microchip TimeProvider TP4100
- BC/SC: Ciena 8112
- GM: Huawei ATN910D-A, Calnex Paragon-neo (emulated Master)

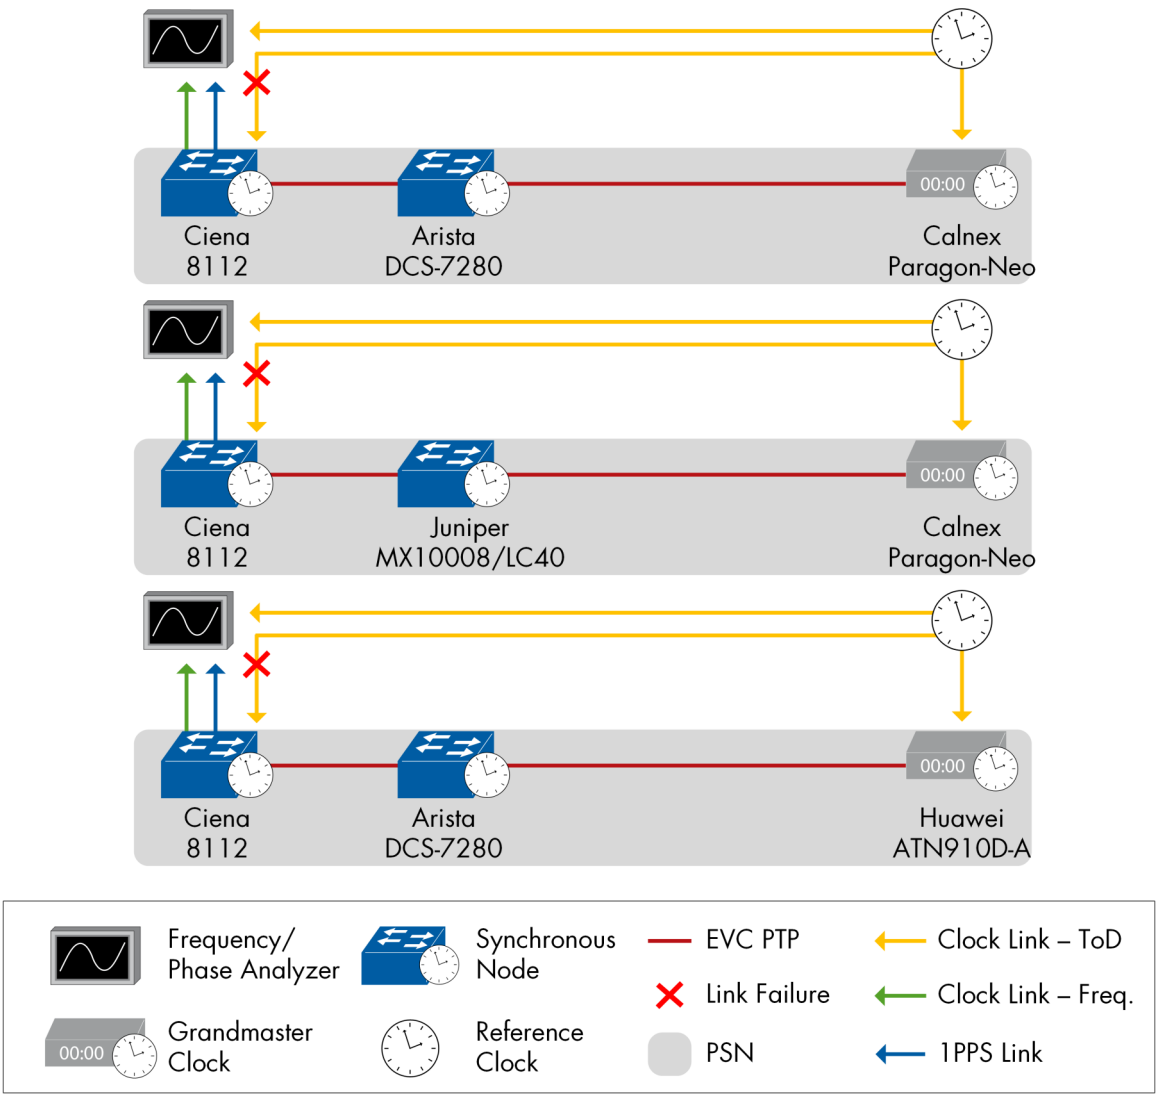

Figure 63: Phase/Time Synchronization Degradation of Primary Source

Phase/Time Full Timing Support over MACsec

The goal of the test was to verify that the timing synchronization quality of a Slave Clock can be maintained when using MACsec between a Boundary Clock and the Slave Clock based on the ITU-T G.8275.1 Telecom Profile.

Originally packet networks were not deemed to be timing-sensitive and as a result the effect on the accuracy of clocks was not considered when provisioning additional security within packet networks. The expectation is that the provision of such network security should not cause problems when deploying security protocols like MACsec. This is especially important as 5G deployments increasingly make use of such security measures.

MACsec is a recommendation based on threat models defined in RFC7384 "Security Requirements of Time Protocols in Packet Switched Networks".

The test procedure involved enabling the MACsec between the Boundary Clock and the Slave Clock and the measurement of the Slave Clock Time Error output. To simulate network asymmetry the Calnex Paragon-X applied an impairment based on the ITU-T G.8261 Test Case 12 whilst measuring the Time Error output of the SC. The test configuration passed G.8271 Accuracy Level 6.

Successful combinations:

- BC/SC: Juniper ACX5448-M
- GM: Calnex Paragon-X
- Impairment Tool: Calnex Paragon-X
- Measurement Analyzer: Calnex Paragon-X

Phase/ Time Full Timing Support: Boundary Clocks Class-C/D Test

With ever-tighter limits imposed on timing synchronization within modern networks, it is important to ensure that equipment used in these networks conforms to the performance limits defined for those networks.

This test aims to verify that a slave clock can maintain its synchronization quality when using the G.8275.1 Telecom profile and SyncE in hybrid mode with a chain of Class D Boundary Clocks that themselves conform to the performance limits defined for boundary clocks in the G.8273.2 Standard.

The latest revision of this standard includes two new high-accuracy clocks (Class C and D) that are subject to tighter performance constraints than existing Class A and B clocks. To ensure that the performance of the Slave Clock in this test configuration matched that of a real-world scenario, the test involved a serie of source failover events.

Such events are used to stress the ability of the Slave clock to cope with such switching and ensure that its performance is not adversely affected, enabling the network to maintain the required performance.

The test used a chain of Clock Class D boundary clocks coupled to two Grandmasters to provide the stimulus to this test. We involved multiple combinations of such configurations to ensure that a valid mix of devices is tested. In each case, we used the Calnex Paragon-t measurement analyzer to accurately analyze the Slave Time Error Performance. All of the test configurations met the Class 6A performance limits.

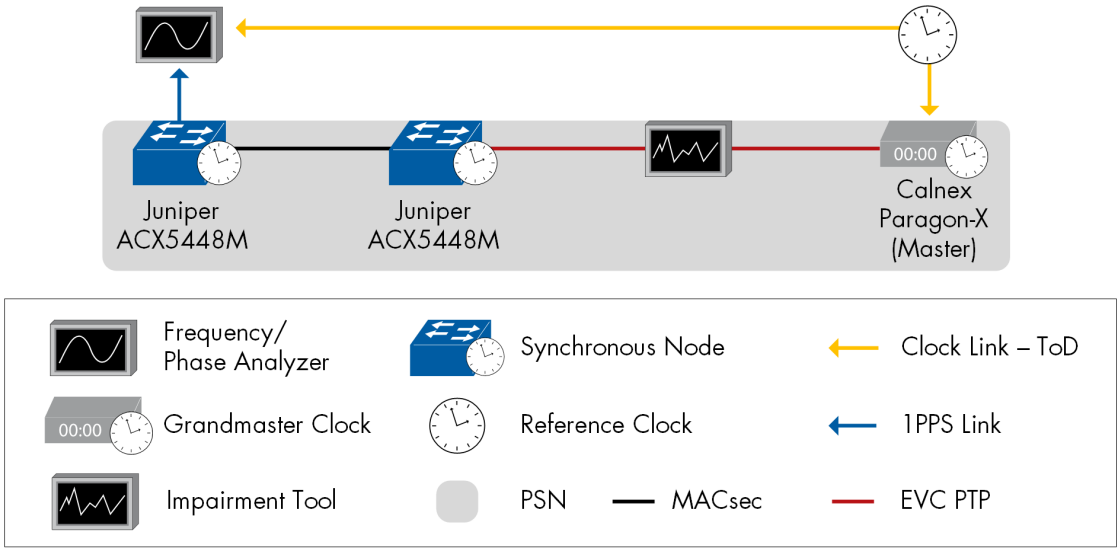

Figure 64: : Phase/Time Full Timing Support over MACsec

The following DUTs successfully participated in the test, as

- T-BC: Arista DCS-7280, Huawei ATN910D-A, Juniper ACX7100-32C
- GM: Calnex Paragon-neo (emulated Master), Ciena 8112, Microchip TimeProvider TP4100
- T-TSC: Ciena 8112, Juniper ACX7100-48L

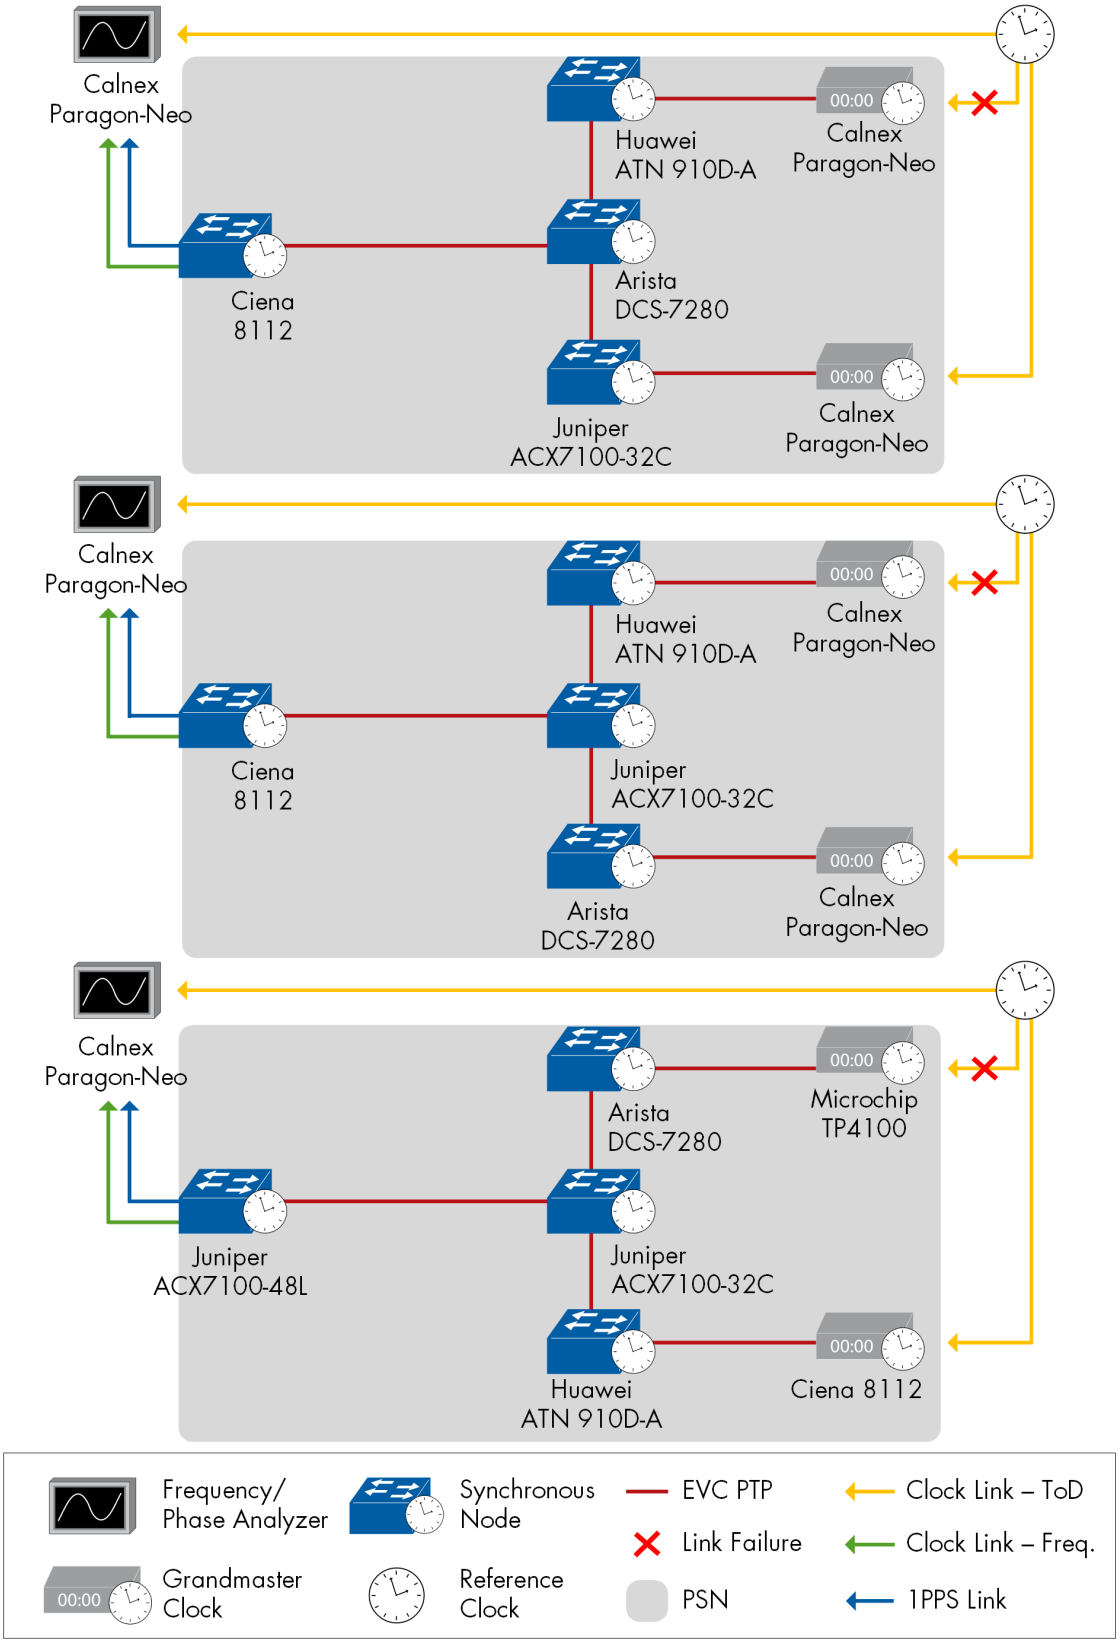

Figure 65: Phase/Time Full Timing Support Boundary Clocks Class-CD Test

Phase/Time Synchronization over FlexE

Maintaining Phase performance when using FlexE to ensure that the impact of using FlexE as a transport container does not impact the inherent timing performance of a Class 6 C/D Boundary clock.

The team used the Calnex Paragon-neo test and measurement instrument in this test to emulate a PTP Master and Slave and to accurately measure the Time Error output of the Boundary Clock to ITU-T G.8273.2 Standard Class C/D limits.

The test used the G.8275.1 ITU-T Telecom Profile carried over 100GbE links using a FlexE transport.

The Boundary Clock was connected to the Paragon-neo Master and Slave and configured to acquire frequency and PTP. PTP was started on the Paragon-neo Master and Slave with SyncE (QL-PRC) generated at the Paragon-neo Master. Once we attained lock at the Boundary Clock, a Time Error measurement was performed on the Paragon-neo for 1000s.

The resulting capture was then analyzed using the Paragon-neo analysis tool to examine the Time Error output of the Boundary Clock.

The 2-way Time Error value was subjected to the G.8273.2 T-BC/T-TSC limits for a Class D clock.

The following DUTs successfully participated in the test, as:

- T-BC: Huawei NE8000 M14
- GM: Calnex Paragon-neo (emulated Master)
- SC: Calnex Paragon-neo (emulated Slave)

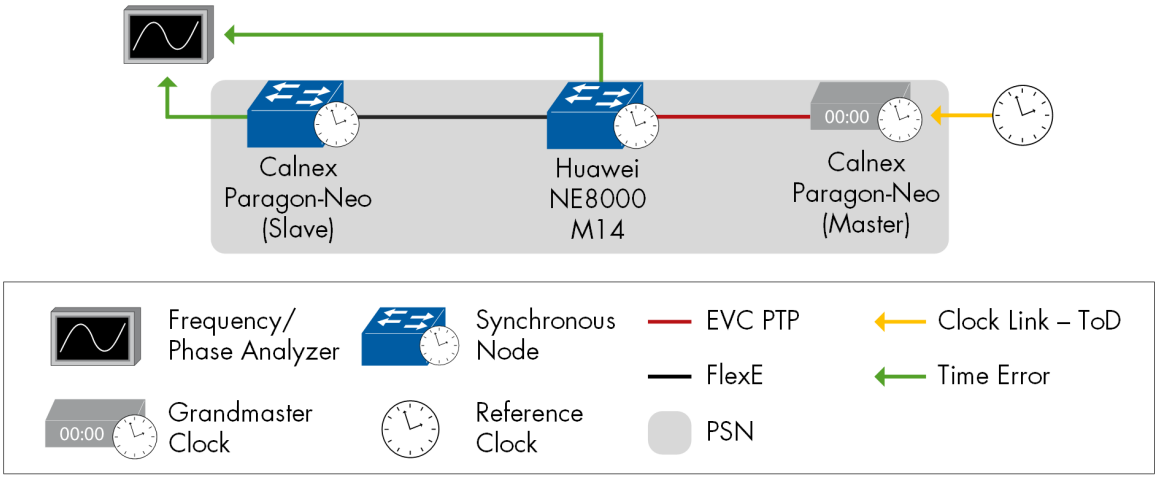

Figure 66: Phase/Time Synchronization over FlexE

Conformance Test Boundary Clock Class C/D—Measuring to G.8273.2 Standards

The migration to 5G has imposed ever-stricter limits on synchronization timing in networks. The ITU-T G.8273.2 T-BC/T-TSC Timing Characteristics standard has introduced tighter limits for devices operating in this environment. This test is designed to determine whether or not vendors' T-Boundary Clock devices conform to these limits.

We used the Calnex Paragon-neo test and measurement instrument in this test to emulate a PTP Master and Slave and to accurately measure the Time Error output of the Boundary Clock to ITU-T G.8273.2 T-BC limits for Class C/D clocks.

The test used the G.8275.1 ITU-T Telecom Profile carried over 1GbE, 10GbE, 25GbE or 100GbE links while simultaneously sending Sync-E in hybrid mode.

The Boundary Clock was connected to the Paragon-neo Master and Slave and configured to acquire frequency and PTP. PTP was started on the Paragon-neo Master and Slave with SyncE (QL-PRC) being generated at the Paragon-neo Master.

Once we attained the lock at the Boundary Clock, a Time Error measurement was performed on the Paragon-neo for 1000s. The resulting capture was then analyzed using the Paragon-neo analysis tool to examine the Time Error output of the Boundary Clock. The 2way Time Error value was subjected to the G.8273.2 T-BC Clock Class D limits.

The following DUTs successfully participated in the test, as

- BC: Arista DCS-7280SR3E, Ciena 8112, Huawei ATN910D-A, Juniper ACX7100-32C, Juniper ACX7100-48L, Microchip TimeProvider TP4100
- GM/SC: Calnex Paragon-neo

There were no issues at 10GbE, 25GbE, or 100GbE, with all devices passing Clock Class D limits.

DUT	GbE
Arista DCS-7280SR3E	25GbE
Arista DCS-7280SR3E	100GbE
Ciena 8112	100GbE
Huawei ATN910D-A	100GbE
Juniper ACX7100-32C	10GbE
Juniper ACX7100-32C	100GbE
Juniper ACX7100-48L	10GbE
Juniper ACX7100-48L	100GbE
Microchip TimeProvider TP4100	10GbE

Table 10: Conformance Test Boundary Clock Class D—DUT, GbE

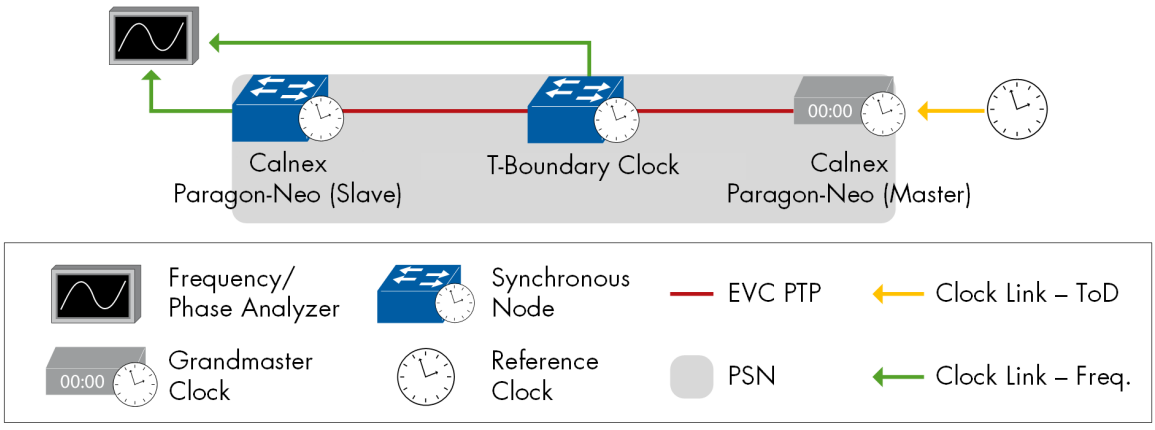

Figure 67: Conformance Test Boundary Clock Class C/D

FlexE

As interoperability for the new technologies and standards are the core of our event, and the scope of our interests includes Datacenter-interconnection, Network Slicing, and moving to 5G setups in the service providers, it's mandatory to test FlexE. FlexE eliminates the one-to-one mapping between the physical interface and the MAC layer, which provides flexibility with using the available bandwidth.

This year we had three main test scenarios for FlexE applications, regarding channelization, FlexE bonding, and Dynamic Bandwidth Adjustment.

FlexE Channelization and Physical Isolation

FlexE allows the channelization usage of the physical links, offering flexibility for service providers to provide client services with different bandwidth in one or more physical links. Users can configure several FlexE Tunnels with different bandwidth for different client services.

In this test, we deployed several FlexE Tunnels to carry different client services and verified the channelization in 2x100G ports between different vendors.

We also verified the physical isolation of different FlexE tunnels by increasing the traffic in one of the FlexE tunnels.

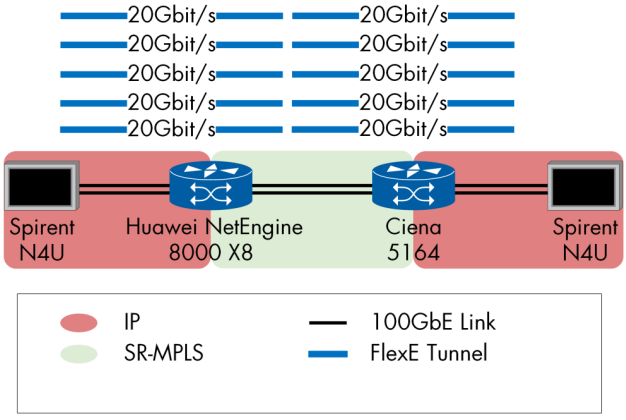

Figure 68: FlexE Channelization and Physical Isolation

FlexE Bonding

Major advantage of FlexE technology, is the ability of bonding multiple links to utilize their bandwidth supporting services requiring higher bandwidth.

The main difference between the traditional Link Aggregation solutions and the FlexE, is that FlexE allows the services to utilize the full bandwidth of the links that are bonded while the LAG utilizes 70-80% of a link under typical conditions.

In this test, we verified the FlexE bonding capability by bonding two 100G ports to a 200G FlexE interface. We also verified the link usage by sending 200Gbit/s traffic from the traffic generator.

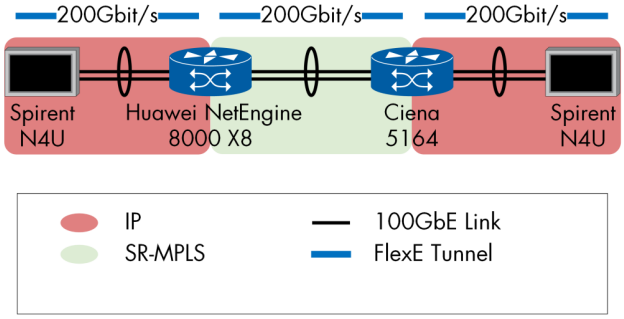

Figure 69: FlexE Bonding

FlexE Dynamic Bandwidth Adjustment

FlexE provides the flexibility of adjusting the client service bandwidth without going on-site physically to switch the physical interface connection. When it comes to the connection between a router and optical transport equipment, service providers can adjust the service bandwidth more efficiently based on the actual requirement of the client service.

In this test, we verified the FlexE capability of dynamic bandwidth adjustment on a 100G interface.

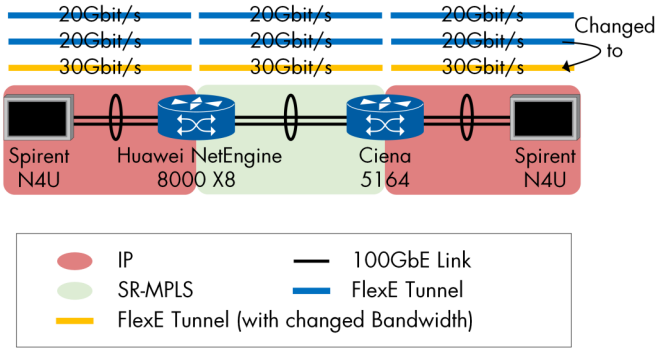

Figure 70: FlexE Dynamic Bandwidth Adjustment

Open RAN Fronthaul Verification

For the first time, EANTC has included an actual mobile use case scenario as an application-layer demo in our test this year. We have tested the readiness of participating vendors' equipment for 5G Open RAN fronthaul connectivity. This has been possible through a collaboration with the **i14y Lab**, a joint research project with Deutsche Telekom and other partners in Berlin. (Please see the i14y-lab.com website for more details.) This project has contributed an O-RAN Radio Unit (O-RU) from Foxconn and an O-RAN Distributed/Central Unit (O-DU/O-CU) from Radisys to verify whether fronthaul connectivity requirements are met.

As it is well-known in the industry, 5G is a fundamental game changer. It will enable more applications through the high data-rate communication, low latency, and massive scale of connectivity. To meet such unprecedented industry demands, all aspects of 5G implementations from core to radio access network (RAN) need to be multi-vendor ready. 5G New Radio (5G-NR) is required to support a much increased scale of radio components due to higher frequency and smaller cell coverage than LTE.

To address the scale and sourcing diversity challenges, mobile operators have founded the O-RAN ALLIANCE with the goal to standardize disaggregated radio access networks. The overarching aim is to reshape the RAN industry towards more intelligent, open, virtualized, and fully interoperable deployments.

There are three main open disaggregated RAN components according to the O-RAN architecture: O-RAN Radio Unit (O-RU), O-RAN Distributed Unit (O-DU), and O-RAN Centralized Unit (O-CU).

The network connection between O-RU and O-DU is called "open fronthaul". There have always been fronthaul connections in the industry (known as CPRI and eCPRI), but they have typically been internal between components of the same radio vendor previously.

When O-RUs and O-DUs from different vendors are connected, standardization is getting crucial. Thus, the O-RAN Alliance has specified an Open Fronthaull (OpenFH) interface. During our tests, we followed the latest version 05.

We focused on the real-time aspects of the Synchronization Plane (S-Plane) and Control/User Plane (CU-Plane). The non-real-time Management Plane (M-Plane) is out of our scope this year.

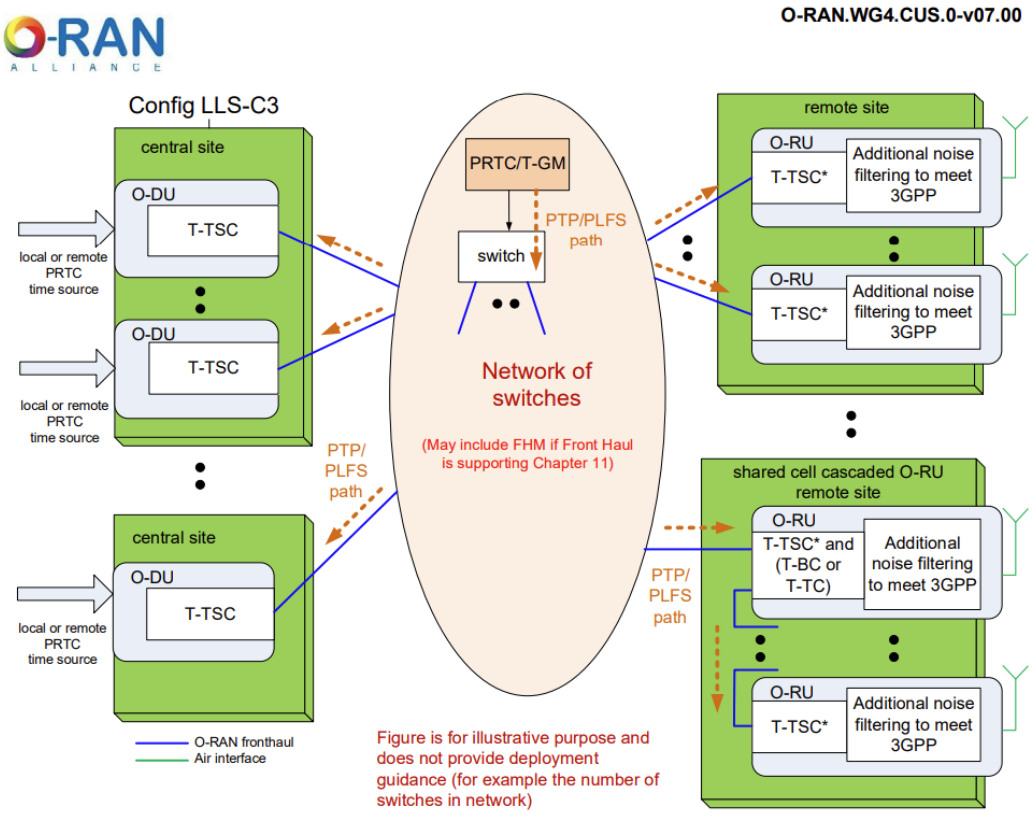

Figure 71: O-RAN Test Topology

On CU-Plane, we verified evolved Common Public Radio Interface (eCPRI) which is unlocked to the vendors and the hardware.

Finally, we verified the interop game-changer. The front haul transport devices in a multi-vendor environment are represented by Calnex, Foxconn, Juniper, Microchip, ng4T, and Radisys.

O-RAN Test Areas

The O-RAN OpenFH Interop tests covered:

- PTP time synchronization using ITU-T G.8275.1 profile in LLS-C3: Functional test and performance test
- eCPRI functional test on CU-Plane

Both O-RU and O-DU work as an ITU-T G.8273.2 Telecom Time Slave Clock (T-TSC); the fronthaul switch works as a Telecom Boundary Clock (T-BC), and one external clock input source is needed as a Telecom Grand Master (T-GM).

We verified the recommended PTP Full Timing Support profile - ITU-T G.8275.1 in which the PTP transport is directly over L2 Ethernet. The testing of the optional PTP Partial Timing Support profile - ITU-T G.8275.2 was out of the scope of this testing, in which the PTP transport is over UDP/IP.

Functional Test of O-DU + Bridged Network + O-RU using ITU-T G.8275.1 Profile (LLS-C3)

As the initial setup test case, we configured the S-Plane only. The S-Plane would typically be enabled before any CU-Plane connection once the O-RU and O-DU become ready for service provisioning. Both the O-RU and the O-DU must be able to synchronize their clocks with the device under test acting as a boundary clock (T-BC). During the testing, the grandmaster clock (T-GM) was connected to the GNSS satellites as the input clock source.

This test verified stable time synchronization via the fronthaul switch of both O-RU and O-DU.

Successful combinations in this test included the following equipment:

- GM: Microchip TimeProvider TP4100
- OpenFH switch Juniper AXC710
- Slave clocks: Foxconn O-RU, Radisys O-DU

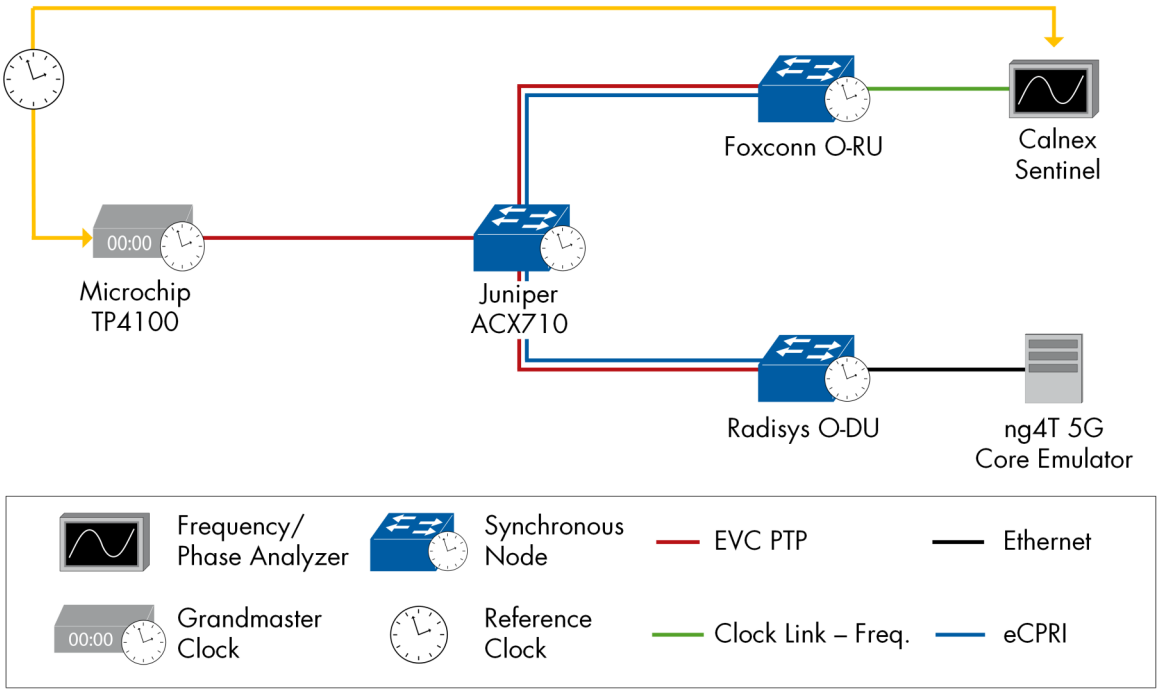

Figure 72: O-RAN EANTC Test Topology

Radio Layer 3 C-Plane Establishment and Initial Radio U-Plane Data Transfer

This is a Radio Frequency (RF) startup use case. After S-Plane functionally works fine, the O-RU would like to start his CU-Plane service, and then wait for the downlink reference signaling from the O-DU. And then, the O-DU would like to start his CU-Plane service for the eCPRI connection. At the same time, the O-CU should work fine and set up the connection to the 5G Core, so that RAN would be ready for reference in the air interface.

We measured the eCPRI packets on the OpenFH switch under the test; the L1 traffic on both O-RU and O-DU; and the air signaling on Calnex Sentinel.

Successful combinations in this test included the following equipment:

- Remote Radio Head (RRH): Foxconn O-RU RPQN7800
- OpenFH switch: Juniper AXC710
- Baseband unit (BBU): Radisys O-DU/O-CU
- Air interface Measurer: Calnex Sentinel Unit
- 5G Core: ng4T 5G Core Emulator

Performance Test of O-DU + Bridged Network + O-RU using ITU-T G.8275.1 Profile (LLS-C3)

On the control plane, signaling data is sent towards the air interface. The performance of the fronthaul connection must be in the acceptable range following O-RAN specification, so that user equipment (UEs) can successfully maintain stable cell connectivity for any coming uplink and/or downlink traffic. We measured the time error (TE) on Calnex Sentinel.

Successful combinations in this test included the following equipment:

- GM: Microchip TimeProvider TP4100
- OpenFH switch (i.e. Boundary Clock): Juniper AXC710
- SC: Foxconn O-RU, Radisys O-DU
- Air interface Measurer: Calnex Sentinel Unit
- 5G Core: ng4T 5G Core Emulator

Test Tools and Emulators Used For O–RAN Testing

To facilitate all tests of the 5G O-RAN use case scenario, the following vendors supported the scenario in addition to the participating vendors:

- Foxconn RPQN7800— O-RAN-standard indoor TDD 5G NR O-RU
- Radisys TDD 5G NR O-DU and O-CU
- ng4T CNF-based Core emulator for both EPC and 5GC

Calnex provided the Sentinel specifically for this test area —a dedicated synchronization tester with 5G Over-the-Air measurements. Special thanks go to Calnex for supporting our events with excellent clock testing equipment.

NETCONF Interoperability Test

It's 2022, the modern Service Provider networks have been and are still being affected by two main challenges. 5G large implementations roll out which are getting hotter in the industry with all the elements needed (management, network slicing, network virtualization, ...). The other challenge is the pandemic which has been affecting the world for two years now and heavily impacted the internet and rocketed its traffic. These are the main reasons that led Service Providers to enhance, extend and improve their infrastructures and services. One of the most important aspects of improving the networks is the management, the whole industry moving towards automated and vendor-agnostic models of Network Management. Here comes the NETCONF as a configuration protocol with YANG data modeling language to form a combination of a management protocol that enables a programable interface to the networking devices from different vendors.

NETCONF offers to facilitate configuration data management and interoperability between devices from different vendors and reduces network faults caused by manual configuration errors. That being said, the automation implementations needs to be validated in realistic environments before they can be entrusted mission critical tasks. Vendor-specific YANG models certainly help network operators to fully and autonomously configure their network fabrics through YANG programmatically. But, mapping the management intent to devices in a multi-vendor network is still difficult because of the diverse management models used in different devices.

When the OpenConfig consortium created their YANG model, that increased the similarity between multiple vendors' management interfaces. Today, various hardware vendors support OpenConfig based management interfaces, but as we noticed from last years testing and this year as well, most OpenConfig implementations still have a long way to go for field usable coverage and interoperability.

Many service providers would be very happy of if the OpenConfig YANG model was the solution everyone agreed to utilize. Device vendors sense this and happily advertise support in their devices. With most vendors, the actual engineering hours are still overwhelmingly spent on the proprietary (native) models. One vendor, however, has demonstrated that the OpenConfig model can indeed be implemented in a highly usable manner.

During our EANTC NETCONF Interoperability tests series, we verified the interoperability between SDN controllers and network devices in a multi-vendor environment represented by Ciena, Cisco, Huawei, Juniper, and Nokia.

This year we performed multiple tests in different areas which represent the bottleneck of service providers' networks management.

Through a careful process of discussing the needs of the service provider and their networks, we agreed with the participated vendors to perform new tests that include Access Control Lists, Routing Policies, and OAM (Operation, Administration, and Management). The new tests will help the service providers to understand the issues they might face while rolling out these configurations with NETCONF/YANG in multiple vendor environments, and to have a solid, better, and wider overview of the status of the interoperability between controllers and network elements from different vendors. Even though we have introduced new tests to perform, which weren't tested before, it was delightful to see how easy and successful the testing was, although we faced some issues; as expected and meant for this event, the great amount of support that we got from the participated vendors was crucial to be able to solve this issues immediately and continue testing.

Aside from the previously mentioned tests, we performed deploying Layer 3 VPN, Layer 2 VPN, MPLS, and Segment Routing MPLS with NETCONF/YANG, as these are essential services in the service providers' networks implementations. It was really great to see these services configurations and provisioning running quickly, and successfully without any major issue, depending on the experiences we gathered from the previous years' events.

Testing Remarks

- The newly introduced tests went smoothly with a minimum amount of inter-op issues.
- The repeated tests were performed directly without any issues, in some cases we have performed tests successfully even without any pre-testing.
- Ciena routers were equipped with YANG models using OpenConfig as core with added extensions, which was super convenient for the controllers to perform the tests smoothly.

- Controllers noticed differences between vendors' YANG models in terms of errors. For one vender they had to correct dozen small errors (mostly originated from the OpenConfig consortium) on the other hand they were able to find more than fifty errors for a different manufacturer.
- The Controllers teams were able to find also some issues in their software, which they corrected.
- One test was marked as failed which is configuring Bidirectional Forwarding Detection with NETCONF, during the test, when the controller pushed the configurations of the BFD session to the router, the router responded with "Malloc failed". This indicates an implementation problem, the responsible team took the issue for further troubleshooting. During the same test, we configured the BFD session manually and tried to delete it with NETCONF, when the deletion call was pushed by the controller, the router responded with "BFD Session Config entry deletion not allowed as an entry in use by bfd session". This is a clear indication that the NETCONF implementation isn't completely transactional in this area. The device should make sure the actions necessary for removing the BFD service are executed in a sequence that the various components can handle. As it is now, the client would have to break up the transaction into multiple calls, removing most of the value with NETCONF.
- A major difference between the controllers was how much transactional the NETCONF calls and the service or devices provisioning could be. During the tests, we noticed one controller was able to deliver a large chunk of basic configurations and provision services all in one transaction, while other controllers needed multiple transactions to deliver the same services.
- Although the controllers' software and capabilities are getting more advanced, we still notice different approaches regarding the NETCONF transaction implementations, and which operation to use when pushing the new configurations. This point led to multiple discussions during the testing about the NETCONF operations and which one is the correct to be used, with different interpretations to the standard.

Future Plans

The success of this testing campaign and the wonderful cooperation from our partners, vendors, and service providers is pushing us to continue this series of NETCONF/YANG interoperability testing events. The future aims for the testing would be concentrating more on the NETCONF itself, its basic operations, and consistency. New test areas will be touching the Network Monitoring, Advanced and Overlapping Services, besides 5G related tests.

Test Setup and Measurement Equipment

EANTC was responsible for setting up the test environment to be used during the test. All of the physical network elements were installed in the EANTC lab beside all the controllers were hosted in EANTC virtual VMware environment with the exception of the Huawei NCE controller which was installed in Huawei premises in France and connected through IPsec tunnels to EANC premises.

In collaboration with VMware, EANTC provided a VIO platform (VMware Integrated OpenStack) to host Cisco and Nokia's controllers and orchestrators. The platform has sufficient resources to host much more large-scale workloads.

EANTC created remote access accounts for all participants to reach their VMs in EANTC's lab to configure their devices and run the tests. We used a Wiki collaboration space to plan the test and document all results. EANTC provided the participants with links for digital meeting rooms to meet, discuss, and run the tests. The collective work requires technical abilities and the ability to adapt the work time and routines to reach the work goals.

Test Areas

The NETCONF Interop Tests 2022 were conducted in three different categorized areas of Networking Management, each area consisted of multiple tests that covered the needs and demands of management and running the modern networks: Device Provisioning, Service Provisioning, and OAM (Operation, Administration, and Management).

We successfully validated 37 test combinations, especially through hierarchical orchestration with the modular provisioning of multi-services. All results were expected and were based on OpenConfig models obtained by Ciena, Cisco, Huawei, Juniper, and Nokia.

Device Provisioning

One of the big advantages of NETCONF is how the protocol works when manipulating a group of related configuration data. NETCONF modifies all or selected parameters on a single primitive operation, and also allows configuration to occur in a transactional manner. NETCONF takes under consideration when a number of the network devices successfully upload the configuration, but others fail. In this case, NETCONF allows a managed device to rollback to a known-state configuration.

Management of IP Implementations

We verified managing IP interfaces with reading and writing operations, as well as updating information for IP changes.

The following systems successfully participated in the test: As controller: Cisco NSO, Huawei NCE, Nokia NSP. As PE router: Ciena 8112, Huawei NetEngine 8000 X8, Juniper MX204.

Management of Interfaces

These were the most basic elements included in the configuration provisioning test. We sent traffic for the created IP interface and verified the configuration change (MTU size or IP address) and the deletion through the controller.

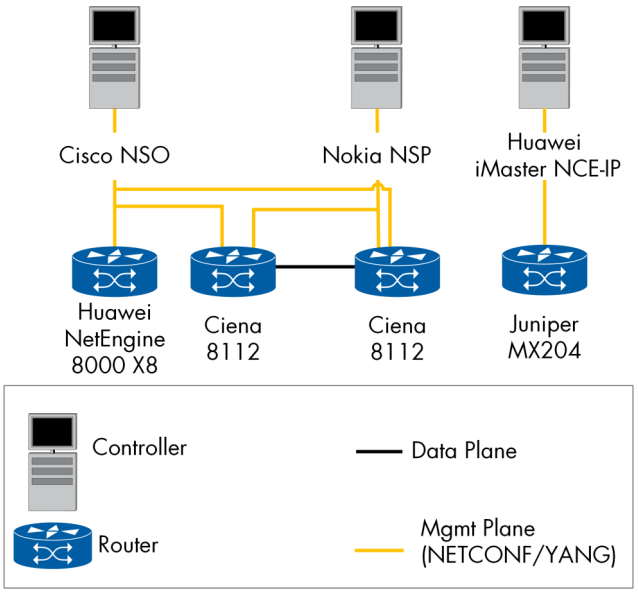

Figure 73: Management of Interfaces and IP Implementation

The following systems successfully participated in the test: As controller: Cisco NSO, Huawei NCE, Nokia NSP. As PE router: Ciena 8112, Huawei NetEngine 8000 X8, Juniper MX204.

Network Access Control Lists (ACLs)

Network Access Control Lists (ACLs) are an ordered-by -user set of rules, used to configure the forwarding behavior in the device. Each rule is used to find a match on a packet, and define actions that will be performed on the packet.

NETCONF uses two models for creating the ACL. One defines generic ACL aspects which are common to all ACLs regardless of their type or vendor, and the second defines the necessary groups for matching fields in the packet. In this test, we verified the creation of ACL on the client device and validate their proper performance. As the common use case of ACL is to achieve a level of security for network access by specifying which parts of the network/service can be accessed by a user and which cannot, we developed an access profile that had matches for an IPv4 destination address and an action of "deny" (<forwarding> drop</forwarding>).

The controller pushed the configuration through NETCONF and we confirmed the specific IP address is in fact unreachable. In the second step, we updated the ACL to include a second match to an IPv4 source address and action of allow. This time previous destination address can be reached only through the specific source address. The following systems successfully participated in the test: Cisco NSO as controller, and Ciena 8112 as PE router.

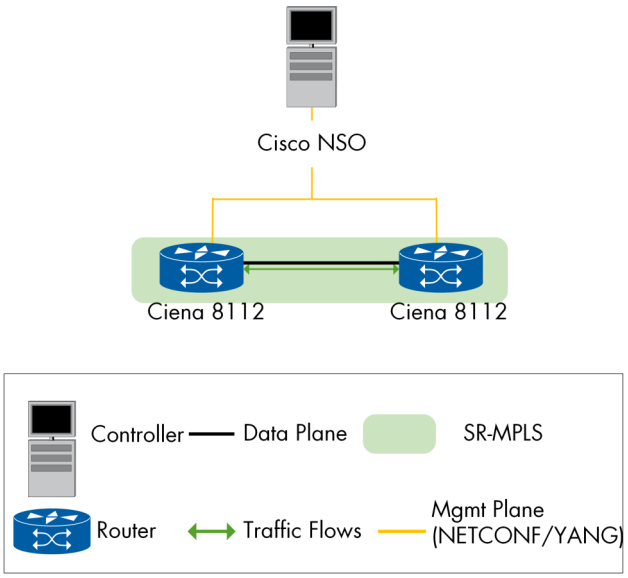

Figure 74: Network Access Control Lists

While trying a different combination for this test, one controller used nc:operation:replace command to apply the configuration instead of nc:operation: merge, and the router didn't support that .It was possible to have a workaround for this but then we decided to skip and not to include it in the report.

Service Provisioning

In general, service provisioning refers to the creation of the service as well as the management of service-related data. New requirements for quick and error-free service turn-up are posing a challenge to network providers. Existing configuration management approaches, such as CLI scripting, and device-specific adapters, are unable to meet these new requirements. NETCONF greatly simplifies device and service configuration management while maintaining high performance.

MPLS Using NETCONF

It looks like the MPLS market is still in demand even with the new rising competition of others technologies (like SD-WAN). With what this technology offers of reliability in delivering packets and high-quality service, service providers will always be required to maintain, implement and manage MPLS networks

So it was only natural to test MPLS with NETCONF this year as well. We verified MPLS configuration based on proprietary YANG models, and we used the MPLS transport as part of the L2/L3VPN service provision test.

OpenConfig defined a data model to configure the Label Distribution Protocol (LDP) and LSP-specific parameters.

This test verified a YANG model that can be used to configure and manage Label Distribution Protocol global and LSP-specific parameters for IGP-congruent LSPs.

The following systems successfully participated in the test: As controller: Cisco NSO, Huawei NCE, Nokia NSP. As PE router: Ciena 8112 and Juniper MX204.

Huawei NetEngine 8000 X8 was manually configured (not provisioned by NETCONF) only to showcase service.

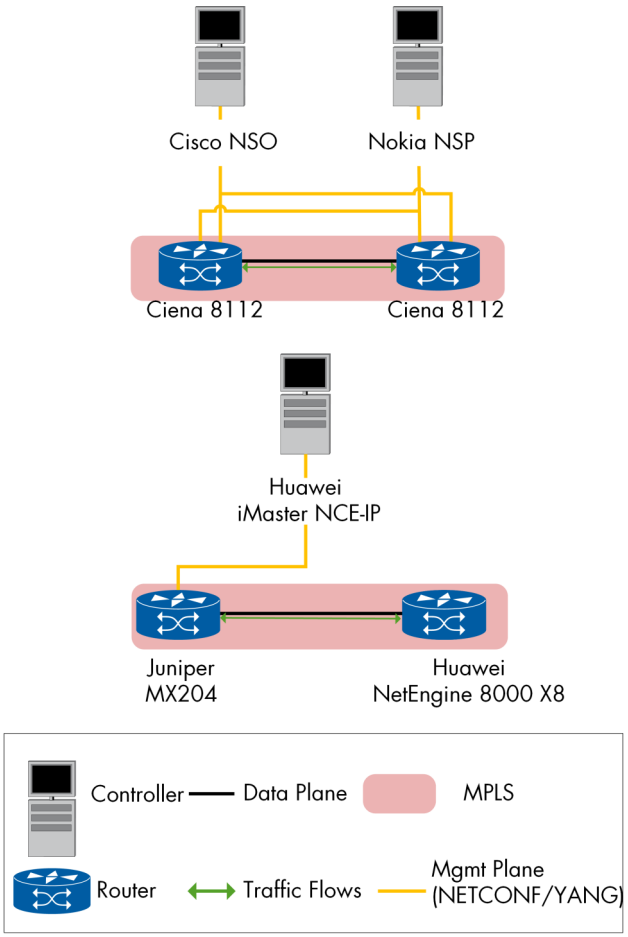

Figure 75: MPLS using NETCONF

Segment Routing MPLS Using NETCONF/YANG

Segment routing as a source-based routing works in a complementary way with MPLS. SR-MPLS aims to simplify networks by assisting service providers in completing service-driven network transformation through less complicated protocols.

This test verified MPLS Segment Routing (SR-MPLS) based on proprietary YANG models. The SR-MPLS transport was part of the L2/L3VPN service provision test. The controller pushed the configuration to the DUTs and we confirmed the labels were learned through SR-ISIS in the MPLS table and checked the operability of MPLS segment routing label-switched path (LSP) connections added by ISIS protocol through performing ping command.

The following systems successfully participated in the test: As controller: Cisco NSO, Huawei NCE, Nokia NSP. As PE router: Ciena 8112 and Juniper MX204. Huawei NetEngine 8000 X8 was manually configured (not provisioned by NETCONF) only to showcase service.

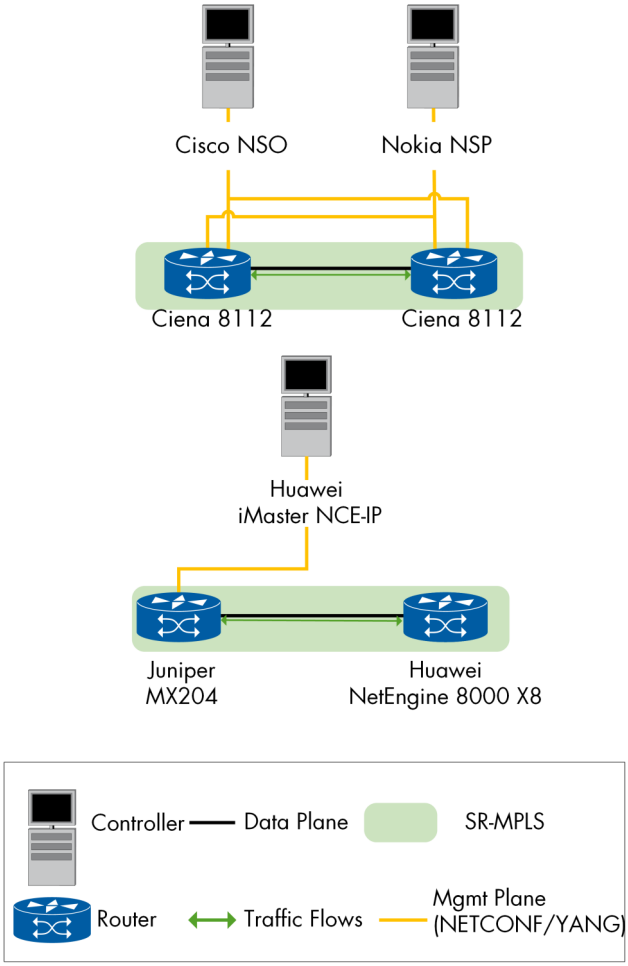

Figure 76: Segment Routing MPLS Using NETCONF/YANG

L3VPN Service Provisioning

The L3VPN technology is a core service every service provider router should offer. We confirmed using NETCONF creating VRF while configuring BGP as a distributer to VPN routing information, and SR-MPLS as a transport network.

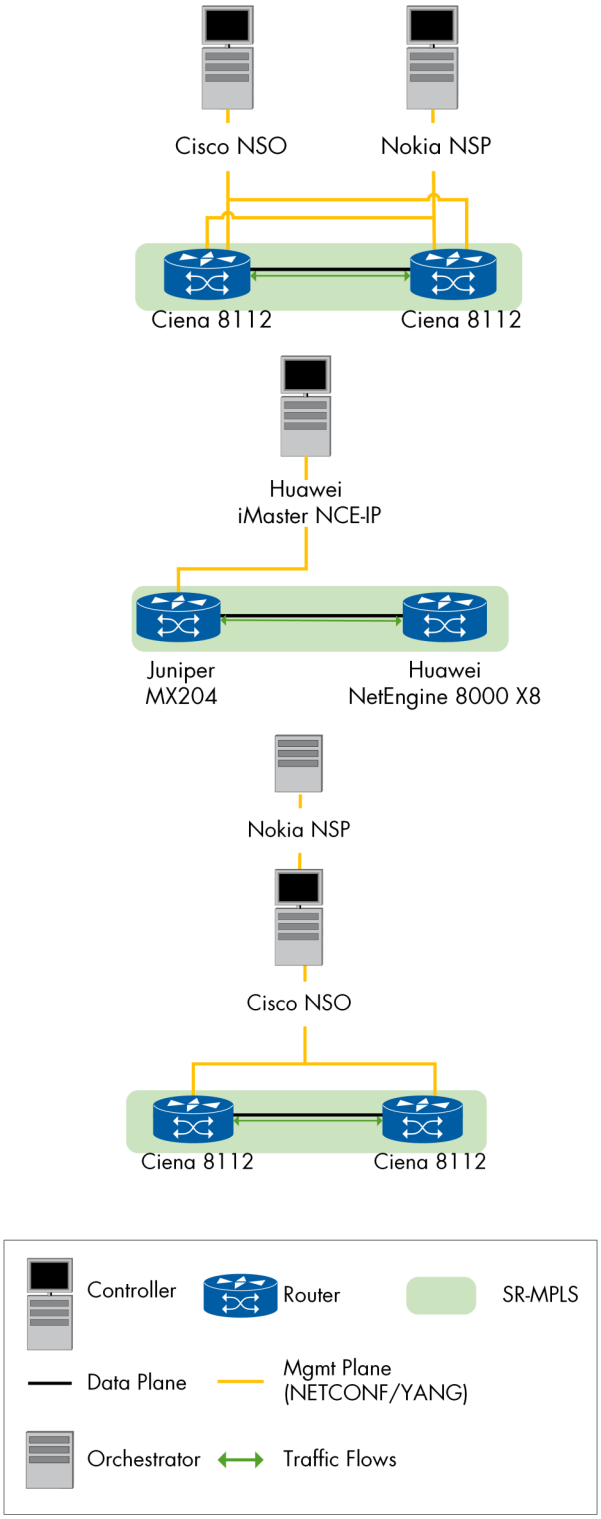

Figure 77: L3VPN Service Provisioning

We verified a hierarchal setup that includes an orchestrator beside the basic topology. There were no L3VPN services on any of the routers and had a 100% drop of the traffic generated between them. The controller pushed the VRF configurations (name, route target, route-distinguisher, AFI, Safi, peer address, and AS number) to the devices and we confirmed the right configurations were applied. Later the traffic was forwarded through the newly created VPN service with no loss.

The following systems successfully participated in the hierarchical orchestration test: Cisco NSO as controller, Nokia NSP acted as orchestrator and Ciena 8112 as PE router. The transport network was based on the SR-MPLS.

The following systems successfully participated in the test: As controller: Cisco NSO, Huawei NCE, Nokia NSP. As PE router: Ciena 8112 and Juniper MX204. Huawei NetEngine 8000 X8 was manually configured (not provisioned by NETCONF) only to showcase service.

Layer 2 EVPN Service Provisioning

We verified EVPN VPWS service EVPN VPWS service provisioning using the standard YANG model northbound in a hierarchical orchestration setup. All test procedures followed the same method as described in the L3VPN service provision with an L2VPN service instance.

The following systems successfully participated in the hierarchical orchestration test: Cisco NSO as controller, Nokia NSP acted as orchestrator and Ciena 5164 as PE router. The transport network was based on the SR-MPLS. The following systems successfully participated in the test:

As controller: Cisco NSO, Huawei NCE, Nokia NSP. As PE router: Ciena 8112, Juniper MX204.

Huawei NetEngine 8000 X8 was manually configured (not provisioned by NETCONF) only to showcase service.

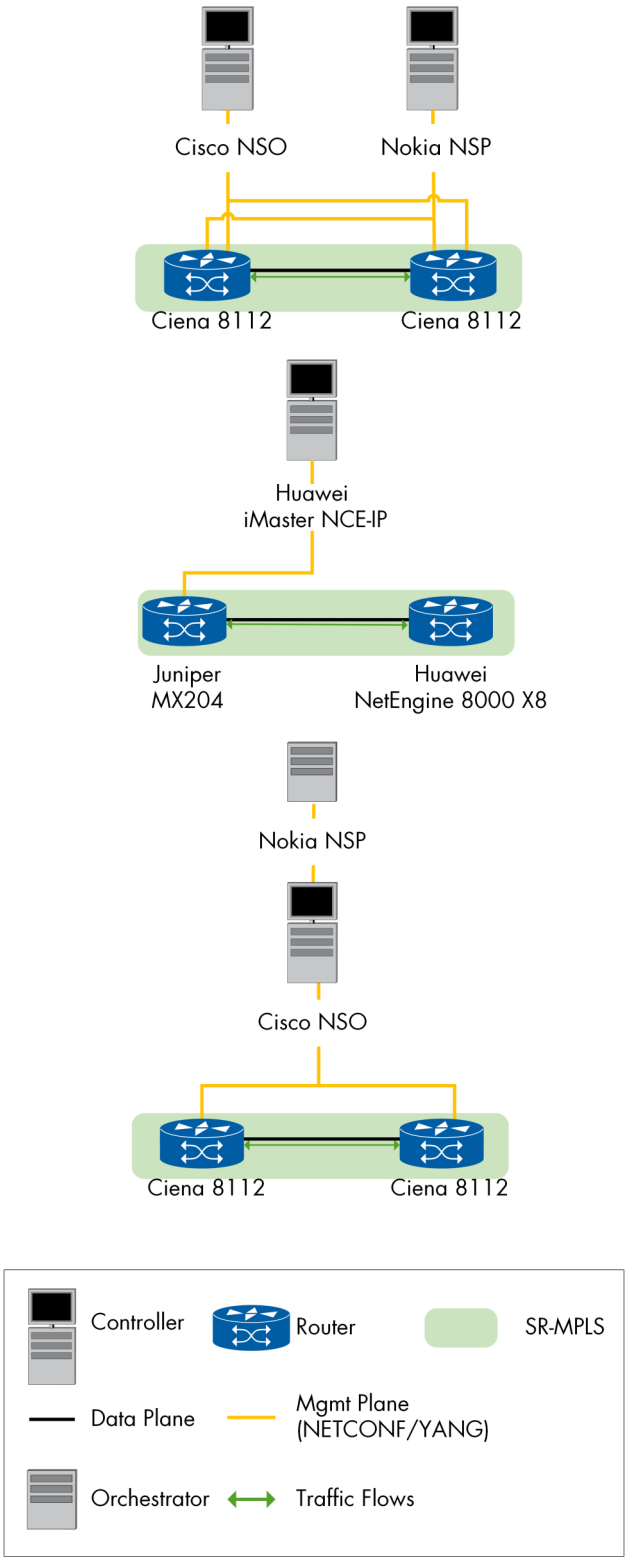

Figure 78: Layer 2 EVPN Service Provisioning

Routing Policies Configuration

A routing policy instructs the router to inspect routes, filter them, and potentially modify their characteristics.

It adds a dynamic feature to the routing table provided by the routing protocol because it sets the rules for importing and exporting from the routing tables and also can manipulate the route attributes.

All that allows the router to control what routes to advertise to its neighbors. A routing policy defines the conditions to use to match a route and the action to perform on the route when a match occurs.

NETCONF/YANG combinations allow the Service Providers to create, edit, or remove routing policies on their edge and all needed devices. The advantage here is not only the easy remotely configuration but also the quick response to the needed changes in case of urgent changes needed.

In this test, we verified the creation of the policy over the client device and we observed the parameters which were administered.

First, the router had BGP protocol configured through the NETCONF with AS number and address family, then the routing policy was applied to deprioritize a route by prepending the AS-PATH attribute. We confirmed that the BGP routing table on the DUT was showing the correct values that were set by the routing policy. As the last step, the controller rolled back the configurations successfully from the routers.

The following systems successfully participated in the test: Cisco NSO as controller, Nokia NSP acted as controller and Ciena 8112 as PE router.

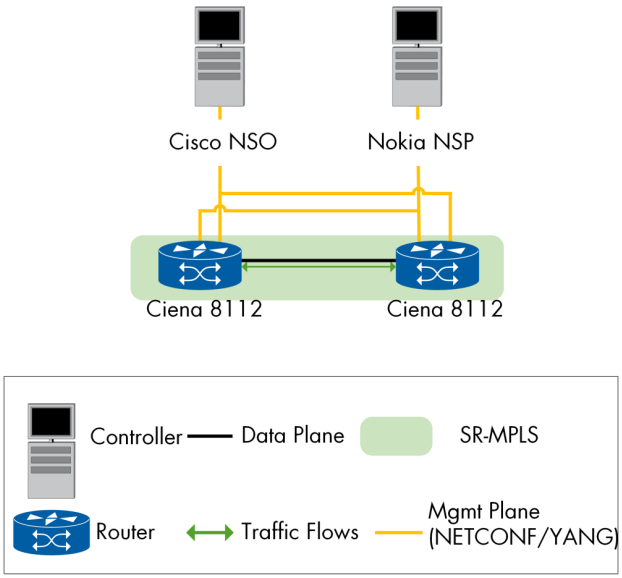

Figure 79: Routing Policies Configuration

Precision Time Protocol Configuration

PTP plays a central role regarding the synchronization and time distribution among networks, especially when it comes to the 5G timing requirements and Time Error strict values. NETCONF gives the chance to configure the PTP profiles and configure the interfaces between multiple devices remotely, quickly and provision the configurations to a large number of devices avoiding the misconfiguration or the errors of configuring this service manually.

This test verified the PTP configuration via the proprietary Ciena YANG model. Before starting the test, none of the clock statistics were shown via CLI on the DUT. We configured a PTP profile based on ITU-T G.8275.1 for unicast through the controller towards the network device. We used the device intern clock from one of the PE as a time source and observed the synchronization status, which changed from free-running to locked on another PE. PTP counters appeared and showed PTP packets exchanged between both PEs

The following systems successfully participated in the test: Cisco NSO as controller, and Ciena 8112 as PE router.

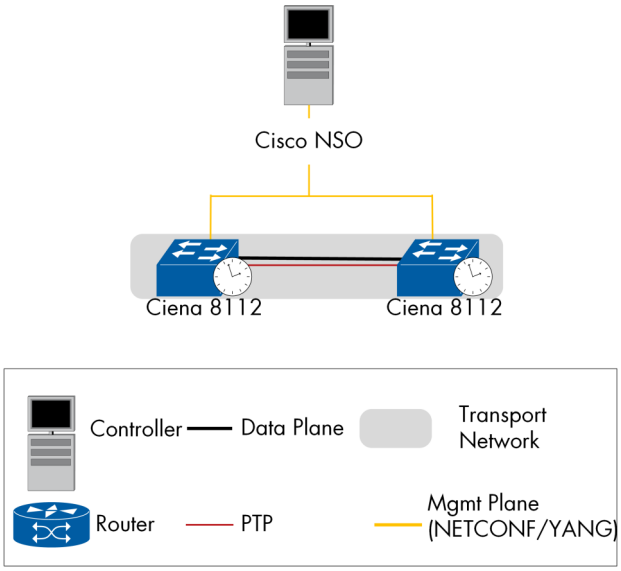

Figure 80: PTP

Operations, Administration, and Maintenance

Operations, Administration, and Maintenance (OAM) are important networking functions that allow operators to monitor networks connections, troubleshoot failures and monitor performance.

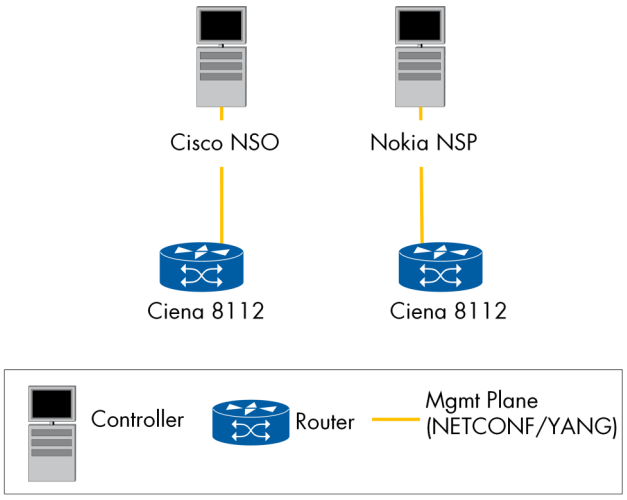

Figure 81: OAM—Test Setup

Proactive OAM

Proactive OAM refers to OAM actions that are carried out continuously to permit proactive reporting of fault.

We used the SOAM (Service OAM) to measure important metrics for Service Level Agreement (SLA) such as availability and delay. This is based on two standards first is the IEEE 802.1Q which defines the Connectivity Fault Management (CFM) capability

and the second one is ITU-T Y.1731 which defines both fault management and performance monitoring.

To enable connectivity fault management the controller configured a maintenance domain and association, and by creating maintenance association endpoints (MEPs) we could generate and respond to the connectivity fault protocol messages.

We set the priority for the CFM messages and lowest priority defect that generates a fault alarm and the delay times of the alarms as well. For the delay measurement, the configuration included frame size, delay intervals, remote-mep-id and others.

The controller was able to collect the live status of the delay measurement between the MEPs.

The following systems successfully participated in the test: Cisco NSO as controller, Nokia NSP acted as controller, and Ciena 8112 as PE router.

On-Demand OAM

Ping and Traceroute are well-known f connectionless-oam-methods used for fault verification and isolation, respectively, for IP networks.

Over the years, different technologies have developed similar toolsets for equivalent purposes.

Since it is the first troubleshooting command in networking, we ran a test to verify the continuity check by issuing a ping command from the controller to the router interfaces.

We tested with IPv4 address ping and also through the VRF tunnel with no reported issues.

The following systems successfully participated in the test: Cisco NSO as controller, Nokia NSP acted as controller, and Ciena 8112 as PE router.

Retrieve Interface Frame Sizes Distribution

Understanding the network and the traffic profiles are major factors in delivering better performance and forwarding capabilities. One tool for this is checking the packet size distribution of the network on any router or interface.

Retrieving this information has many use cases. Perhaps the most essential one is detecting a potential DOS attack using small packets going through routers.

This test verified the read capabilities of the packet counter with OpenConfig YANG models. Once the session was established successfully the controller would be able to retrieve the interface statistics for the frame distribution.

We generated traffic using IMIX packet size and observed the same distribution of frame sizes on the interface as expected.

The following systems successfully participated in the test: Cisco NSO as controller, and Ciena 8112 as PE router.

System Inventory

We verified the network device's system information with reading and writing operations.

Through NETCONF we displayed the device information of location, chassis id, number, etc. We pushed changes to the device and confirmed the change on the router.

We also returned the intended configuration after deleting them directly from the router and observed the controller rollback the old configuration.

The following systems successfully participated in the test: Cisco NSO and Huawei NCE as controller, Ciena 8112 and Juniper MX204 as PE router.

Hardware Management

We used this test to confirm the collection of the hardware information and counter statistics from the client device.

We were able to identify common properties of the hardware components like vendor data (name, part number, serial number, manufacture date) and diagnostics properties such as temperature and voltage supply.

Telemetry sensor also was observed after specifying the subset of the data that we want to stream from the router using sensor paths.

The following systems successfully participated in the test: Cisco NSO as controller, and Ciena 8112 as PE router.

NETCONF Consistency

Our contacts in the service providers have reported some problems regarding the consistency of the NETCONF transactions and responses from the network elements when the NETCONF messages contain some deviation or unexpected values. The reported issues contained some CLI freeze or device reboot, which was the motivation for this test.

We observed the response and behavior of network elements upon receiving illegal NETCONF injections commands and verifying the device's capability to report the error conveniently.

According to RFC 6241, all NETCONF messages must be well XML formatted to be readable by the client. Upon receiving a NETCONF message, the client evaluates the parameters and either executes the command or, in case of an incorrect expression, an error should be raised.

NETCONF engines (like NSO and NSP) don't allow sending invalid values in the first place, so in order to test the router's response to such cases, we used different tools.

One was the netconf-console to send a file with an illegal value. the other was using raw SSH Client to add the required message.

When the router received a mismatched value to a parameter, like a character or invalid value to MTU, it replied with a "bad-value" error.

But when we sent a message containing binary sequences that could not be decoded into UTF-8 like "Ä,Ö.." (which is not encoded in UTF-8) we noticed that the router have excepted the character without issuing any error.

The PE responded with "malformed-message" error only when received special characters from the following list (!,@,#,$,%,^,&,*,(,),_,{,}, [,],:,;,",',<, >,`,/,?,."",',&lt;,&gt;,`,/,?,~,.).

The following systems successfully participated in the test: Cisco NSO, Nokia NSP acted as controller and Ciena 8112 as PE router.

YANG Model Retrieval

This test considered as simple, nevertheless from the minor reported issues of the NETCONF/YANG protocol, the errors during the fetching of the YANG models from the elements.

Controllers performed this step ahead of the testing, and a full list of the router YANG modules was fetched.

The following systems successfully participated in the test: Cisco NSO, Huawei NCE and Nokia NSP acted as controller, Ciena 8112, Huawei NetEngine 8000 X8 and Juniper MX204 as PE router.

Conclusion

EANTC Interoperability Testing Event 2022 is the most recent episode of the long multi-vendors environment testing series, yet to be one of the most highlighted testing events we've ever organized and hosted.

The 2022 testing event came when the industry was in a whole new era. 5G, IoT, AI, and more new life-changing technologies are starting to occur in each aspect of our lives. Vendors, Service Providers, and Technology development bodies work together to adapt to the new world and new capabilities and possibilities these new technologies could achieve.

During the event this year, we achieved outstanding results in all the testing areas which were touched, which led to having some important headlines in each area.

Segment Routing and Flex Algo: We were able to test The interworking between different domains (SR-MPLS, SRv6, and VXLAN) between different vendors, in addition, to testing Prefix summarization in SRv6 and tests of resiliency that included TILFA (topology-independent loop-free alternate) and S-BFD (Seamless BFD).

EVPN: The concentration this year was on the large Data centers technologies like MAC Mobility and the interworking between EVPN VXLAN and VXLAN, which showed how we could extend the bridge domain via EVPN VXLAN environment between data centers with Seamless Stitching of VXLAN tunnels coming from LAN level to the VXLAN tunnels which is dedicated to data center interconnect while avoiding EVPN VXLAN tunnels to grow exponentially in real-world scenarios.

SDN: Once again, we were able to feel the evolution and increasing robustness of the SDN technologies across the vendors. Utilizing the capabilities of BGP-TE and PCEP for managing the routing policies and testing in an SRv6 network was the highlight of this area.

Clock Synchronization: It is always delightful to present new aspects for the clock synchronization and distribution through the networks. Different approaches of PTP Topologies and Profiles were tested, PTP over FlexE, beside the Boundary Clocks Class C/D conformance test which was introduced in our events previously last year. Having said that, we are now used to introduce first time industry tests in our events, which leads to the next topic.

Open Radio Access Network Fronthaul (O-RAN FH): EANTC is an active member of the O-RAN Alliance together with multiple vendors who participated in our event. This special relationship led to first industry tests O-RAN Fronthaul reference topology in our lab. EANTC' unique expertise in testing and organizing blitz, efficient, and valid testing events combined with dedication and the enthusiasm for new technologies from Juniper, Microchip and Calnex led to achieve PTP time synchronization using ITU-T G.8275.1 profile in LLS-C3 topology. These tests and we are planning to expand such kind of tests for the next years.

NETCONF/YANG: Through the years of EANTC testing series, we are noticing the wider support for NETCONF/YANG from both vendor side and service providers. Although we face always some interoperability issues—which is the goal of this event—we observed how the interworking between the different vendors is becoming easier and more efficient. Integrating the different YANG models which are supported from the networks elements, and the controllers are getting better for OpenConfig YANG models, IETF models, or even the proprietary models. The easy integration of the proprietary YANG models has a downside, which discourages the networks elements vendors to support the standardized models (vendor agnostic models). We did not observe a noticeable increase in supporting the vendors' agnostic models from the vendors which is a major requirement and would be a huge benefit for the service providers.

In the end, we thank our participants enough for the hard work and support, from creating and developing the tests until finishing all the test execution and delivering the brilliant work.

EANTC AG
Salzufer 14, 10587 Berlin, Germany
info@eantc.de, https://www.eantc.de/
[v1.2 20220414]