EANTC

Multi-Vendor Interoperability Test

White Paper 2020

Table of Contents

Editor's Note

Among the longstanding series of annual EANTC multi-vendor interoperability events at Upperside's MPLS + SDN + NFV World Congress, this year's event has turned out to be special in multiple ways. Most importantly, the novel Coronavirus (COVID-19) has not managed to stop the hot staging. For two weeks in early March, more than 100 attendees from 16 vendors teamed up to validate the interoperability and integration of state-of-art network technologies. More than two-thirds of the participants joined remotely this time. At EANTC, we have conducted distributed interop and performance tests for virtualization technologies (NFV) since 2017, which typically took much longer due to the availability constraints of individual teams. The coordinated testing of physical equipment with such a large group of remote participants in real-time is unprecedented in the industry. Within two weeks, we successfully completed more than 500 test combinations thanks to the enthusiasm and dedication of local and remote vendor teams, some of them in distant time zones with inconvenient working hours.

For the first time, we are publishing the white paper in advance of the actual showcase. The results of our testing are ready for deployment now, and we would like to share the details of the multi-vendor technology solutions already now. We will publish a series of eight short videos on the event page covering the full range of test areas, starting on March 31st, co-presented by the participating vendors, many with live multi-vendor demos.

Technically, our tests evolved from last year's coverage, confirming that many vendor solutions are further solidified and extended. We have focused Segment Routing, EVPNs, and SDN orchestration, covering more advanced scenarios with more participating vendors than before. From our point of view, it is the best time right now for service providers to migrate to Segment Routing if you haven't done so yet. More vendors joined the SRv6 tests as well, growing the ecosystem. Given the sometimes mutually exclusive support, it is likely that SR over MPLS, SR over VXLAN, and SRv6 will coexist in the future. Vendors who support all three modes might play a key role in future networks.

Open source and white box solutions take an increasingly important role in our tests. One vendor provided an open-source SDN controller based on OpenDaylight. Two other vendors jointly provided a white-box edge router solution. In previous years, we had already seen some open source and/or whitebox-related solutions (not all of them returned for participation this time). Clearly, the functionality and maturity of the participating solutions are maturing.

There are many more results and successes to be shared, for example in the clocking area and with regards to flexible Ethernet (FlexE). We hope that the white paper will provide all the details you are looking for!

One final word: The MPLS + SDN + NFV World Congress has been rescheduled to June 30 to July 3, 2020. (Check for updates at https://www.upperside-conferences.com/). I really hope that the conference will take place in beautiful Paris, and looking forward to celebrating the technology innovations, our community and life in general with all of you. Until then, stay safe and healthy!

Introduction

This year, we designed the interoperability test cases to probe the maturity of transport network solutions to support 5G networks, data center networking evolution, and multi-vendor domain orchestration. The event covered more than 60 test cases evaluated in many multi-vendor combinations. We classified the test cases in four main categories; Ethernet VPN (EVPN), Segment Routing, Software-Defined Networking (SDN), and packet networks clock synchronization.

In the EVPN section, the participating vendors broadly supported the basic test scenarios for symmetric and asymmetric integrated routing/bridging (IRB). New test cases were added focusing on EVPN redundancy and

Interworking aspects (i.e., Active-Active Proxy MAC-IP Advertisement). The Segment Routing (SR) section covered various deployment flavors of SR. For the first time, we successfully tested the interoperability of the SR-MPLS data plane using OSPFv2 with equipment from nine vendors. SR resilience solutions took the lion's share of attention this year: We tested numerous TI-LFA scenarios including SR-MPLS and SRv6 data plane. Additionally, we verified Seamless BFD implementations for SR Traffic Engineering tunnels. Our tests of the SRv6 data plane progressed well, with one new vendor joining.

As FlexE support is growing rapidly across the industry, we included it for the first time. The technology is specifically suited for 5G transport network slicing. We verified the basic interoperability and bandwidth adjustment of multi-vendor FlexE pipes over a 100GbE link.

Our SDN orchestration tests included two main areas: Controller-based traffic engineering and network automation using a combination of NETCONF and RESTCONF protocols with a range of YANG models. We examined services creation including L3VPN and L2VPN over MPLS data plane. Also, we tested EVPN service provisioning over MPLS data plane for the first time.

The multi-vendor interoperability tests of packet network clock synchronization showed an excellent level of maturity in terms of high-precision clock source failover, GNSS security, and boundary clock performance. All these aspects directly relate to the requirements of 5G networks and Cloud-RAN deployments.

Participants and Devices

Participants	Devices
ADVA Optical Networking	Activator
Arista Networks	7050SX2 7050X3 7280R2 7280R3
Arrcus	ArcRR QuantaMesh T4048-IX8A QuantaMesh T7080-IXAE
Calnex Solutions	Paragon-T Paragon-X
Ciena Communications	5171 6500 T12
Cisco Systems	3100-V 3600-R Nexus 9300-FX Nexus 9300-FX2 Nexus 9300-GX NSO Server PC
Delta Electronics	AGC7648SV1 AGCV208S
ECI Telecom	NPT-1800
Huawei Technologies	ATN980C NetEngine 8000 F1A NetEngine 8000 M8 NetEngine 8000 X4 NCE Controller
Juniper Networks	ACX5448-D cRPD HealthBot MX204 NorthStar Controller QFX10002-72Q QFX5110-48S QFX5120-32C QFX5120-48Y
Keysight Technologies	Ixia Ixia IxNetwork Ixia XGS2/Novus
Lumina Networks	SDN Controller Suite

Participants	Devices
Metaswitch Networks	NOS Toolkit
Microchip	BlueSky TimeProvider 4100 TimeProvider 5000
Nokia	7750 SR-1 NSP Server
Spirent Communications	TestCenter N4U

Table 1: Participants and Devices

Interoperability Test Results

This white paper documents only positive results (passed test combinations) individually with 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 proceed to solve interoperability issues after our test so there is no point in punishing them for their willingness to learn by testing. Confidentiality is vital to encourage manufacturers to participate with their latest - beta - solutions and enables a safe environment in which to test and to learn.

Terminology

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

Test Equipment

With the help of participating test equipment vendors, we generated and measured traffic, emulated and analyzed control and management protocols and performed clock synchronization analysis. We thank Calnex, Ixia and Spirent for their test equipment and support throughout the hot staging.

EVPN

EVPN and MAC Learning Intelligence

There was no dispute about the techniques of learning MAC addresses how significant these are. The IETF protocols were all within the scope of EVPN testing. Just like MAC mobility, integrated routing and switching (IRB) and loop prevention, effective solutions to the headaches of MAC addresses over the control plane have gradually become test standards.

As all-active multi-homing moved into the fore of testing, various EVPN-related benefits received the test features, such as EVPN with heterogeneous integration like with a single-homed remote peer or self-multi-homed in a multi-vendor environment, enabled with load balancing, and protection.

Routing with EVPN: Symmetric Model

Initially, the data center overlay was viewed as a layer 2 domain. For example, between hosts separated by an EVPN, the only communication, if any exists, remains in layer 2. However, if these hosts are on different subnets, such inter-communication required a Layer 3 data center gateway. Since 2013, the IETF released the Integrated Routing and Switching draft to replace the centralized gateway approach with an integrated mechanism - IRB. The IRB functionality is needed on the PE nodes to avoid inefficient forwarding of overlay traffic. We tested the IRB for the sixth consecutive year, based on the IETF specification which was latest updated in 2019. We wanted to witness that the rapid growth of the data center architecture remains interoperable given by the different options. On the base of the single-homed PE solutions with IRB, the all-active multi-homed PE moved into sight. The mode of how ingress PE and egress PE interacted with IRB included the symmetric as well as asymmetric semantic. The EVPN encapsulation allowed both the EVPN-MPLS as well as EVPN-VXLAN. The layer 2 service types included the VLAN-based and the VLAN-bundle-aware EVPN. Each of the options received in return to the test, a versatile of rich combinations.

In the symmetric IRB test, the PEs established the EVPN with each other running symmetric IRB. In this mode, both IP and MAC lookup were required at both ingress and egress NVEs. We checked at the PEs and expected via BGP to exchange RT 2 (host IP address routing) and RT 5 (IP prefix) between the PEs.

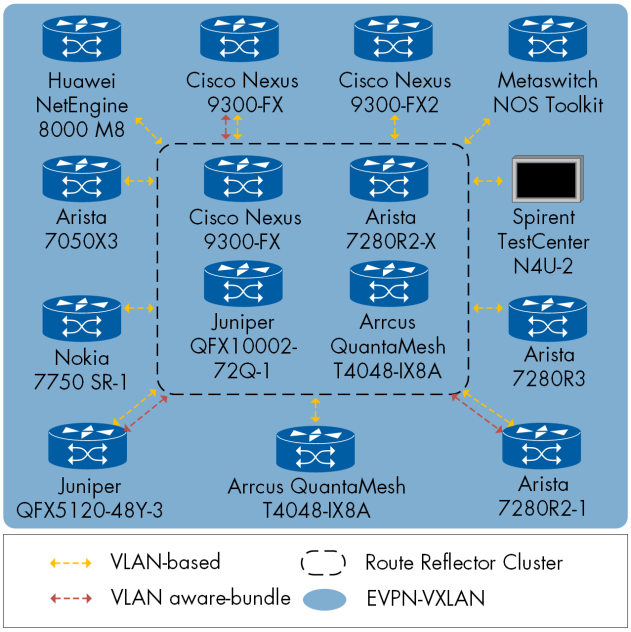

Figure 1: Symmetric IRB for VLAN-based EVPN with VXLAN Encapsulation

We observed that different subnet prefixes of emulated hosts were learned via RT5 from each other PEs. In all tests, we sent IPv4 test traffic between hosts established. We observed in the test there was no packet loss.

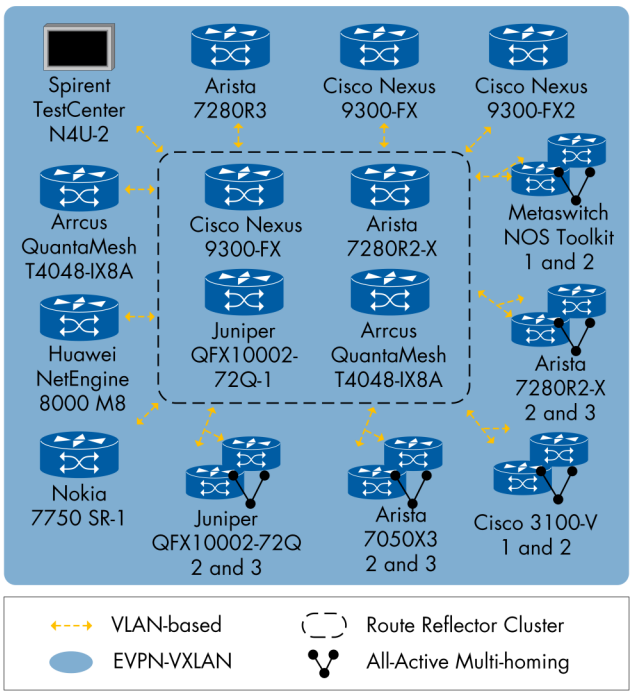

Figure 2: Symmetric IRB for Multi-Homed EVPN with VXLAN Encapsulation

In an all-active multi-homing scenario, the multi-homing pair advertise A-D per ES and A-D per EVI routes with the same ESI value. Also any MAC address learned by the pair is advertised in a MAC/IP route with the same ESI advertised in the A-D route.

This allowed the remote PEs to install the MACs that came from the multi-homing pair so that the next-hop contained both all-active PEs. The single-home PEs do not advertise A-D routes and their MAC/IP routes contain a zero ESI. We observed there was no packet loss.

The following devices successfully participated in the test:

- All-active Multi-Homed PE: Arista 7050X3, Arista 7280R2, Cisco 3100-V, Juniper QFX10002-72Q, Metaswitch NOS Toolkit
- Single-Homed PE: Arista 7050X3, Arista 7280R2, Arista 7280R3, Arrcus QuantaMesh T4048-IX8A, Cisco Nexus 9300-FX, Cisco Nexus 9300-FX2, Huawei NetEngine 8000 M8, Juniper QFX5120-48Y, Nokia 7750 SR-1, Spirent TestCenter N4U
- Spine: Arista 7280R2, Arrcus QuantaMesh T4048-IX8A, Cisco Nexus 9300-FX, Juniper QFX10002-72Q

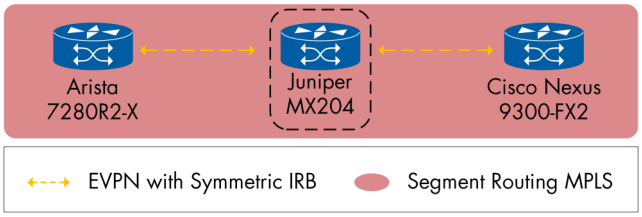

Figure 3: Symmetric IRB for EVPN-MPLS

The following devices successfully participated in the test:

- PE: Arista 7280R2, Cisco Nexus 9300-FX2
- Route Reflector: Juniper MX204

Routing and Switching with EVPN: Asymmetric Model

The asymmetric IRB semantic requires both IP and MAC lookups at the ingress NVE with only MAC lookup at the egress NVE. In the asymmetric test, we checked at the PE that remote MAC addresses that came via RT 2 (MAC/IP routes) were installed in the local ARP table.

Before starting with the test, we verified that there was no flooding in the network with test traffic. The 3 lookups from ingress require the encapsulation with the destination MAC address, and flooding from ingress means an unknown MAC address, which was only expected at the MAC learning phase. This step was not needed in the symmetric test, because lookup ends at layer 3 both sites.

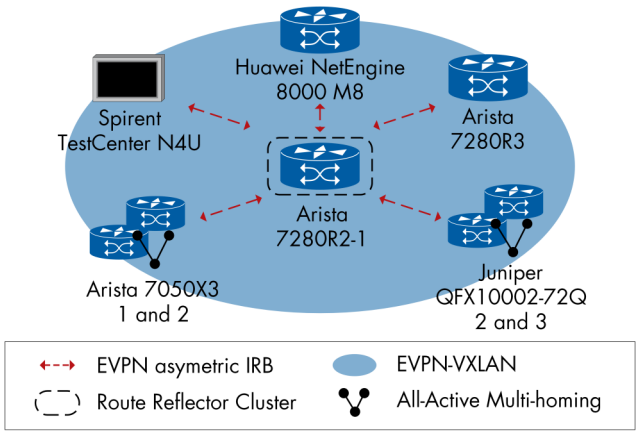

Figure 4: Asymmetric IRB for VLAN-based EVPN with VXLAN Encapsulation

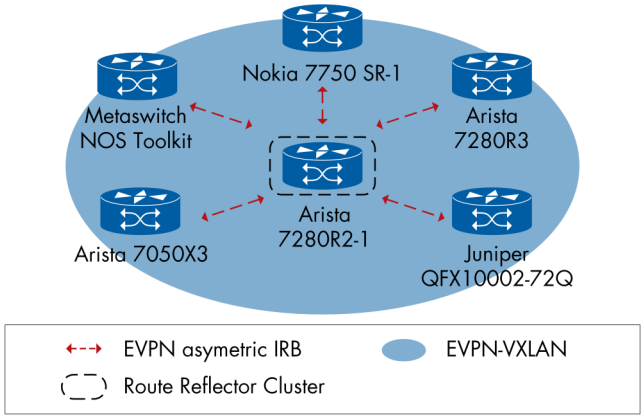

Figure 5: Asymmetric IRB for VLAN-based EVPN with VXLAN Encapsulation

The following devices successfully participated in the test:

- All-Active Multi-Homed PE: Arista 7050X3, Juniper QFX10002-72Q
- Single-Homed PE: Arista 7280R3, Huawei NetEngine 8000 M8, Juniper QFX10002-72Q, Metaswitch NOS Toolkit, Nokia 7750 SR-1
- Spine: Arista 7280R2

MAC Mobility

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 RFC7432. 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.

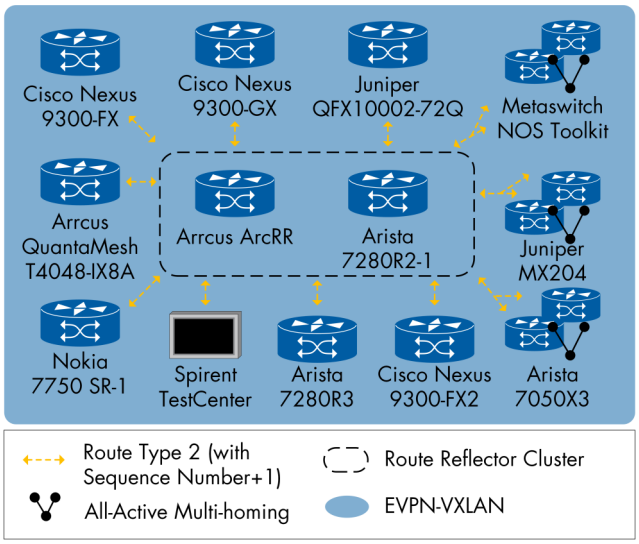

Figure 6: MAC Mobility at EVPN-VXLAN

The following devices successfully participated in the EVPN-VXLAN test:

- All-active Multi-Homed PE: Arista 7050X3, Juniper MX204, Metaswitch NOS Toolkit
- Single-Homed PE: Arista 7280R3, Arrcus QuantaMesh T4048-IX8A, Cisco Nexus 9300-FX, Cisco Nexus 9300-FX2, Cisco Nexus 9300-GX, Juniper QFX10002-72Q, Nokia 7750 SR-1, Spirent TestCenter N4U
- Spine: Arista 7280R2, Arrcus ArcRR

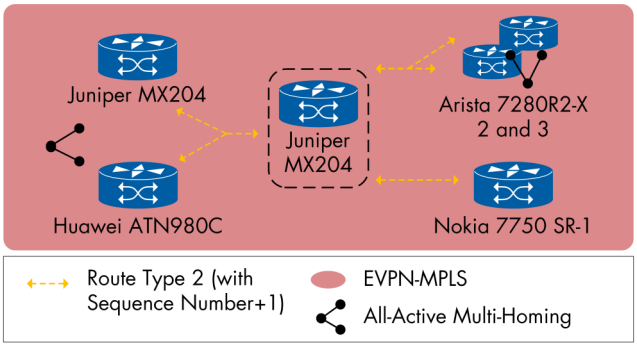

Figure 7: MAC Mobility at EVPN-MPLS

In one case, the sequence value sent by a PE was not interpreted by all PEs in the same way. As explained by the vendor, there are two types of routes RT2 and RT5 (the draft IETF Extended Mobility Procedures for EVPN-IRB) that support MAC mobility, and the DUT responded to both types, however, the added value was not equal, which caused the same MAC with different sequence value appeared in the network, which bounced a false host-move back to other PEs causing a loop in EVPN. We removed the DUT from the EVPN and repeated the test.

The following devices successfully participated in the EVPN-MPLS test:

- All-Active Multi-Homed PE: Arista 7280R2, Huawei ATN980C, Juniper MX204
- CE: Huawei NetEngine 8000 F1A
- Route Reflector: Juniper MX204
- Single-Active Multi-Homed PE: Huawei ATN980C, Juniper MX204
- Single-Homed PE: Nokia 7750 SR-1

EVPN Loop Detect

IETF specified the draft Loop Protection in EVPN networks. It is suitable for layer 2 as an optional loop protection operation, which is applied on duplicate MAC addresses. This additional mechanism resolves the loop caused by accidental backdoor links. Recalling the scenario from the MAC mobility test, where a rapid host-moves presented a duplicated MAC address among PEs, due to the continuous MAC movement. The MAC damping functionality within the EVPN standard would result in the MAC addresses being black-listed from the EVPN control-plane, but any loop in the forwarding plane would remain. We connected a cable between the ACs on both sides of the EVPN to make the EVPN to a loop.

We verified at the PE that the repetition of the MAC moves is set to 3 times. Then we generated Ethernet frames from an AC attached to the EVPN. We observed on the PE that the MAC black hole has been detected and the loop is broken. There were two solutions to break the loop in this test. In one solution there was a scheme to block the AC port which was on a lower priority than the PE side.

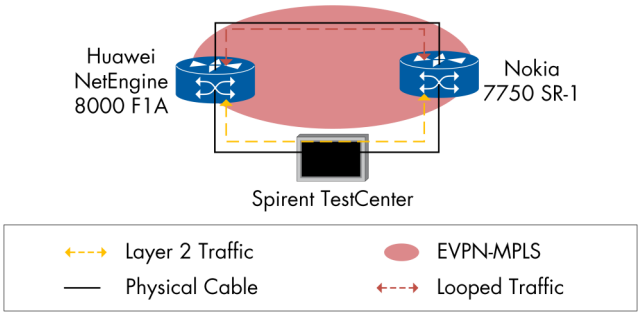

Figure 8: EVPN Loop Detection

On the other solution, the duplicate MAC addresses were installed as blackhole MACs and frames with source MAC matching the blackhole MACs were discarded. No traffic drop was observed after that the loop has been blocked.

The following devices successfully detected and blocked the loop:

- Huawei NetEngine 8000 F1A, Nokia 7750 SR-1

Proxy MAC-IP Advertisement

All-Active multi-homed EVPN allows incoming traffic shared among the active links based on the hashing algorithm. This results in not all source MAC addresses will be learned through the same PE. The PE's which belong to the same ESI will learn a different set of MAC addresses and advertise different sets of EVPN route-type 2. In case of link or node failure, EVPN type 2 routes will be withdrawn and these addresses will not be reachable, till the traffic is sent out through a different active link from the host.

Proxy (IP-MAC) allows all the PE's in the same ESI to re-advertise the same EVPN route-type 2 even it's not learned locally, provided that the proxy bit is set. This enables the traffic to flow in and out of the multihomed CE during the transient time of link or node failure.

We sent test traffic and verified it was load-balanced between all active links of the all-active multi-homed PE. We performed the link failure test at the PE which was ready for the failover setup. The out of service time was 300 ms.

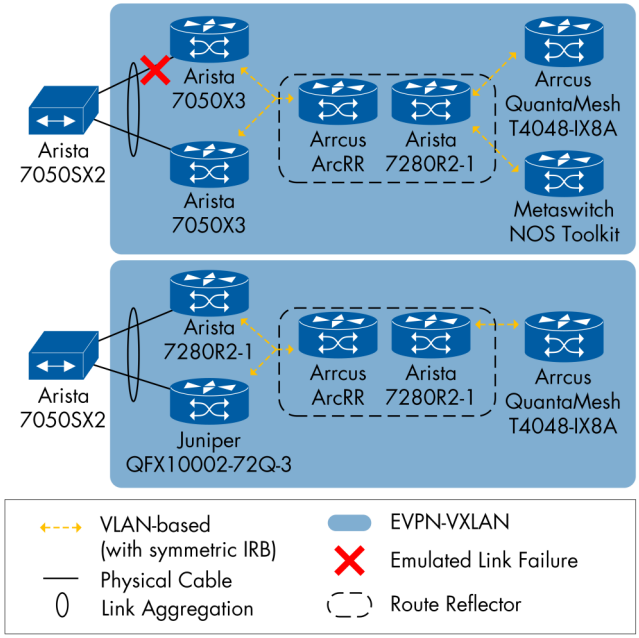

Figure 9: Proxy MAC-IP Advertisement

The following devices successfully participated in the test:

- All-Active Multi-Homed PE: Arista 7050X3, Arista 7280R2, Juniper QFX10002-72Q
- Single-Homed PE: Arrcus QuantaMesh T4048-IX8A, Metaswitch NOS Toolkit
- Spine: Arista 7280R2, Arrcus ArcRR

Carrier Ethernet

EVPN lighted up into Carrier Ethernet once again at the hot staging. Each service type found its named place on the BGP control plane inherited from EVPN and enabled unicast, multicast and broadcast to the MPLS-based data plane. The IETF L2VPN working group, as one of the EVPN's leading standard bodies in the Carrier Ethernet, is currently standardizing the framework. We looked forward to a big surprise, neither in EVPN nor in Carrier Ethernet, as both are well-understood protocols.

E-Line Service

EVPN inherits VPWS technology which is a natural choice for E-Line. The framework is currently under standardization by the IETF L2VPN Working Group. The control plane requires per EVI Ethernet Auto-Discovery route as per RFC 4760. EVPN defines a new BGP Network Layer Reachability Information (NLRI) used to carry all EVPN routes. BGP Capabilities Advertisement used to ensure that two speakers support EVPN NLRI (AFI25, SAFI70).

Among all EVPNs appeared on the test stage, the multi-homed setup enabled interaction through AD route in a multi-vendor scenario, and performed redundancy. We verified the all-active multi-homed EVPN with ESI tag 0 on PE, with ESI tag 1 for single-active multi-homed EVPN, and other unique values for the single-homed EVPN. We sent IPv4 traffic to the EVPN and did not observe any packet loss. In the multi-homed scenario, we introduced a link failure between CE and PE while traffic was running, then measured the convergence time of the EVPN service. All-active multi-homing scenario:

- Out of service time during failover (max): 226 ms
- Out of service time during recovery (max): 0.1 ms

Single-active multi-homing scenario:

- Out of service time during failover (max): 474 ms
- Out of service time during recovery (max): 0.1 ms

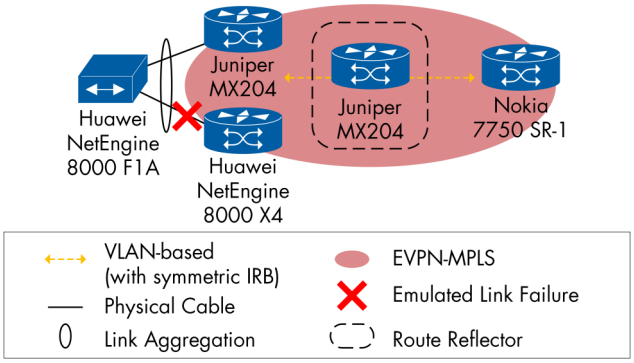

Figure 10: E-Line Service

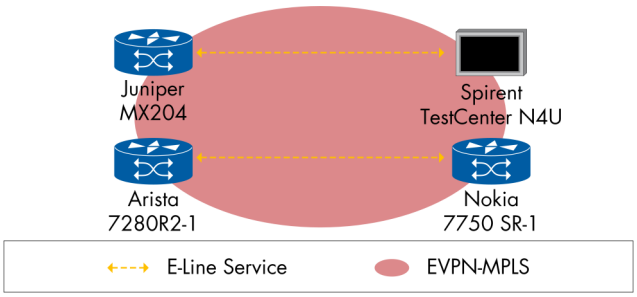

Figure 11: E-Line Service

The following devices successfully participated in the test:

- All-Active Multi-Homed PE: Huawei NetEngine 8000 X4, Juniper MX204
- Route Reflector: Juniper MX204
- Single-Active Multi-Homed PE: Huawei NetEngine 8000 X4, Juniper MX204
- Single-Homed PE: Arista 7280R2, Juniper MX204, Nokia 7750 SR-1, Spirent TestCenter N4U

Flexible-Cross Connect

EVPN introduces the EVPL (Ethernet Virtual Private Line) service type, and IETF drafted the VPWS flexible cross-connect (FCX) since 2018. By setting the Announce bit in AD route to 1, the peer binds ACs into a single EVPN to reduce the number of routes and resources bound to it.

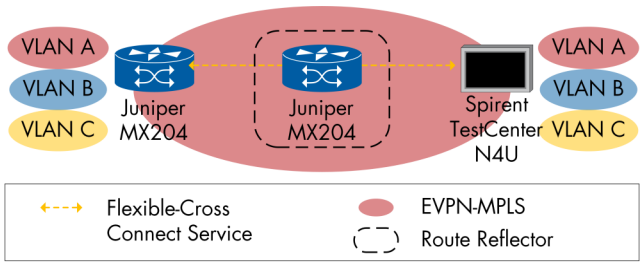

Figure 12: Flexible-Cross Connect

We verified at the PE that AD with Announce bit 0 was learned. We also checked in the routing table that the ACs represented by multiple VLANs were only propagated with a single EVPN route. We sent IPv4 traffic and did not observe any packet loss.

The following devices successfully participated in the test:

- PE: Juniper MX204, Spirent TestCenter N4U

E-LAN

We also verified multi-point to multi-point EVPN service in terms of E-LAN service. The setup was taken from the MAC mobility test.

- Out of service time during failover (max.): 74ms
- Out of service time during recovery (max.): 0 ms

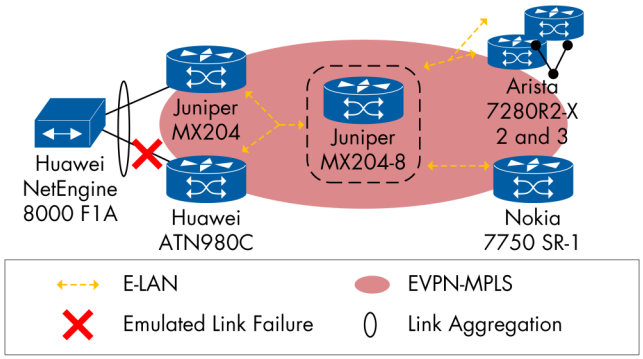

Figure 13: E-LAN

The following devices successfully participated in the test:

- All-Active Multi-Homed PE: Huawei ATN980C, Juniper MX204
- Single-Homed PE: Arista 7280R3, Nokia 7750 SR-1

E-Tree Service

In an E-Tree service E-Tree (rooted-multipoint), endpoints are labeled as either Root or Leaf sites. Root sites can communicate with all other sites. Leaf sites can communicate with Root sites but not with other Leaf sites. The RFC 8317 specified three scenarios to distinguish the mesh relationship between the roles. Given the binding Leaf or Root Site(s) per AC, PEs can receive traffic from the root and leaf ACs of a given EVI. The EVPN introduces multicast capability with Inclusive Multicast Ethernet Tag route (RT3).

We sent IPv4 multicast traffic to the established E-Tree service first from the root to any other endpoints. All leafs and roots from other ACs received packets without any packet loss. Then we sent unicast IPv4 packets from each leaf to all other endpoints. We expected at all roots to receive the test traffic without any packet loss, as well as 100% packet drop between the leafs.

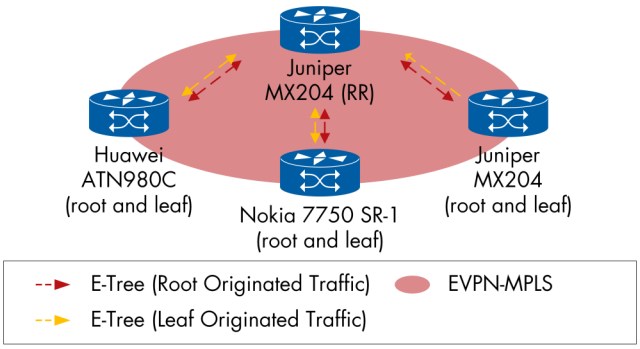

Figure 14: E-Tree

The following devices successfully participated in the test with root and leaf roles:

- Huawei ATN980C, Juniper MX204, Nokia 7750 SR-1
- Route Reflector: Juniper MX204

In one case there was at one root unexpected packet loss at where all incoming packets were discarded. The reason lied in VLAN strip. The sender removed the VLAN tag from the EVPN packet and used labels, but the receiver could not handle the traffic.

EVPN with Multicast

EVPN has always been committed to resource-saving since the layer 2 broadcast domain is based on the control plane. Two significant multicast protocols came into test focus.

IGMP Proxy

The IGMP proxy function optimizes network usage through a selective multicast mechanism. The host expresses their interests in multicast groups on a given subnet/VLAN by sending IGMP membership reports (Joins) for their interested multicast group(s). An IGMP router periodically sends membership queries to find out if there are hosts on that subnet still interested in receiving multicast traffic for that group. The goal of the IGMP proxy mechanism is to r educe the flood of IGMP messages (query and report) in EVPN instances between PE routers. The IETF draft (evpn-igmp-mld-proxy) defines a set of multicast routes, among which Selective Multicast Ethernet Tag (Type 6: SMET) Route enabled the basic setup of an EVPN IGMP proxy in this test.

To verify that proxy function has been enabled we first sent multicast traffic without any Join messages generated from an emulated host. As expected, all traffic was dropped. Then we generated the Join messages to the PE from the emulated host, and checked at each PE that the multicast group has been learned and the RT 6 route with the "IGMP Proxy Support" flag enabled was installed in the routing table. We sent multicast traffic and did not observe any packet loss.

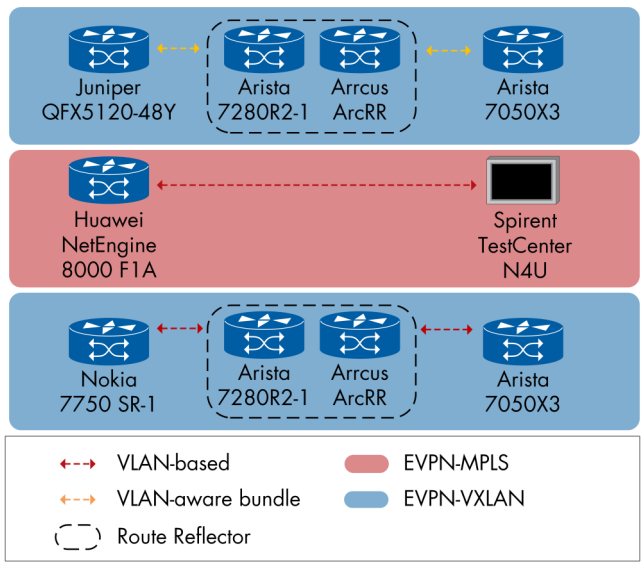

Figure 15: IGMP Proxy

While traffic was running, we generated leave messages from the emulated host. As expected, then all traffic was dropped.

The following devices successfully participated in the test as IGMP proxy:

- Emulated host with IGMP joins and leaves: Spirent TestCenter N4U
- PE: Arista 7050X3, Huawei NetEngine 8000 F1A, Juniper QFX5120-48Y, Nokia 7750 SR-1, Spirent TestCenter N4U
- Route Reflector: Arista 7280R2, Arrcus ArcRR

One vendor separated itself from the test. The PE in this test did not install any IGMP flag into the RT6 route, and the routes were discarded from other PEs. Therefore, this solution was not compliant to the standard.

Assisted Replication

Assisted Replication (AR) is under standardization in the IETF for multicast distribution, it updates the ingress replication mechanism and introduces the remote node replication mechanism. The driven motivation behind is that this helps software-based EVPN nodes to efficiently replicate large amounts of broadcast and multicast traffic to remote nodes without affecting their replication performance. With Assisted Replication feature, ingress passes all broadcast and multicast traffic to an AR-replicator that replicates the traffic further in the EVPN service on its behalf. In other words, the replicator represents powerful resources, which will greatly simplify the layout requirements of ingress nodes which need to be enabled with assisted replication.

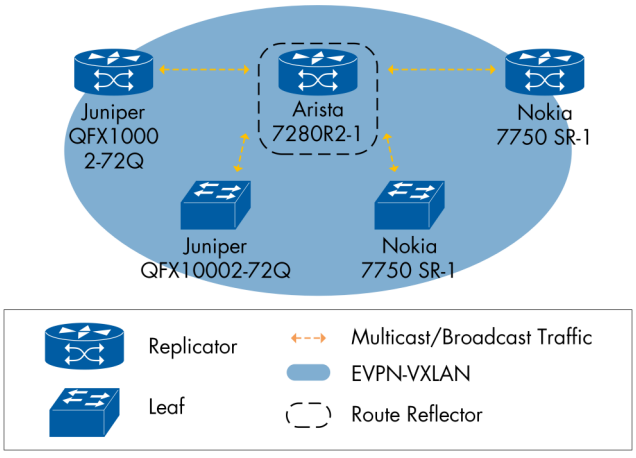

Figure 16: Assisted Replication

We verified at the ingress node with AR enabled that a route of the replicator is shown. Through port statistic, we observed that both multicast and broadcast were traffic directed to the AR-replicator from the ingress, and over there traffic was replicated on the AR-replicator to the other leafs as expected.

The following devices successfully participated in the test:

- Leaf: Juniper QFX10002-72Q, Nokia 7750 SR-1
- Replicator: Juniper QFX10002-72Q, Nokia 7750 SR-1
- Route Reflector: Arista 7280R2

Multi-Domain Interconnection

Just like the test richness of the EVPN inside the data center, how to show it from the outside also brought us suspense expectations, and further prospects became the vision during the hot staging to bring further testing ideas. See the results of verified EVPNs stationed at the border that provided interactions with various VPN third parties.

EVPN-VXLAN and EVPN-MPLS Interworking

Traversing EVPN data planes filled the first gap in the puzzle. Anyone accessing a VXLAN data center can go through a EVPN-MPLS network. When the gateway provides dual VXLAN and MPLS data planes, it also inherits the advantages of the EVPN control plane. No new signaling is required, the MAC learning is given in the control plane as shown in many other tests during the hot staging. Therefore, the gateway carries out routes from one domain to the other. We did not observe any packet loss in the end-to-end traffic.

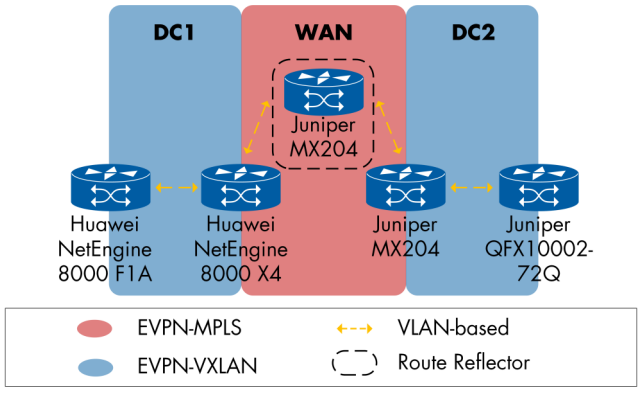

Figure 17: EVPN-VXLAN and EVPN-MPLS Interworking

The following devices successfully participated in the test:

- Data Center Edge: Huawei NetEngine 8000 F1A, Juniper QFX10002-72Q
- Data Center Gateway: Huawei NetEngine 8000 X4, Juniper MX204
- Route Reflector: Juniper MX204

EVPN and IP-VPN Interworking

Expanding the scope of VPNs, an MPLS core network will emerge between data centers. The gateway requires dual data planes and shall provide translations between EVPN routing and MPLS labels.

In view of the multiple successful results of this test, EVPN inspired our imagination and proposed a further testing idea for dual-homed gateways 2021. The IETF discussed the options for providing dual-homed gateways for EVPN and IPVPN domains in the "EVPN and IPVPN Interworking" draft updated in 2019. As a prospect for multi-homing solutions across data centers, we are looking for tests to prevent loop such as D-PATH capabilities. Although the current test scope was apparently a single-homed gateway, and the idea has not yet been added, but Arista demonstrated the D-PATH from a single-homed gateway via CLI before the test was completed, let us have an outlook for the next event.

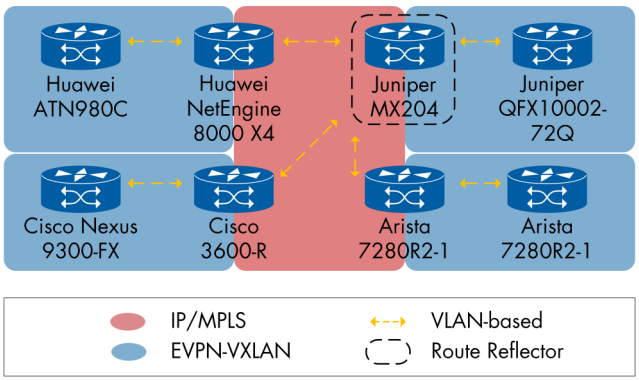

Figure 18: EVPN and IP-VPN Interworking

The following devices successfully participated in the test:

- Data Center Edge: Arista 7280R2, Cisco Nexus 9300-FX, Huawei ATN980C, Juniper QFX10002-72Q
- Data Center Gateway: Arista 7280R2, Cisco 3600-R, Huawei NetEngine 8000 X4, Juniper MX204
- Route Reflector: Juniper MX204

EVPN-VXLAN and EVPN-VXLAN Interworking

For EVPNs that traverse multiple data centers, each data center consists of a different AS, and the gateway will be responsible for exporting routes from one network segment to another.

We did not observe any packet loss.

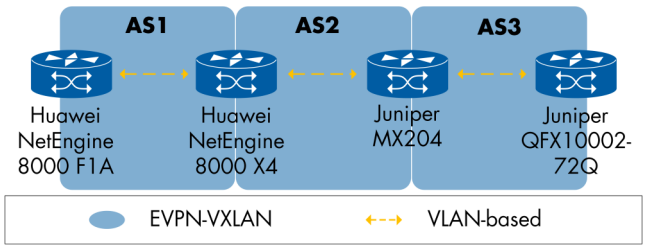

Figure 19: EVPN-VXLAN and EVPN-VXLAN Interworking

The following devices successfully participated in the test:

- Data Center Edge: Huawei NetEngine 8000 F1A, Juniper QFX10002-72Q
- Data Center Gateway: Huawei NetEngine 8000 X4, Juniper MX204

Segment Routing

The momentum of Segment Routing is magnifying nowadays. This is happening because of technology maturity which is provided in the market by the network vendors supported by the deep belief of the customers in the capabilities of the segment routing paradigm.

Standing on the source routing paradigm, segment routing with its versions (SR-MPLS and SRv6) natively enriches the current transport networks with more features spanning the network resiliency and network automation.

This year, the SR-MPLS interoperability testing includes various topological combinations, by involving different interior gateway protocols (IGP); ISIS and OSPF. The SRv6 data plane has a good portion of test cases that are related to L3VPN, E-Line, TI-LFA and Egress Node Protection.

Business VPNs with Segment Routing

Both layer 2 and layer 3 VPNs determined their capabilities in segmented routing which in return provided a rich set of protocol features for testing.

The traditional IPv4/IPv6 BGP-based L3VPN certainly found its place in the test over the SR-MPLS data plane as well as the SRv6 data plane. EVPN services are located in the SRv6 data plane.

IPv4/IPv6 BGP-based L3VPN over Segment Routing MPLS Data Plane

For many years, SR-MPLS was frequently listed at the top of our interoperability testing areas. Previously, the participated vendors selected IS-IS protocol as the preferred option to build the SR-MPLS topology. This year we expanded the available link-state protocols that can be used to build the SR-MPLS networks using OSPFv2. The IETF draft (ospf-segment-routing) describes the required extensions in OSPFv2 to support segment routing.

The physical topology for this test case was simplified as a leaf-spine architecture. The provider (P) router represented the spine node, while all the other provider edge (PE) routers represented the leaf nodes. We physically connected each PE with one link to the traffic generator which emulated the customer edge (CE) router.

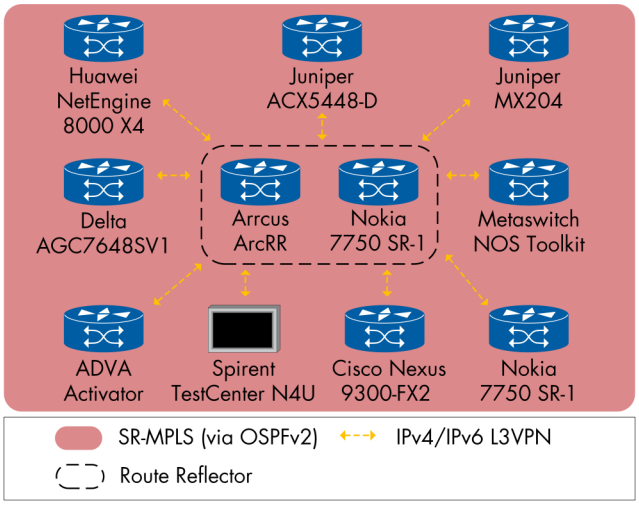

Figure 20: L3VPN over Segment Routing MPLS (Underlay IGP with OSPFv2)

Across the different IGP (IS-IS, OSPFv2) scenarios, all participated vendors configured IPv4/IPv6 BGP-VPN L3VPN services. We started the test by sending bidirectional IPv4 and IPv6 traffic between each pair of PE routers. The generated traffic (emulated customer traffic) was encapsulated in two SR-MPLS labels; VPN & transport segments. We verified the traffic flow between all the PE pairs without any packet loss.

The following devices successfully participated in the L3VPN for both IPv4 and IPv6 services over MPLS-based Segment Routing. The underlay IGP was OSPFv2.

- P: Arrcus ArcRR, Nokia 7750 SR-1
- PE: ADVA Activator, Cisco Nexus 9300-FX2, Delta AGC7648SV1, Huawei NetEngine 8000 X4, Juniper ACX5448-D, Juniper MX204, Spirent TestCenter N4U, Metaswitch NOS Toolkit, Nokia 7750 SR-1

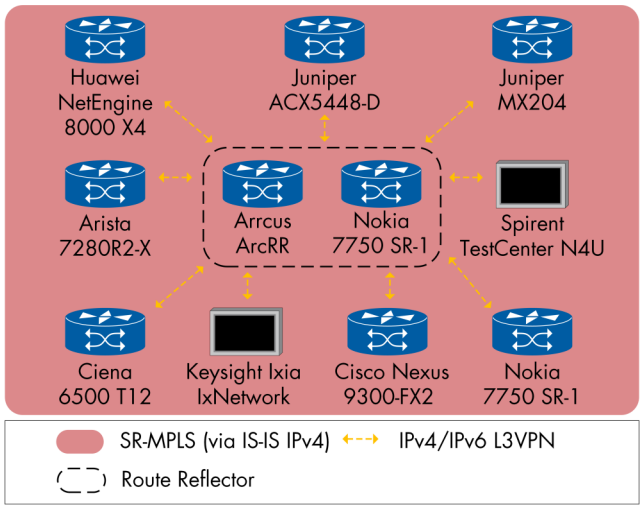

Figure 21: L3VPN over Segment Routing MPLS (Underlay IGP with IS-IS IPv4)

The following devices successfully participated in the L3VPN for both IPv4 and IPv6 services over MPLS-based segment routing. The underlay IGP was IS-IS IP.

- P: Arrcus ArcRR, Nokia 7750 SR-1
- PE: Arista 7280R2, Ciena 6500 T12 , Cisco Nexus 9300-FX2, Huawei NetEngine 8000 X4, Juniper ACX5448-D, Juniper MX204, Keysight Ixia IxNetwork, Nokia 7750 SR-1, Spirent TestCenter N4U

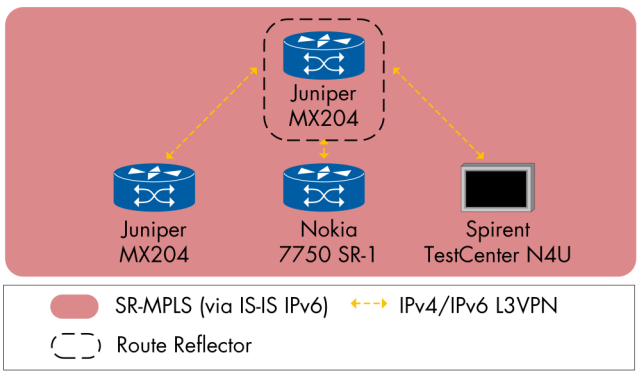

Figure 22: L3VPN over Segment Routing MPLS (Underlay IGP with IS-IS IPv6)

The following devices successfully participated in the L3VPN for both IPv4 and IPv6 services over MPLS-based Segment Routing. The underlay IGP was IS-IS IPv6.

- P: Juniper MX204
- PE: Juniper MX204, Nokia 7750 SR-1, Spirent TestCenter N4U

EVPN and L3VPN over SRv6 Data Plane

IETF RFC 5120 defines an optional M-IS-IS topology (with multiple topologies inside an IS-IS segment) that requires several extensions to packet encoding and other SPF procedures. Therefore, we added a new physical topology beside the basic IS-IS topology. The IS-IS standard topology was named as MT-0 for backward compatibility.

The topologies included PE routers connected to P routers in the transit network to support the different IS-IS implementations. The P routers were configured to support plain IPv6 traffic forwarding. While the PE routers were configured to create the L3VPN service at where the IPv6 header was a source address of PE Ipv6 loopback, the destination address of PE2's END.DT4 (or 6)/128 IPv6 address. This was from the locator range. the P routers still needed to install this destination address entry and forward on it.

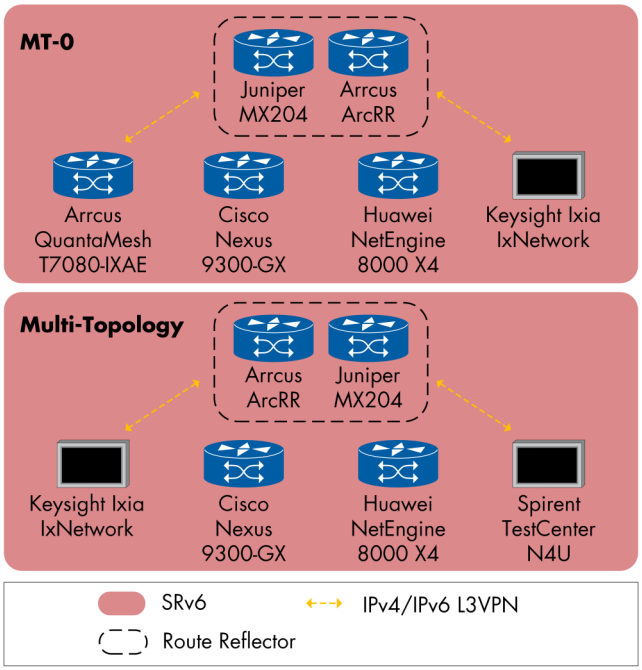

Figure 23: L3VPN over SRv6

In this case, the PE routers were determining from ISIS and BGP VPN next-hop, what destination address to use. We generated a combination of IPv4/IPv6 bidirectional traffic and verified the traffic flow between the PE pairs and we didn't observe any packet loss between the PE pairs.

The following devices successfully participated in the tests:

- P in MT-0: Arrcus ArcRR, Juniper MX204
- P in MT: Arrcus ArcRR, Juniper MX204
- PE in MT-0: Arrcus QuantaMesh T7080-IXAE, Cisco Nexus 9300-GX, Huawei NetEngine 8000 X4, Keysight Ixia IxNetwork
- PE in MT: Cisco Nexus 9300-GX, Huawei NetEngine 8000 X4, Keysight Ixia IxNetwork, Spirent TestCenter N4U

We tested EVPN over SRv6 between Huawei and Keysight. There were two service types the E-Line service and EVPN L3VPN. We also generated Inter-subnet traffic for EVPN L3VPN using both MAC/IP Route as well as IP prefix route. The physical topology was as simple as direct connectivity between the PEs.

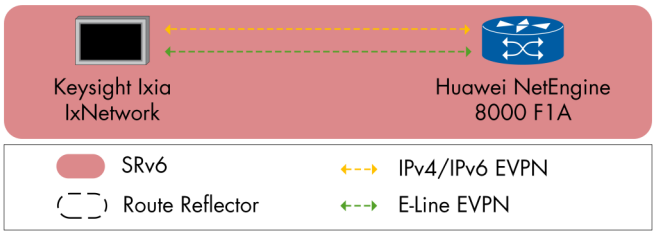

Figure 24: EVPN over SRv6

The following devices successfully participated in the test:

- PE: Huawei NetEngine 8000 F1A, Keysight Ixia IxNetwork

Topology Independent Loop Free Alternative

Topology Independent Loop Free Alternative (TI-LFA) provides full coverage of link and node protection in less than 50 msec. TI-LFA prevents micro-loops and traffic congestion due to the sub-optimal routing during a link or node failure. By enabling TI-LFA on the protecting node (aka Point Local Repair PLR), the IGP computes a backup path for each protected link or node. The backup path matches the loop-free post-convergence path and will be activated after detecting any failure against the protected link or node.

SR-MPLS and SRv6 facilitate the TI-LFA deployment by encoding the loop-free post-convergence path into the segment list. The segment list steers the traffic toward a specific repair node (aka PQ node) using Node-SID or a specific link by Adjacency-SID.

The main target of the test was to verify the functionality and interoperability of PLR to detect a local link failure and accordingly, activate a pre-computed loop-free backup path. The PLR inserts extra segments to steer the packets into the backup path. The 2020 test campaign involved a variety of TI-LFA scenarios. We run the test across different protection mechanisms like link and Shared Risk Link Group (SRLG), and with different data planes including SR-MPLS and SRv6.

TI-LFA for SR-MPLS

We built a physical full-mesh topology that consisted of four different network nodes to test link and SRLG TI-LFA over the SR-MPLS data plan. The participated vendors configured the network nodes with an L3VPN BGP-VPN service. We used Spirent TestCenter to generate IPv4 traffic from the emulated CEs.

Prior to the link failure, the ingress PE+PLR (network node 1) forwarded the traffic to the directly connected egress PE (network node 4). To simulate the link failure, we asked the vendor of network node 4 to shut down the link between network node 4 and network node 1 (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 22ms - 32 ms. The expected value was 50 ms.

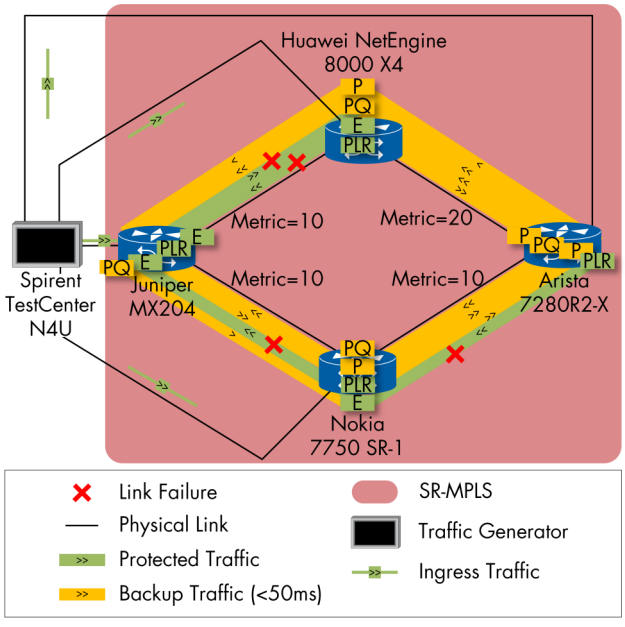

Figure 25: TI-LFA for SR-MPLS

The following devices successfully participated in the TI-LFA over SR-MPLS:

- Egress node: Huawei NetEngine 8000 X4, Juniper MX204, Nokia 7750 SR-1
- P node: Arista 7280R2, Huawei NetEngine 8000 X4, Nokia 7750 SR-1
- PLR: Arista 7280R2, Huawei NetEngine 8000 X4, Juniper MX204, Nokia 7750 SR-1
- PQ: Arista 7280R2, Huawei NetEngine 8000 X4, Juniper MX204, Nokia 7750 SR-1

One pair showed 63 ms out of service time. The expected value was 50 ms. During the TI-LFA with the SRLG test, we observed 20 ms out of service time for the L3VPN service.

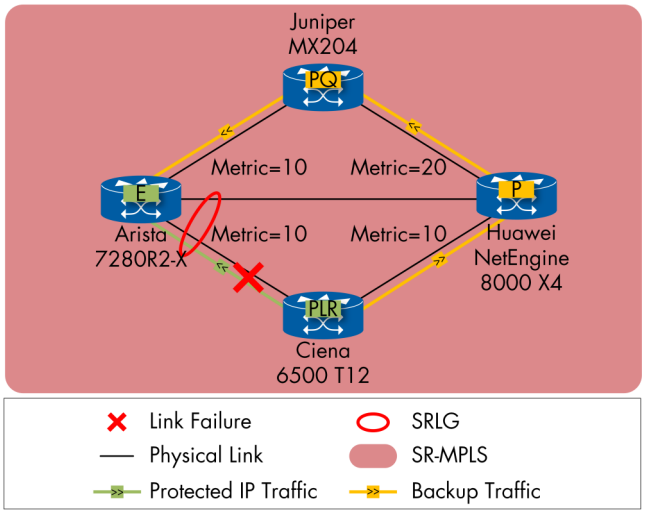

Figure 26: TI-LFA with SRLG for SR-MPLS

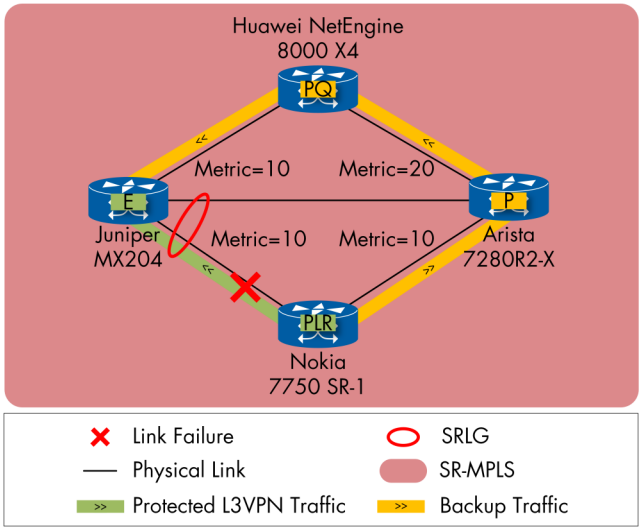

Figure 27: TI-LFA with SRLG for SR-MPLS

The following devices successfully participated in the TI-LFA with SRLG over SR-MPLS test:

- Edge node: Arista 7280R2, Juniper MX204
- P: Arista 7280R2
- PLR: Ciena 6500 T12, Nokia 7750 SR-1
- PQ: Huawei NetEngine 8000 X4, Juniper MX204

Two L3VPN pairs showed 1 - 4 s out of service time which was not included in the report. One pair with IP service showed 3 s out of service time which was due to convergence of IP and therefore understandable.

TI-LFA over SRv6

For the next-generation networks that are built on top of the SRv6 data plane, TI-LFA plays a fundamental role in service protection against link, node, and SRLG failures. TI-LFA SRv6 data plane applies the same concept of inserting extra SID to the SRH of the original packet by the PLR to steer the traffic toward the repair or PQ node. The SRv6 SID should be inserted or processed by SRv6 capable device. Thus, the implementation of TI-LFA in the SRv6 data plane requires the PLR and PQ nodes to be SRv6-aware network nodes.

This year, we built the topology for SRv6 TI-LFA SRLG testing to include SRv6 capable nodes as much as possible. All participated network nodes support SRv6 forwarding as a baseline function. Afterwards, the vendors configured their devices according to the assigned role as described below.

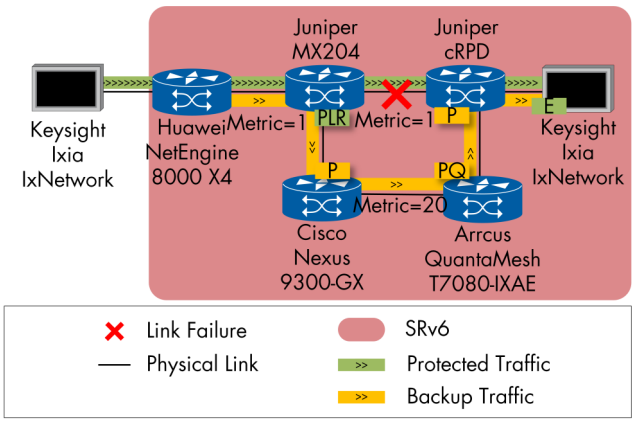

Figure 28: TI-LFA over SRv6

The following devices successfully participated in the test:

- Egress node: Keysight Ixia IxNetwork
- P node: Cisco Nexus 9300-GX, Huawei NetEngine 8000 X4, Juniper cRPD
- PLR: Juniper MX204
- PQ and PQ node and PSP (Penultimate Segment Popping) function: Arrcus QuantaMesh T7080-IXAE

In the SRLG test, we configured the link between PLR and two of the P nodes with the same SRLG ID, the PLR shall compute and install the backup path through the link between itself and the PQ node. To emulate the link failure, we ask the vendor to shut down the link between PLR and the P node while the traffic was flowing.

We recognized packet loss for a period of less than 50 ms, then the traffic was restored by directing the traffic toward PQ node. We identified that by monitoring the packet counters of the interface connecting the PQ node.

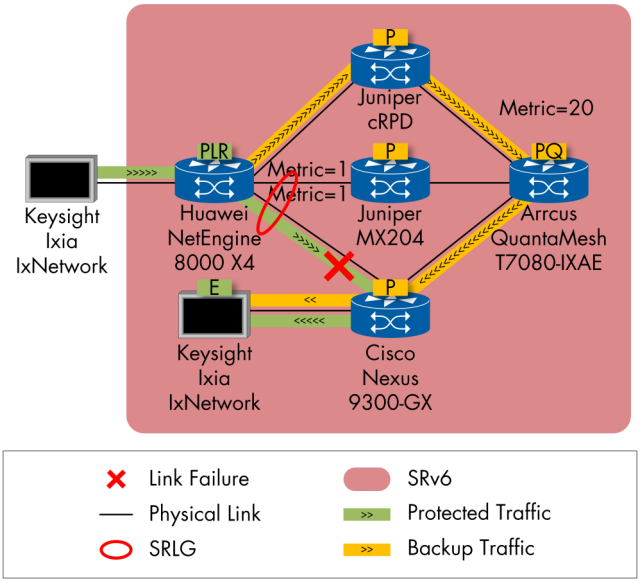

Figure 29: TI-LFA with SRLG over SRv6

The following devices successfully participated in the test:

- Egress node: Keysight Ixia IxNetwork
- P node: Cisco Nexus 9300-GX, Juniper cRPD, Juniper MX204
- PLR: Huawei NetEngine 8000 X4
- PQ node and PSP (Penultimate Segment Popping) function: Arrcus QuantaMesh T7080-IXAE

SR-TE and Seamless BFD

Seamless Bidirectional Forwarding Detection (S-BFD) is a light version of the classical BFD protocol. S-BFD designed to work with the Segment Routing Traffic Engineering (SR-TE) sessions for SR-TE path monitoring. If the S-BFD session fails, the S-BFD process brings down the SR-TE session. But, if the S-BFD session is recovered, the S-FBD process brings up the SR-TE session and preempts the primary role quickly. The efficiency of S-BFD sources from the lean mechanism of SR-TE path monitoring by implementing two logical nodes; the initiator and the reflector. The initiator sends the S-BFD packets toward the reflector through the monitored SR-TE policy. The reflector listens to the S-BDF packets and reflects back the packets toward the initiator.

For S-BFD testing purposes, we built a square physical topology which consisted of ingress PE, egress PE, and two P routers. We asked each pair of PEs to configure two SR-MPLS TE policies; one was the primary SR-MPLS TE and the other one was the backup. Using Keysight IxNetwork, we started generating IPv4 unidirectional traffic from PE1 (Initiator) to PE2 (Reflector). As expected, we observed the traffic was flowing through the primary SR-MPLS TE path. To emulate the SR-MPLS TE session tear down, we configured a simple packet filter on the node P1 to filter and drop the SR-MPLS frames. In less than 50 ms, the S-BFD session was down and brought down the primary SR-MPLS TE and we witnessed the traffic was switched over into the backup SR-MPLS TE.

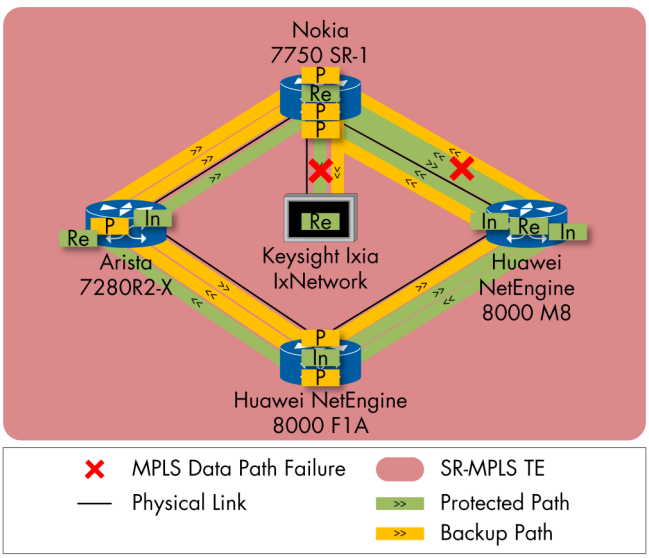

Figure 30: S-BFD over SR-MPLS with TE

After a while, we deactivated the packet filter and the S-BFD frames return to flow without any blocking. We checked the status of the S-BFD session and the primary SR-MPLS TE session, both of them were up. Hereafter, we confirmed the switch over of the traffic to the primary SR-MPLS TE path again.

The following devices successfully participated in the test:

- Initiator: Arista 7280R2, Huawei NetEngine 8000 F1A, Huawei NetEngine 8000 M8
- P node: Arista 7280R2, Huawei NetEngine 8000 F1A, Huawei NetEngine 8000 M8, Nokia 7750 SR-1
- Reflector: Arista 7280R2, Huawei NetEngine 8000 M8, Keysight Ixia IxNetwork, Nokia 7750 SR-1

SRv6 Traffic Engineering

SRv6 simplifies TE by steering traffic along any desired path in the network. The benefit differs from the traditional way to reserved resources on each node of RSVP-TE, the path is installed directly into the data packet.

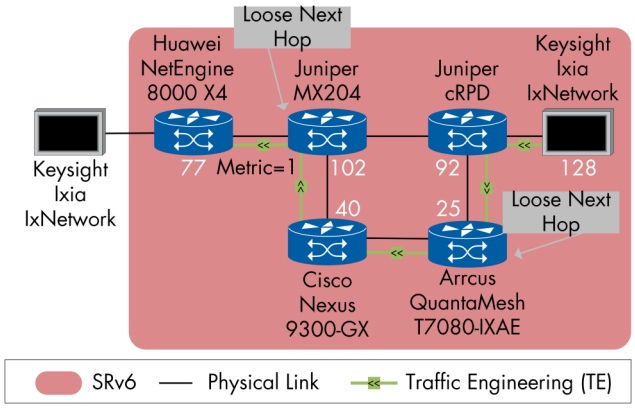

Figure 31: SRv6-TE

We performed a path selection test with SRv6 traffic engineering. Ixia generated packets with SRH, describing the loose next-hops that the packet shall traverse across the network. In this test case, Arrcus and Juniper shall be traversed, thus the END SIDs of Arrcus (fd01:0:25::1) and Juniper (fd01:0:102::1) are listed in SRH. As the packet traversed through Arrcus, END SID action was performed whereby SRH header along with outer IPv6 header is updated and the packet is forwarded onto Juniper. And as the packet traversed through Juniper, END SID action with PSP was performed, and the SRH header was removed from the packet. Therefore, each loose next-hop (Arrcus and Juniper) must have in its routing table information about how to handle packets received with END SID, and how to reach the SRv6 locator of the next loop hop on the path to the final destination.

The following devices successfully participated in the test:

- P with PSP function: Arrcus QuantaMesh T7080-IXAE, Juniper MX204
- P: Cisco Nexus 9300-GX, Juniper cRPD
- PE: Huawei NetEngine 8000 X4, Keysight Ixia IxNetwork

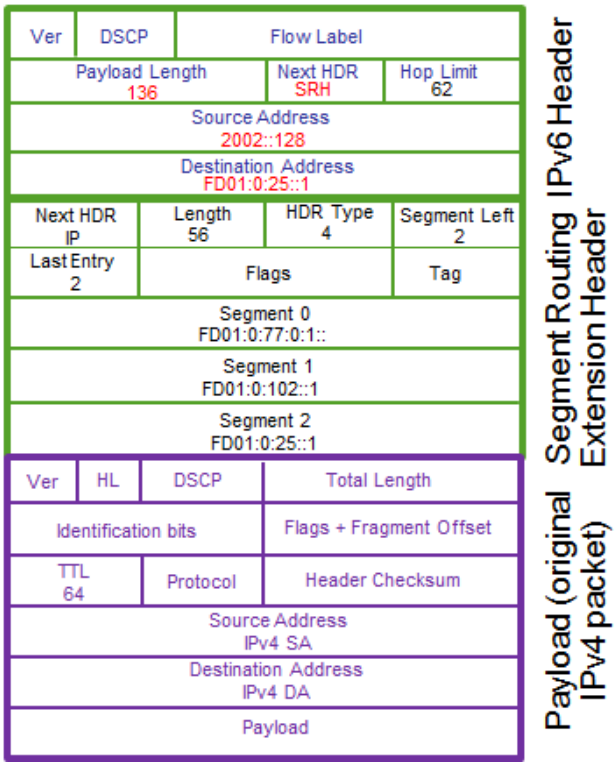

Figure 32: Incoming SRv6 Packet

Segment Routing Label Switched Path Ping/Traceroute

The RFC 8287 defines the LSP ping and traceroute method for Segment Routing with MPLS data plane. Similar to conventional LSP ping/traceroute, the SR fault detection and isolation tools are also based on MPLS echo request and echo reply. But Segment Routing LSP ping/traceroute include a new TLV type, the Segment ID sub-TLV.

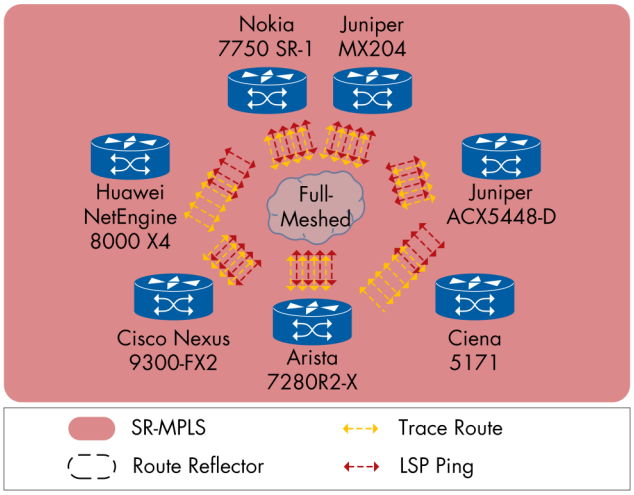

Figure 33: LSP Ping/Trace Route

On receipt of the sub-TLV carried in an MPLS echo request sent by the sender LSR, the LSR responder needs to check the segment ID obtained from the sub-TLV with the local advertised segment ID, to determine if the MPLS echo request has been forwarded from the correct path. The LSP ping/traceroute response is carried in an MPLS echo reply.

Based on the fully connected network between different vendors during the test, we tested the Segment Routing LSP ping/traceroute between vendors.

The following devices successfully participated in the tests:

- PE: Arista 7280R2, Ciena 5171, Cisco Nexus 9300 -FX2, Huawei NetEngine 8000 X4, Juniper ACX5448-D, Juniper MX204, Nokia 7750 SR-1

One pair failed on the traceroute without any re-test since the reason was unclear until the end of the test session.

BGP Segment Routing: BGP-Label Unicast

Segment Routing can be used in large scale Data Centers as a simple solution to provide traffic engineering and fast re-route capabilities in the DC fabrics. In this test, we verified that the overlay can be built using Multi-hop eBGP peering between Endpoints, and can use BGP-signaled MPLS LSPs as transport.

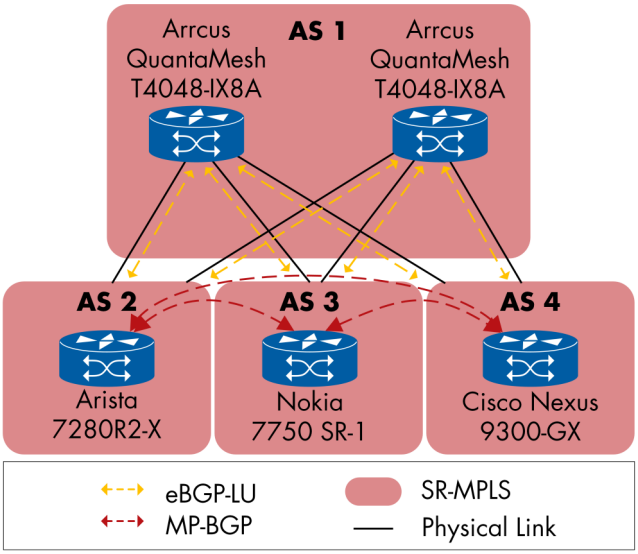

Figure 34: BGP-Label Unicast

We tested BGP Segment Routing using BGP Labeled Unicast (BGP-LU) NLRI in a typical Clos topology with two Spines and three Leaves. Vendors configured the Leaf nodes (DUTs) to advertise the optional and transitive BGP attribute; Prefix-SID attribute in the BGP-LU NLRI. Spine nodes were enabled with BGP-LU capability to forward the BGP update messages with MPLS labels. Additionally, Arrcus Spine node enabled BGP Segment Routing capability to generate MPLS labels for the BGP updates received from Leaf nodes. Using Spirent TestCenter, we generated full-mesh traffic between all Leaf nodes.

The following devices successfully participated in the test:

- PE: Arista 7280R2, Cisco Nexus 9300-GX, Nokia 7750 SR-1
- Spine: Arrcus QuantaMesh T4048-IX8A

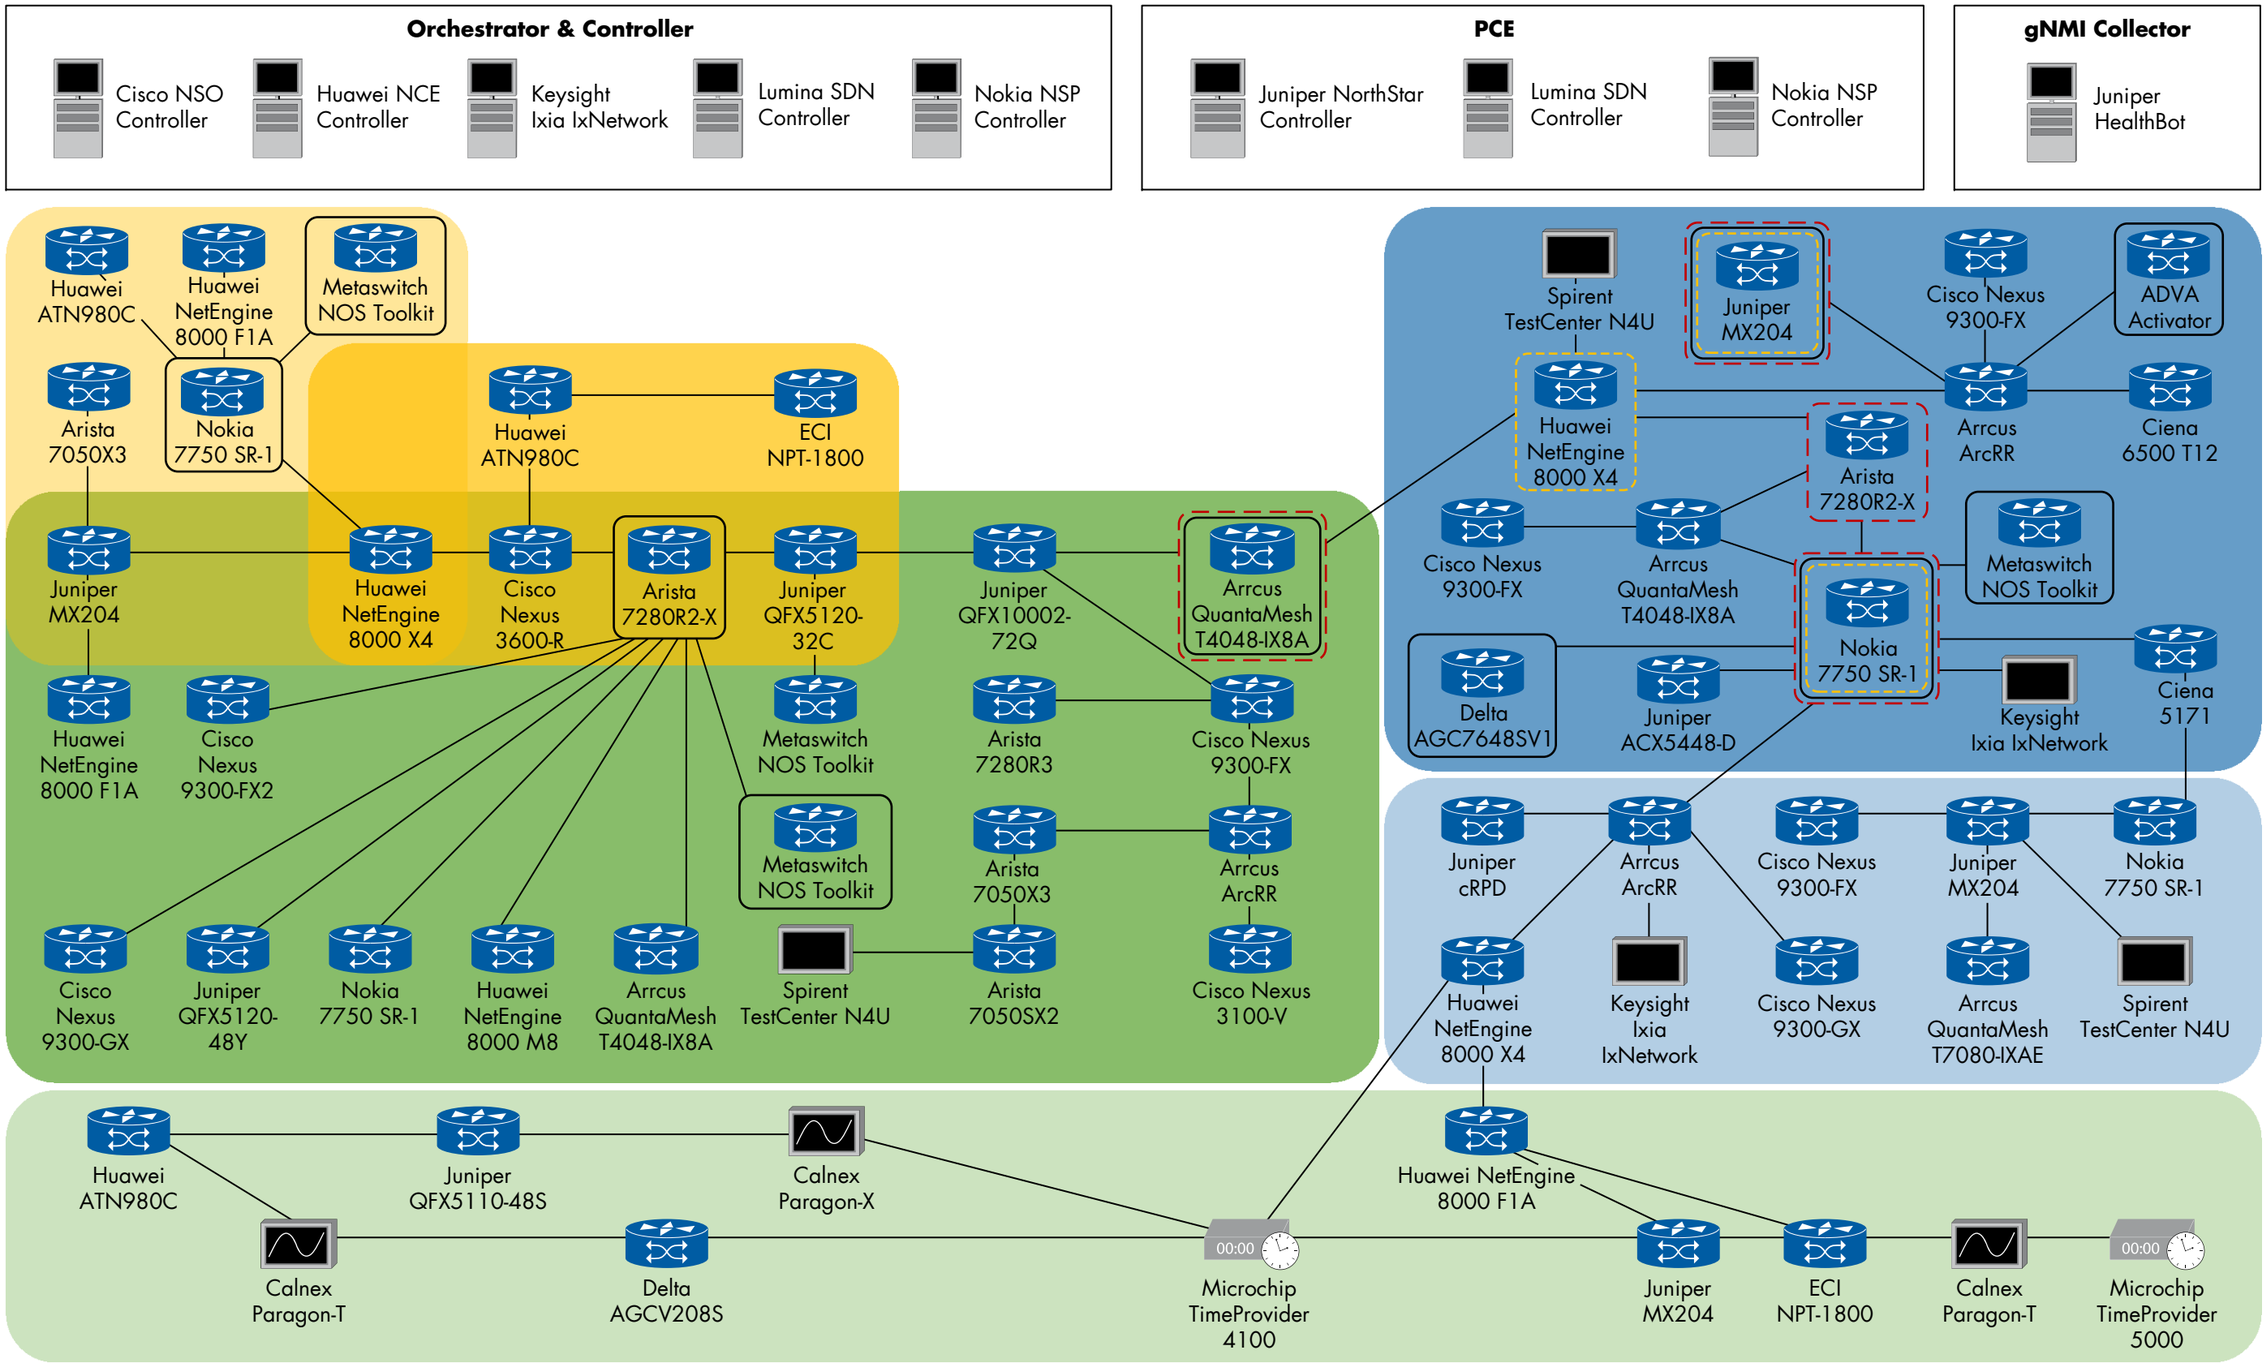
Orchestrator & Controller
Cisco NSO Controller
Huawei NCE Controller
Keysight Ixia IxNetwork
Lumina SDN Controller
Nokia NSP Controller
PCE
Juniper NorthStar Controller
Lumina SDN Controller
Nokia NSP Controller
gNMI Collector
Juniper HealthBot
Huawei ATN980C
Huawei NetEngine 8000 F1A
Metaswitch NOS Toolkit
Arista 7050X3
Nokia 7750 SR-1
Huawei ATN980C
ECI NPT-1800
Juniper MX204
Huawei NetEngine 8000 X4
Cisco Nexus 3600-R
Arista 7280R2-X
Juniper QFX5120-32C
Juniper QFX10002-72Q
Arrcus QuantaMesh T4048-IX8A
Huawei NetEngine 8000 F1A
Cisco Nexus 9300-FX2
Metaswitch NOS Toolkit
Arista 7280R3
Cisco Nexus 9300-FX
Metaswitch NOS Toolkit
Arista 7050X3
Arrcus ArcRR
Cisco Nexus 9300-GX
Juniper QFX5120-48Y
Nokia 7750 SR-1
Huawei NetEngine 8000 M8
Arrcus QuantaMesh T4048-IX8A
Spirent TestCenter N4U
Arista 7050SX2
Cisco Nexus 3100-V
Spirent TestCenter N4U
Juniper MX204
Cisco Nexus 9300-FX
ADVA Activator
Huawei NetEngine 8000 X4
Arrcus ArcRR
Ciena 6500 T12
Arista 7280R2-X
Cisco Nexus 9300-FX
Arrcus QuantaMesh T4048-IX8A
Metaswitch NOS Toolkit
Delta AGC7648SV1
Juniper ACX5448-D
Nokia 7750 SR-1
Keysight Ixia IxNetwork
Ciena 5171
Juniper cRPD
Arrcus ArcRR
Cisco Nexus 9300-FX
Juniper MX204
Nokia 7750 SR-1
Huawei NetEngine 8000 X4
Keysight Ixia IxNetwork
Cisco Nexus 9300-GX
Arrcus QuantaMesh T7080-IXAE
Spirent TestCenter N4U
Huawei ATN980C
Juniper QFX5110-48S
Calnex Paragon-X
Huawei NetEngine 8000 F1A
Calnex Paragon-T
Delta AGCV208S
Microchip TimeProvider 4100
Juniper MX204
ECI NPT-1800
Calnex Paragon-T
Microchip TimeProvider 5000
Core Router
Grandmaster
Physical Link
NETCONF/YANG
Data Center 1 - VxLAN
Clocking
Emulator
Slave Clock
Controller/ Management server/PCE
PCEP
Data Center 2 - MPLS
Segment Routing MPLS
SDN Edge Router
IP/MPLS
Segment Routing – SRv6

SDN

Nowadays, business requirements are changing rapidly. Service providers have to adjust to these changes to accommodate market needs. Having a centralized network management protocols and service orchestration is a key point to achieve this flexibility. The following section describes the Path Computation Element Protocol and NETCONF/YANG interoperability tests, results and interoperability findings. The tests were chosen to adhere to market needs and serve as proof that SDN provides a credible approach to current challenges. In short, some interoperability issues were found. However, the vendors managed to solve most of them. Some vendors are still missing some features that prevented the execution of some combinations. In general, the test results presented in this section show a wide range of interoperability between vendors in multiple scenarios including some advanced cases.

Path Computation Element Protocol

A Path Computation Element (PCE) is a centralized controller that is capable of computing a network path and applying computational constraints. A Path Computation Client (PCC) is a client application requesting a path computation to be performed by a PCE. The communication mechanism between a PCE and a PCC is the TCP-based PCEP protocol, as defined in RFC5440 and extended in RFC8231 (Stateful PCE Support), RFC8281 (PCE Initiated LSP/Path Support) and RFC8664 (Segment Routing LSP/Path Support).

This test area confirms the interoperability between different PCE and PCC solutions during the establishment of traffic-engineered LSPs (TE LSPs) in the MPLS domain.

PCE-initiated Paths in a Stateful PCE Model

In an MPLS domain, the dynamic Lable Switched Path (LSP) creation or teardown plays an important role in application-based traffic engineering. A possible use case is where an application can request a path with certain constraints between two network nodes by contacting the PCE. In the case of a failure scenario between two nodes, the backup path needs to be created dynamically by the PCE. In this test, we verified the LSP setup, state synchronization, update and deletion of PCE-initiated LSP without needing any local configuration of PCC. The test topology included one centralized PCE and two PE nodes acted as PCC. Additionally, two transport nodes are also used in the test to re-optimize the LSP.

To provide routing for the internal provider network, PE nodes and the transport nodes (node 3 and node 4) are configured with ISIS-SR in the MPLS domain.

As the LSP Path Setup Type (PST), two PCE implementations supported both SR-TE and RSVP-TE and one PCE implementation supported only RSVP-TE. The nodes -PCCs and Network nodes- in the MPLS domain were configured with ISIS-SR and ISIS-TE and RSVP respectively. The routing and traffic engineering topology information were advertised from the network nodes to the PCE using the BGP Link-State (BGP-LS) address-family. This information was used by the PCE to calculate and set up LSPs/Paths in the PCC nodes. To verify the LSP setup, the L3VPN service was configured with mp-BGP service signaling protocol. The lowest IGP cost was set as constraining LSP to follow the optimal path.

After the successful PCEP session establishment between the PCE and the PCCs, the PCE computed the LSP path based on the obtain information via BGP-LS and sent LSP initiate request message to PCCs. Based on the information, each PCC in the topology created one transport path in the direction facing the next PCC. We verified created path information in the PCE and PCC side and confirmed the successful traffic flow between the PCCs via the L3VPN service. In the test for one of the PCE, the original path was configured through node 3. For other PCE solutions, the original path was configured directly between PCCs as in the figure below.

For the LSP state synchronization, we terminated the PCEP session first and waited until the timeout for clearing the LSP from both PCCs. Once the LSP status was cleared form the both PCCs -after the State Timeout Interval expired- we re-established the PCEP session between PCE and PCCs. During the synchronization phase, the PCCs sent a report message to the PCE to inform about the LSP state in the PCCs. The PCE identified from the report that no LSP was available in the PCCs -as the LSP state was removed after the State Timeout Interval expired- and proceed re-instantiate the missing LSPs that were previously created in both PCCs. The traffic flow between the PCCs was seamless. This procedure confirmed that the PCE and PCCs are synchronized correctly via the PCEP session.

For the LSP update, we increased the IGP cost in the original path between the two nodes. Meanwhile, PCCs delegated the LSP to the PCE for the LSP update. In this case, PCE identified the path constraint automatically and re-optimize the path and sent the optimized path details to the PCCs. Both PCCs created a new transport path in the direction facing the next PCC

through the node 3 and node 4 for once PCE solution as in the figure below. For other PCE solution, the PCC created the path through the node 3. We confirmed the successful traffic flow between the PCCs.

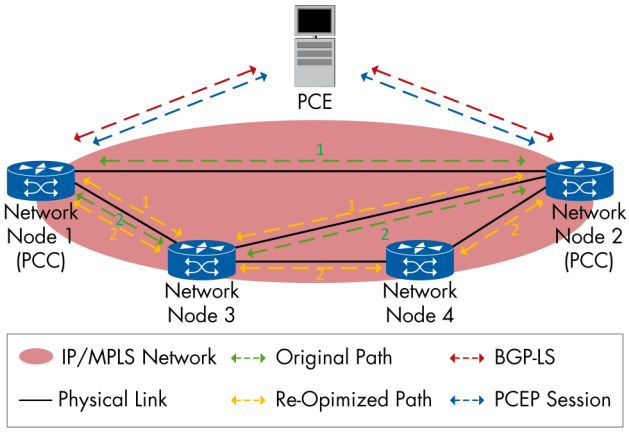

Figure 35: PCE-initiated Paths in a Stateful PCE model

During the LSP state synchronization test with Nokia's PCE, one PCC solution sent the LSP state report with the delegation flag with 0 after the PCEP session was back again and the PCC solution expects from PCE to send a PCinitiate message for setting the delegation in the LSP state report. Since the delegation flag was set to zero from PCC in the LSP state report, the PCE was not able to update or delete the LSPs in the PCC. When the PCEP session was continuously up and running, the LSP state report was successfully reported with the delegation bit one. In this test, Huawei PCC was able to set the delegation bit to one in the LSP state report after the PCEP session was up.

During the LSP deletion with Nokia PCE, Nokia PCE sent the PCinitiate message with setting the R flag to one. After deleting the LSP, one PCC sent the LSP state report with non zero LSP-ID. The LSP-ID should be zero in the LSP state report. This information will be used by the PCE to remove the LSP in the PCE. Since the LSP-ID was not set to zero, PCE was not able to delete the LSP in the PCE side.

In this test, Juniper PCC was successfully sent the LSP-ID with zero value. In this test, Juniper and Nokia PCE were successfully deleted the LSP in the PCCs with setting the delegation bit to one. During the path re-optimization, Nokia PCE and Lumina SDN Controller were able to optimize the path automatically once we change the metric of the path.

PCC-initiated Paths in a Stateful PCE Model

In some useful scenarios, PCC requests the PCE for setting up an LSP and after getting a path from the PCE, it may delegate the control later to the active stateful PCE. According to RFC8231, the PCC sends a path computation element request message (PCReq). The PCReq message has been extended in RFC5440 to include the LSP object. When the PCE receives the PCReq message, it will compute the LSP and send it to the PCC.

This test verified the PCC capability for requesting an SR-TE/RSVP-TE path in a single IGP domain. Furthermore, we tested LSP updates with LSP delegation, path optimization with LSP revocation.

The test setup included one centralized PCE and two PE nodes acted as PCC. Meanwhile, one transport node is also used in the test for the path re-optimization purpose. PE nodes and the transport nodes are configured with ISIS-SR in the MPLS domain to provide routing for the internal provider network. This routing information was used to generate labels. The routing information was synchronized with PCE using the BGP-LS mechanism. As the LSP signaling protocol, two PCE solutions configured with SR-TE. To verify the installed LSP, the L3VPN service was configured with mp-BGP service signaling protocol. The lowest IGP cost was set as constraining LSP to follow the optimal path.

The test started with the PCEP session establishment between the PCE and the PCCs. After the successful PCEP session establishment, PCCs requested to PCE for

PCE	PCC1	PCC2	Node3	Node4
Nokia NSP Server	Juniper MX204-6	Huawei NetEngine 8000 X4	Juniper MX204-4	Nokia 7750 SR-1-2
Juniper NorthStar Controller	Nokia 7750 SR-1-2	Huawei NetEngine 8000 X4	Juniper MX204-4	N/A
Lumina SDN Controller	Juniper MX204-4	Huawei NetEngine 8000 F1A	Nokia 7750 SR-1-2	N/A

Table 2: PCE-initiated Paths in a Stateful PCE Model - Successful Combinations

the path computation. Upon reception of the request, PCE computed the LSP path based on the obtain information via BGP-LS and sent the path details to the PCC. Using the PCE computed path information, each PCC in the topology created one transport path in the direction facing the next PCC. We confirmed the successful traffic flow between the PCCs via the L3VPN service. After successful path creation, we increased the IGP cost in the links between the PE nodes. Meanwhile, PCCs delegated the LSP to the PCE for the LSP update. In this case, PCE identified the path constraint automatically and re-optimize the path and sent the optimized path details to the PCCs. Both PCCs created a new transport path in the direction facing the next PCC through the node 3. The traffic flow was successful between the PCCs via the node 3. In the case of LSP revocation, PCCs sent the LSP revocation (no delegation) report to the PCE. After the LSP revocation, we modified the IGP cost on the current path. Since the LSP delegation was inactive, PCE was not able to optimize the path itself. We performed the re-optimization command on the PCC and Nokia PCC sent a PCReq to Nokia PCE for the path re-optimization. Upon reception of the PCC request, PCE updated the LSP and sent the update message to PCCs. The PCCs created LSP in the direction facing the next PCC as in the initial stage. The traffic flow was seamless between the PCCs.

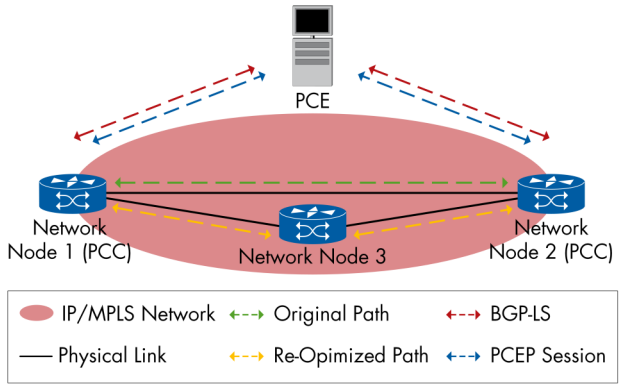

Figure 36: PCC-initiated Paths in a Stateful PCE model

During the LSP revocation, one PCC solution achieved the LSP revocation by sending the LSP state as a PC_report only massage which accepted by the PCE to revocate the LSP delegation.

BGP Flowspec for IPv6

BGP Flowspec defines a new Multi-Protocol BGP (MP-BGP) Network Layer Reachability Information (NLRI) as specified in RFC5575. The new NLRI collects layer 3 and layer 4 details that are used to define a Flow Specification. The actions are assigned on the router based on the flow specification.

The test topology consists of one Flowspec controller and two PE nodes. As the first step, the BGP session was established between the PE nodes. After the BGP session establishment, we generated 4 different UDP traffic streams between the same source and destination IPv6 address with different ports and DSCP values.

All the four traffic streams were successfully forwarded without any loss. In the next step, we established the BGP session between the controller and the PE nodes and configured the two Flowspec rules in the controller. The first rule is defined to match a destination IPv6 address and a UDP port ID for rate-limiting scenario and the second rule is defined to match a destination IPv6 address and a DSCP value for full traffic drop. After applied the Flowspec policies by the controller, we generated the same traffic streams again. Arista 7280R2-X node met the two Flowspec conditions and limited the traffic rate to 128 Kbit/s for one traffic stream and dropped all the packets for another stream. The remaining two streams were not affected as expected. Since the Arrcus supports only drop Flow-spec policy, the packet drop condition was met in the streams as specified by the controller.

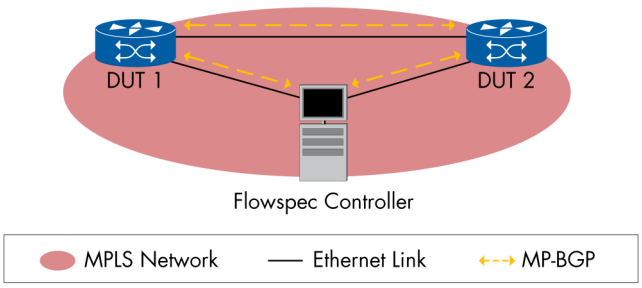

Figure 37: BGP Flowspec for IPv6

PCE	PCC1	PCC2	Node3
Nokia NSP Controller	Huawei NetEngine 8000 F1A	Nokia 7750 SR-1-2	Juniper MX204-4

Table 3: PCC-initiated Paths in a Stateful PCE Model - Successful Combinations

Flowspec Controller	DUT1	DUT
Keysight Ixia IxNetwork	Arrcus QuantaMesh T4048-IX8A	Arista 7280R2-X
Keysight Ixia IxNetwork	Arista 7280R2-X	Arrcus QuantaMesh T4048-IX8A

Table 4: BGP Flowspec for IPv6 - Successful Combinations

Egress Peer Engineering with Segment Routing

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 SDN controller to program any egress peer policy at ingress border routers or 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 to compute efficient BGP Peering Engineering policies and strategies.

This test verifies that the EPE Segment Routing could be used to allocate MPLS label for each engineered peer and use a Lable stack to steer traffic to a specific destination. The test topology included a centralized EPE controller and three Autonomous Systems (AS). In the AS 65001, the DUT1 was configured as a PE router, and DUT2 was configured as ingress Autonomous System Boundary Router (i_ASBR). Meanwhile, DUT3 and DUT4 are located in different AS, namely AS 65002 and AS 65003. The ingress port of DUT 1, egress ports of the DUT3 and DUT4 were connected to the traffic generator.

The reachability information was provided by a peer AS on all data-plane interconnection links of DUT2 using an eBGP Network Layer Reachability Information (NLRI) advertisement.

The EPE controller learned the BGP Peering SID's and the external topology of DUT2 via BGP-LS EPE routes.

DUT1 and DUT2 are configured with ISIS-SR in the MPLS domain to provide routing for the internal network. In the DUT2, the paths to the DUT3 and DUT 4 are configured with different SR color policy and set the DUT3 as the best transport path. We generated the IPv4 service traffic flow between the DUT1 and DUT3 as well as DUT1 and DUT4. The traffic flow routed through the DUT3 as expected. During the traffic flow, the EPE controller pushed the color based SR policy on the DUT1 using BGP-SR Policy to choose the DUT4 as the transport port. After successfully applied the policy, traffic was re-routed through the DUT4. This test confirmed that DUT1 successfully encapsulated the traffic for delivery over the designated data-plane interconnection link.

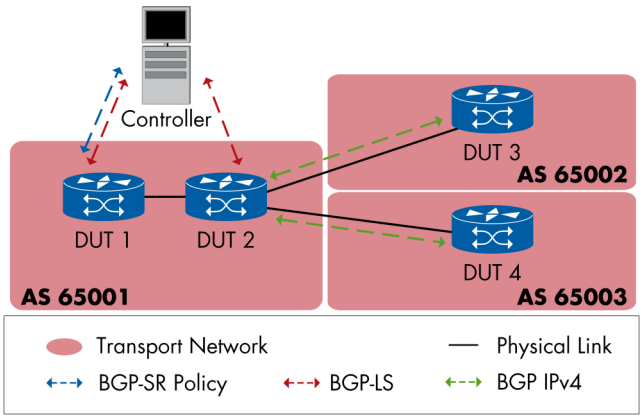

Figure 38: Egress Peer Engineering with Segment Routing

EPE Controller	DUT1	DUT2	DUT3	DUT4
Keysight Ixia IxNetwork	Nokia 7750 SR-1-2	Juniper MX204-4	Arrcus QuantaMesh T4048-IX8A	Arrcus QuantaMesh T4048-IX8A
Nokia NSP Controller	Arista 7280R2-X	Juniper MX204-4	Arrcus QuantaMesh T4048-IX8A	Arrcus QuantaMesh T4048-IX8A
Nokia NSP Controller	Nokia 7750 SR-1-2	Juniper MX204-4	Arrcus QuantaMesh T4048-IX8A	Arrcus QuantaMesh T4048-IX8A

Table 5: Egress Peer Engineering with Segment Routing - Successful Combinations

NETCONF/YANG

To simplify and automate network device configuration the IETF has developed a Network Configuration Protocol (NETCONF) and a modeling language (YANG). This approach helped service providers and enterprises to cut the time, cost and the manual steps needed for network configuration.

In this test section, we provided a combination of NETCONF and RESTCONF protocols with different YANG modules. Our main focus was to test L2VPN, L3VPN and EVPN network services.

Multi-Vendor/Multi-Domain Controllers Orchestration

Network operators fragment their transport networks into multiple vendor domains and each vendor offers its SDN controller to manage their network components. Multi-domain controller's orchestrator allows operators for simpler networking control and provision of end-to-end services across the multi-domain networks regardless of the control plane technology of each vendor. In this test, we provisioned different services using the multi-domains controller and verified end-to-end connectivity via the traffic flow. NETCONF/RESTCONF was used as a management protocol between domain controllers and between the controllers and the Orchestrator.

This test topology included one Multi-Domain Orchestrator and two controllers, one for each domain. Each domain had two network nodes. To provide routing for the internal provider network, the nodes in MPLS Domain 1 were configured with OSPF as IGP and t-LDP as the service signaling protocol, meanwhile, the nodes in MPLS Domain 2 were configured with ISIS-SR. We verified the NETCONF session between the Domain controller and the DUTs.

Between Cisco's Multi-Domain Orchestrator and the Cisco's Domain controller (Domain controller1), the NETCONF session was established. Meanwhile, a RESTCONF-like connection was enabled between the Cisco's Multi-Domain Orchestrator and the Lumina's Domain controller (Domain controller 2). The Cisco's Domain controller was configured to establish the L2VPN service between the DUT1 and DUT2. At the same time, Lumina's Domain controller was configured to establish an EVPN service between the DUT3 and DUT4. The Multi-Domain Orchestrator was set up to trigger the defined services in the Domain controller. After triggering Multi-Domain Orchestrator, we verified the service creation in the DUTs of each domain and also verified successful service removal from all nodes and stopped end-to-end traffic flow. This test confirmed the integration of Multi-Domain Orchestrator with domain controllers to established different services in different domains.

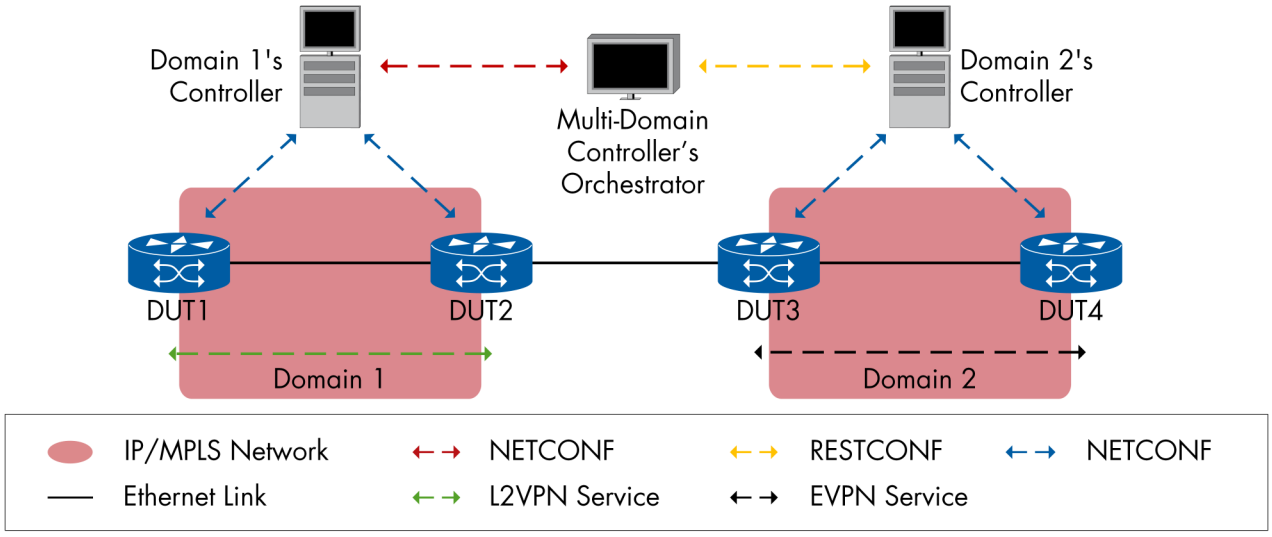

Figure 39: Multi-Vendor/Multi-Domain Controllers Orchestration

Multi-Domain Controller's Orchestrator	Domain1's Controller	Domain2's Controller	DUT1	DUT2	DUT3	DUT4
Cisco NSO Orchestrator	Cisco NSO Controller	Lumina SDN Controller	Metaswitch NOS Toolkit	Nokia 7750 SR-1	Nokia 7750 SR-1-2	Juniper MX204-6

Table 6: Multi-Vendor/Multi-Domain Controllers Orchestration - Successful Combinations

Device Configuration Using NETCONF/YANG

The NETCONF protocol defines a simple mechanism through which a network device can be managed, configuration data information can be retrieved and new configuration data can be uploaded and manipulated. The protocol allows the device to expose a full and formal application programming interface (API). Applications can use this straightforward API to send and receive full and partial configuration data sets.

In this test, We defined a set of configurable elements on the DUTs and used NETCONF protocol from a compliant client to change the parameters on the DUTs, which runs the NETCONF server. The NETCONF client was connected to the DUTs (NETCONF servers) and established the NETCONF sessions between them. After the session establishment, the NETCONF client successfully retrieved the device running configuration and also NETCONF client was able to retrieve specific information like interface details using subtree filtering. Then we verified some configuration change on the DUTs which was performed by the NETCONF client. We also confirmed that the NETCONF client was able to retrieve some operational status and roll back the configuration to its previous state. Finally, the NETCONF client terminated the NETCONF session successfully.

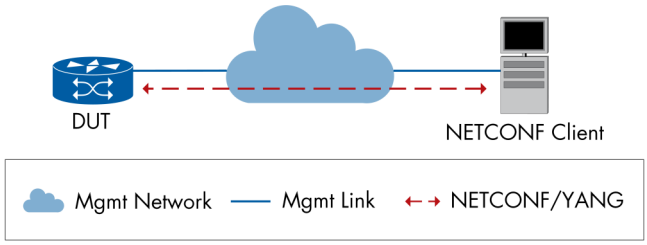

Figure 40: NETCONF/YANG

L2VPN Service Creation Using NETCONF/YANG

YANG is a data modeling language that was introduced to define the contents of a conceptual data store that allows networked devices to be managed using NETCONF. In this test, we verified that the IETF L2VPN service YANG model (RFC 8466) can be used to configure and manage L2VPNs. It verified VPWS specific parameters as well as BGP specific parameters applicable for L2VPNs. A NETCONF compliant client was used as a centralized controller to configure a group of PE nodes and provision L2VPN services.

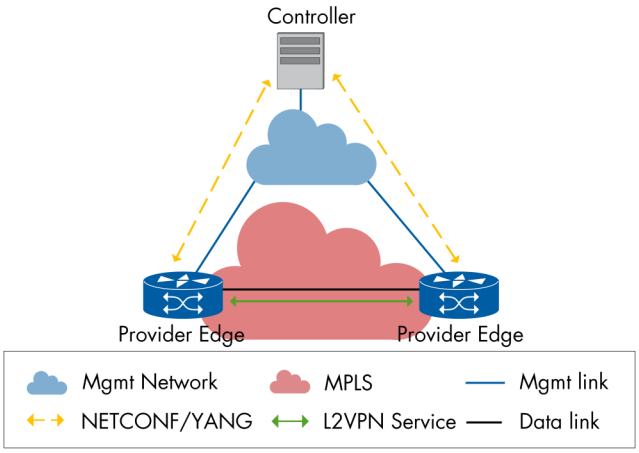

Figure 41: L2VPN Service Creation Using NETCONF/YANG

Controller	PE1	PE2
Cisco NSO Controller	Nokia 7750 SR-1	Metaswitch NOS Toolkit
Cisco NSO Controller	Juniper MX204-6	Nokia 7750 SR-1

Table 8: L2VPN Service Creation Using NETCONF/YANG - Successful Combinations

Controller	DUT1	DUT2	DUT3	DUT4	DUT5	DUT6
Huawei NCE Controller	Nokia 7750 SR-1	N/A	N/A	N/A	N/A	N/A
Lumina SDN Controller	Juniper MX204-6	Metaswitch NOS Toolkit	Nokia 7750 SR-1	N/A	N/A	N/A
Cisco NSO Controller	ADVA Activator	Arista 7280R2-X	Arrcus QuantaMesh	Juniper MX204-6	Metaswitch NOS Toolkit	Nokia 7750 SR-1

Table 7: Device Configuration using NETCONF/YANG

In this test, we verified YANG model that can be used to configure and manage L2VPN service. Meanwhile, we verify VPWS specific parameters as well as BGP specific parameters applicable for L2VPNs. The topology consists of one NETCONF client and two NETCONF servers (DUTs) in the MPLS domain. The NETCONF session was established between the NETCONF client and server. After the session establishment, we verified that the complete configuration from the DUTs was retrieved and synchronized with the NETCONF client's config database. Using the NETCONF client, we initiated a VPWS-L2VPN instance configuration on the DUTs and NETCONF client showed the service status as up. We confirmed the status of the service creation on the DUTs as well. The traffic flow was seamless between the DUTs. We requested the NETCONF client to delete the previously configured service and verified that the current configuration is identical to the initial configuration. Finally, we confirmed that no traffic flow was between the DUTs.

L3VPN Service Creation Using NETCONF/YANG

This test verifies that the IETF L3VPN YANG model (RFC 8299) can be used to configure and manage the L3VPN service. Meanwhile, we verified VRF specific parameters as well as BGP specific parameters applicable for L3VPNs. The procedure of the test is very similar to the previous test. In this, we verified L3VPN service creation and deletion while verifying that the required traffic was flowing as expected.

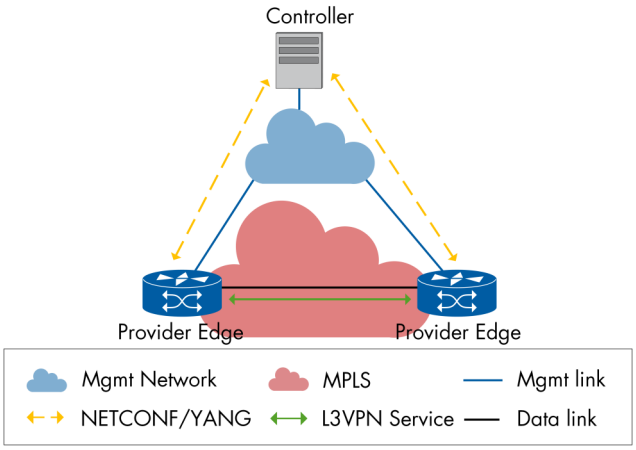

Figure 42: L3VPN Service Creation Using NETCONF/YANG

Controller	PE1	PE2
Cisco NSO Controller	Nokia 7750 SR-1	Juniper MX204-6
Cisco NSO Controller	Nokia 7750 SR-1	Metaswitch NOS Toolkit

Table 9: L3VPN Service Creation Using NETCONF/YANG - Successful Combination

EVPN Service Creation Using NETCONF/YANG

This test verifies that a service YANG model can be used to configure and manage EVPN service. The service model does not contain most of the device level connection details, so the controller must compute and allocate resources in order to set up the service.

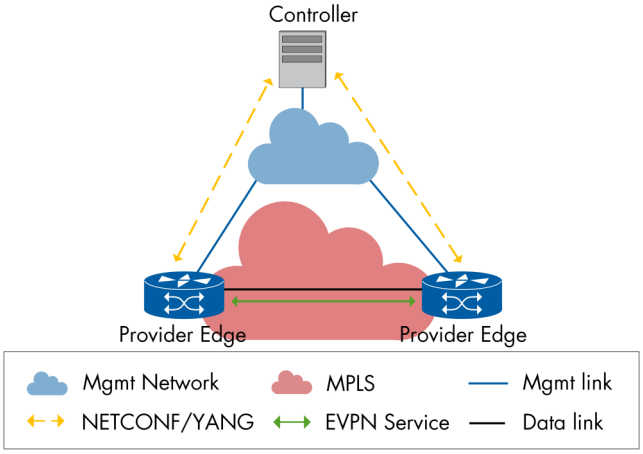

Figure 43: EVPN Service Creation Using NETCONF/YANG

Controller	PE1	PE2
Cisco NSO Controller	Arrcus QuantaMesh T4048-IX8A	Metaswitch NOS Toolkit
Cisco NSO Controller	Metaswitch NOS Toolkit	Nokia 7750 SR-1
Cisco NSO Controller	Nokia 7750 SR-1	Juniper MX204-6
Lumina SDN Controller	Nokia 7750 SR-1-2	Juniper MX204-6

Table 10: EVPN Service Creation Using NETCONF/YANG - Successful Combination

This test verifies the YANG model that can be used to configure and manage EVPN service. Meanwhile, we verified MPLS/VXLAN specific parameters as well as BGP specific parameters applicable for EVPN. The procedure of the test is very similar to the L2VPN service creation test. In this, we verified EVPN service creation and deletion while verifying that the required traffic was flowing as expected.

OpenConfig and Streaming Telemetry

The Streaming telemetry is the data model to facilitate operational data monitoring with higher efficiency. The network devices work in push mode to send the network operation data to the collector, instead of pull mode comparing with SNMP or CLI. The OpenConfig data models work with different transport protocols, such as NETCONF/RESTCONF, gNMI.

This test verifies the Streaming telemetry for interface monitoring using OpenConfig data models with the transport protocol gRPC Network Management Interface (gNMI). The gNMI defines a particular set of gRPC operations such as Capability Request, Get Request, Set Request, and Subscribe Request. The test topology consists of one collector and two PE nodes. The PE nodes are configured with ISIS in the MPLS domain to provide routing for the internal provider network. The gRPC service was configured between the collector and the PE nodes. In this test, the YANG model defined in openconfig-if-ethernet.yang: version 2.7.2 was used to collect the bandwidth data reported by the PE nodes. After subscribing to the telemetry service of the PE nodes, PE nodes streamed interface statistics to the collector via gRPC and collector got a stream to read a sequence of messages back. The Juniper Collector (Healthbot) successfully collected the interface's incoming/outgoing octets of ADVA and Delta PE nodes and calculated the octets difference with the time. These statistics will be used for the alarming with the threshold values.

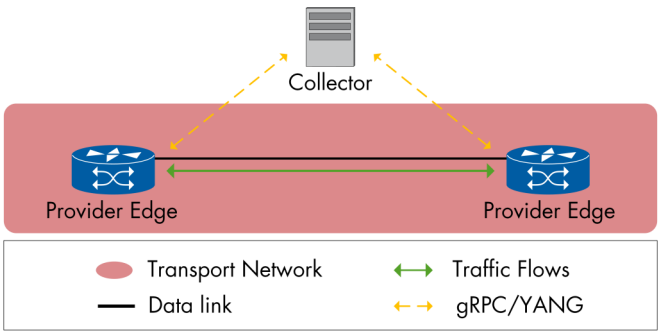

Figure 44: Openconfig and Streaming Telemetry

Controller	PE1	PE2
Juniper HealthBot	ADVA Activator	Delta AGC7648SV1-2

Table 11: Openconfig and Streaming Telemetry - Successful Combination

FlexEthernet (FlexE) Channelization

FlexE Channelization and Physical Isolation

Channelization was striking with VPN tests, particularly noticeable with Flexible Ethernet (FlexE) as defined in the implementation Agreement by OIF (Optical Internetworking Forum). FlexE allows channelization of one or more physical links to support different Ethernet MAC rates (e.g. 10G, 40G, nx25G) over the single or multiple Ethernet links. Especially, users can configure several FlexE tunnels with different bandwidth for different client services.

In this test, we deployed 5 FlexE tunnels to carry 5 different channels with a granularity of 50, 20, 15, 10 and 5 Gbit/s L2VPN traffic and verified the channelization in a 100G back-to-back Flex scenario. We first performed a frame loss test to all channels with a full load (overheads Bytes excluded) to verify the baseline physical isolation of different FlexE tunnels, where 0 frame loss was expected. Then, we performed the frame loss test under an overload condition. This traffic included the same traffic as before and increased one flow in a selected FlexE tunnel to verify that the overload shall not affect other non-participating channels. Only overloaded tunnels shall show frame loss. All test results were expected.

The following devices successfully participated in the test:

- PE: ECI NPT-1800, Huawei ATN980C

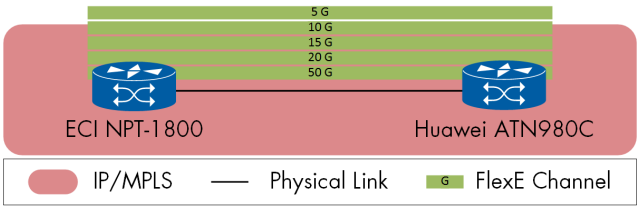

Figure 45: FlexE Channelization and Physical Isolation

Name	Rx Rate (Mbps)	
100G Port VLAN-ID1 50G BiDi traffic	94,440.985	0 Frame Lost
100G Port VLAN-ID4 20G BiDi traffic	38,496.239	
100G Port VLAN-ID3 15G BiDi traffic	28,872.18	
100G Port VLAN-ID2 10G BiDi traffic	19,248.119	
1G BiDi oversubscription	1,889.688	1.8% Frame Lost
100G Port VLAN-ID1 5G BiDi traffic	9,624.06	0 Frame Lost

Figure 46: Keysight Traffic Statistic for Overload Scenario

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.

We verified the FlexE capability of dynamic bandwidth adjustment on 100G ports in a back-to-back FlexE scenario. The setup included two FlexE channels: FlexE tunnel 1: 10G; FlexE tunnel 2: 20G; Just after that, we determined no frame loss in the baseline setup (tunnel 1: 10G; tunnel 2: 20G), we first increased the tunnel 2 traffic to 30 Gbit/s and expected 10 Gbit/s traffic drop, to ensure that tunnel 2 was configured with 20G. As expected, 10 Gbit/s traffic was dropped.

Then we asked the vendors to increase the bandwidth of tunnel 2 to 30G. We observed the bandwidth adjustment status via CLI then sent the same traffic as described in the previous step. We observed 0 packet lost as expected.

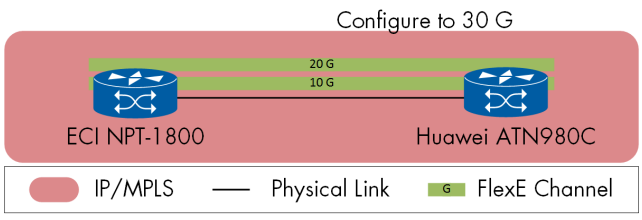

Figure 47: FlexE Dynamic Bandwidth Adjustment

The following devices successfully participated in the test:

- PE: ECI NPT-1800, Huawei ATN980C

Clocking

As the growing importance of the time synchronization in the industry, and the crucial role it plays in modern technologies like 5G; clocking tests become very essential, not only regarding time accuracy and error but also failover scenarios and security.

The industry is going towards 5G and its standards, the synchronization requirements are becoming tighter and harder to achieve, which points out the need for very accurate, modern, wise and reliable testing and test scenarios. Here comes the role of the EANTC MPLS event, providing highly reliable, accurate, and independent testing.

This year's event presented new test plans covered: 5G synchronization requirements, GNSS security, ITU-T performance requirements e.g. Boundary Clocks class C/D, in addition to tests have done before covering resiliency scenarios and PTP implementations; We tested the behavior of the time signal delivery in optimal and sub-optimal conditions: hold-over performances, source failover between two grandmaster clocks with high precision clocking and we reached 27 successful combinations.

For most of the tests we defined the accuracy level of ± 260 ns (ITU-T recommendation G.8271 accuracy level 6A), in other cases we defined the accuracy level of ±1.5 µs (ITU-T recommendation G.8271 accuracy level 4) as our end-application goal, with 0.4 µs as the phase budget for the air interface. Therefore, the requirement on the network limit, the last step before the end-application, had to be ±1.1 µs.

EANTC used the Calnex Paragon suite of products for both measurement and impairment scenarios. Paragon-X was used to generate the network impairment characteristics (G.8261 Test case 12) and in providing accurate measurement, Paragon-T used to provide measurements, and was very valuable having 4 possible ports to take the measurements at the same time which helped us performing multiple tests in parallel. The Calnex Analysis Tool was our analysis and reports generation tool, it provided all what we needed to apply masks or calculating the Time Error with all its forms (Maximum Absolute Time Error, Constant Time Error,...) and also reporting against the 5G network limits and clock mask performance.

The primary reference time clock (PRTC) was GPS using an L1 antenna located on the roof of our lab. The synchronization test team tested brand new software versions, products, and interface types, including PTP over 100 GbE.

Our tests helped to discover several small issues, but the R&D departments of the vendors reacted quickly providing patches and troubleshooting support.

Phase/Time Partial Timing Support

As a respond to the need for phase/time synchronization solution that is more feasible in non-greenfield deployments, PTP telecom profile for phase/time synchronization with partial timing support from the network has been developed.

We performed the test with the ITU-T G.8275.2 profile (PTP telecom profile for Phase/Time-of-day synchronization with partial timing support from the network) between the Grandmaster and the Boundary Clock, without any physical frequency reference – such as SyncE; the grandmaster clock was provided with GPS input, while the slave and boundary clock started from a free-running condition.

This setup emulates Partial Timing Support scenario, using the impairment between the GM and the BC; the impairment was applied as per the PDV profile according to G.8261 test case 12 using Calnex Paragon-X device.

For this test case, we have 3 vendors Microchip, Juniper, and Huawei. The devices swapped the roles to get as many pairs as possible. Calnex Paragon-X was used to generate the impairment between the GM and BC, while Paragon-T was used to take the measurements.

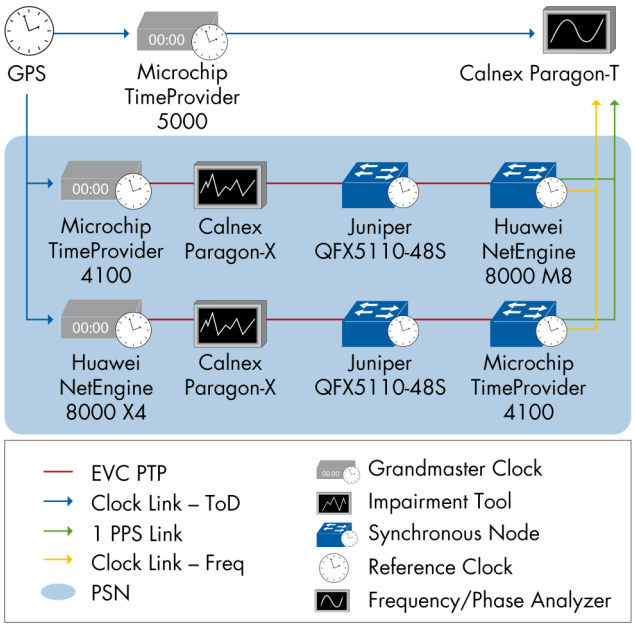

Figure 48: Phase/Time Partial Timing Support

The following devices successfully participated in the test:

- Boundary Clock: Juniper QFX5110-48S
- Grandmaster: Huawei NetEngine 8000 X4, Microchip TimeProvider 4100
- Impairment: Calnex Paragon-X
- Phase and Frequency Analyzer: Calnex Paragon-T
- Slave Clock: Huawei NetEngine 8000 M8, Microchip TimeProvider 4100

We had to wait the GM to lock on the BC - for one combination - before the applying the impairment, because they were not able to lock and continued the remaining steps successfully, some of the devices had the limitation of having a GM on 10Gbps link and a slave on 1Gbps which caused PTP locking issues, then we had to use same speed for both.

Phase/Time Synchronization, Source Failover

Deploying a reliable time delivery method is a concept of IEEE 1588-v2, and to achieve this, the clock source redundancy is required with primary and secondary grandmasters provided to a clocking system. For PTP the whole chain includes a boundary clock and its slaves next to the clock source. The boundary clock determines the primary grandmaster with the best clock quality and interacts with the slaves via PTP to deliver precise time over the network.

In this setup, we tested a real-life resiliency with two Grandmasters, Boundary Clock, and a Slave clock. The Boundary clock was locked on the primary GM, and then we degraded the GM A quality by unplugging the GNSS antenna. We verified that the boundary clock switched over to the secondary grandmaster and measured the slave clock's transient response.

The test was performed using the ITU-T G.8275.1 between the Grandmasters, Boundary clock and Slave clock, while the SyncE was enabled through the whole chain.

It's critical to calibrate the GMs and to compensate the cable delays between the GMs and the GNSS antenna to guarantee that these delays will not affect the test results. The following combinations achieved the G.8271 level 6 accuracy.

All the passed combinations, have reached the accuracy level 6 (130 ns - 260 ns), with a total of four test results.

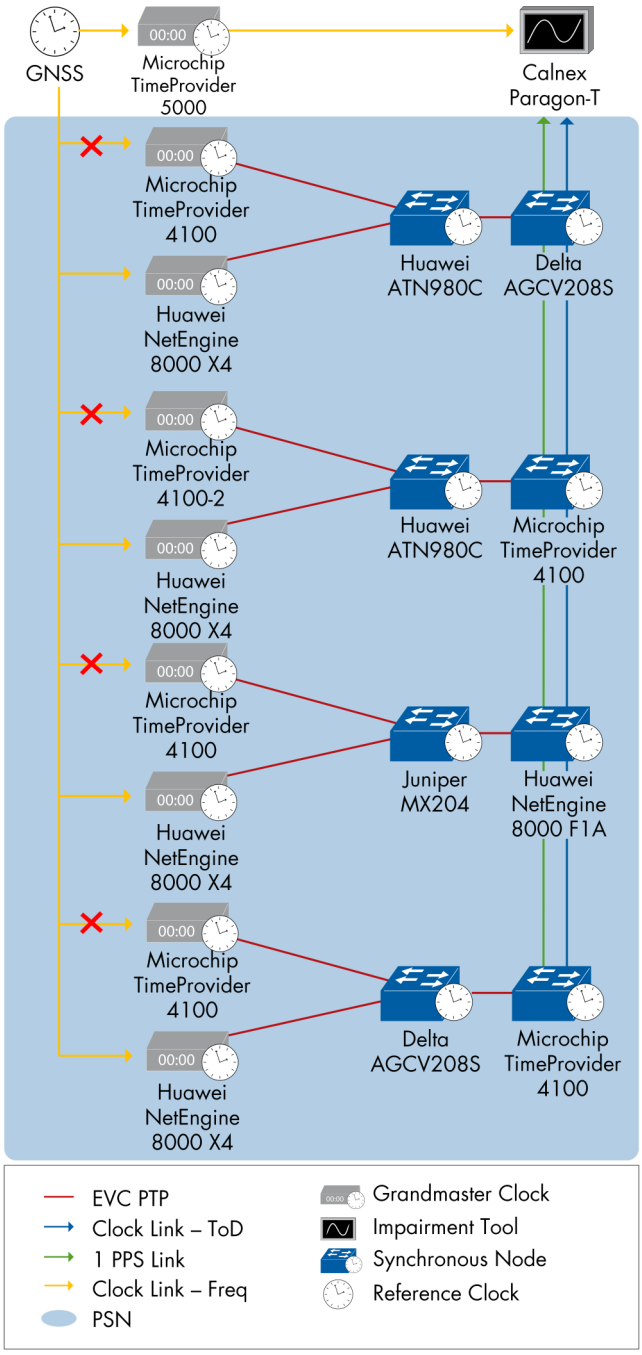

Figure 49: Phase/Time Synchronization, Source Failover, 1-4 pair - Level 6 Accuracy

The following devices successfully participated in the test:

- Boundary Clock: Delta AGCV208S, ECI NPT-1800, Huawei ATN980C, Juniper MX204
- Grandmaster: Huawei NetEngine 8000 X4, Microchip TimeProvider 4100
- Phase and Frequency Analyzer: Calnex Paragon-T
- Reference Clock: Microchip TimeProvider 5000
- Slave Clock: Delta AGCV208S, ECI NPT-1800, Huawei NetEngine 8000 F1A, Microchip TimeProvider 4100

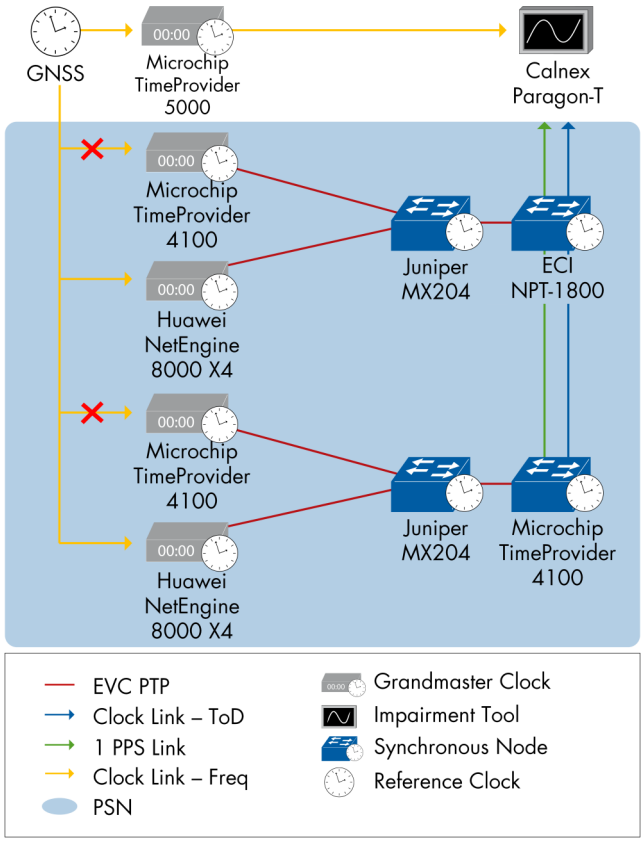

Figure 50: Phase/Time Synchronization, Source Failover, 5-6 Pair - Level 6 Accuracy

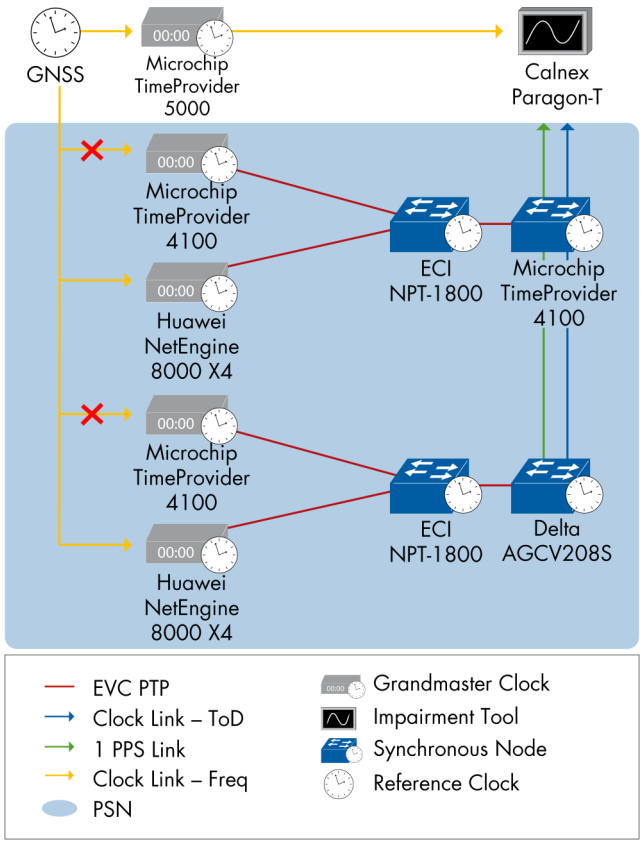

Figure 51: Phase/Time Synchronization, Source Failover, 7-8 Pair - Level 4 Accuracy

Finding the right match in all pairs could be a tough technical task in an end-to-end network service. To find the inaccuracies from failed tests, we summarized the following issues.

We observed in one case a huge drift and time errors (20k ns). We took apart the system to a minimum with the boundary clock and found the gap. It was generated by the boundary clock when both GMs lose the GNSS signal which leaves the BC in Holdover state. Although the accuracy maintain during the Holdover state without any time and frequency reliable source would be unlikely to happen, but the idea of this step is checking an extreme case.

One slave clock did not send any PTP packets, but it was unclear if the boundary clock sent any unexpected sync message. Therefore, the test could not be locked.

High-Precision Clocking Source Failover

Slave and boundary clocks shall be provided with a primary and secondary grandmaster for resiliency. In this test, the boundary clock was mounted in a spotlight to show the best time measurement in a lightweight way. We particularly did not add other slave clocks for overloads, in order to enhance accuracy high precision reached.

Both grandmasters were provided with a GPS signal. We allowed the boundary clock to lock to the primary grandmaster and then degraded the primary grandmaster's quality by disconnecting its GPS input and measure the slaved clock's transient response. We also verified the correct clock class is signaled by the grandmasters.

The test was performed using the ITU-T G.8275.1 between the Grandmasters, and the Boundary clock while the SyncE was enabled through the whole chain. All the depicted combinations passed the G.8271 level 6 accuracy (Maximum Absolute Time Error < 260 ns) - some of them achieved G.8271 level 6C accuracy.

We reached up to 65 ns precision accuracy.

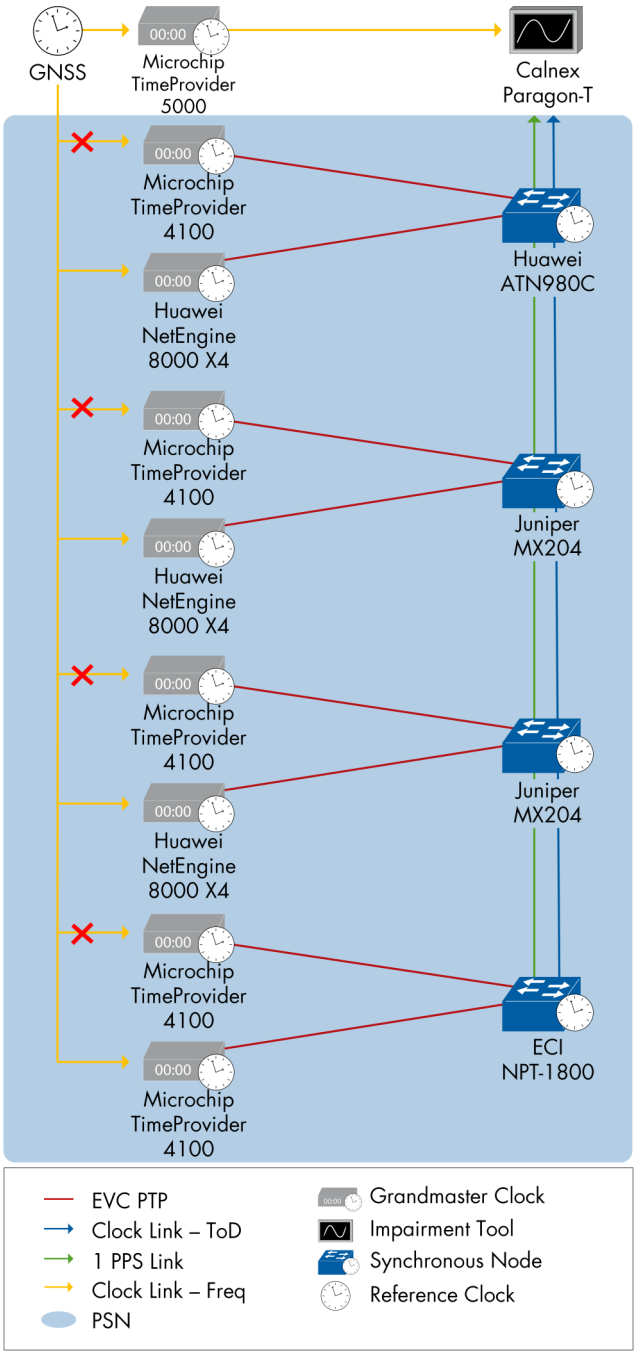

Figure 52: High-Precision Clocking Source Failover

The following devices successfully participated in the test:

- Boundary Clock: ECI NPT-1800, Huawei ATN980C, Huawei NetEngine 8000 F1A, Juniper MX204
- Frequency and Phase Analyzer: Calnex Paragon-T
- Grandmaster: Huawei NetEngine 8000 X4, Microchip TimeProvider 4100
- Reference Clock: Microchip TimeProvider 5000

Phase/Time Partial Timing Support over MACsec

The goal of the test was to verify that a slave clock can maintain the required synchronization quality when it was using the G.8275.1 profile and while using MACsec between Boundary and Slave clocks.

Packet network was originally not meant to be sensitive, so the clock accuracy was aforethought, even less credence in the security for clock synchronization. It usually works well without causing noticeable problems. However, 5G latest has the demand for applications and security turns to be considered. To avoid the DoS attack on PTP since messages were readable in the clear text for everyone between the boundary clock and the slave clock.

MACsec is a recommendation based on the requirements of the threat model defined in RFC 7384 "Security Requirements of Time Protocols in Packet Switched Networks", provided guidance to the standards a community for PTP like protocols in developing frameworks and guidance for securing network-based protocols.

We enabled the MACsec between the BC and SC, we measured the time error on the output of the Slave Clock, then we performed a resiliency scenario, disconnected the GM A from the GNSS, we observed the failover to the GM B and measured the time error.

During the test, we used the Calnex Paragon-X to capture the packets between the BC and SC to make sure they are encrypted and the MACsec is enabled.

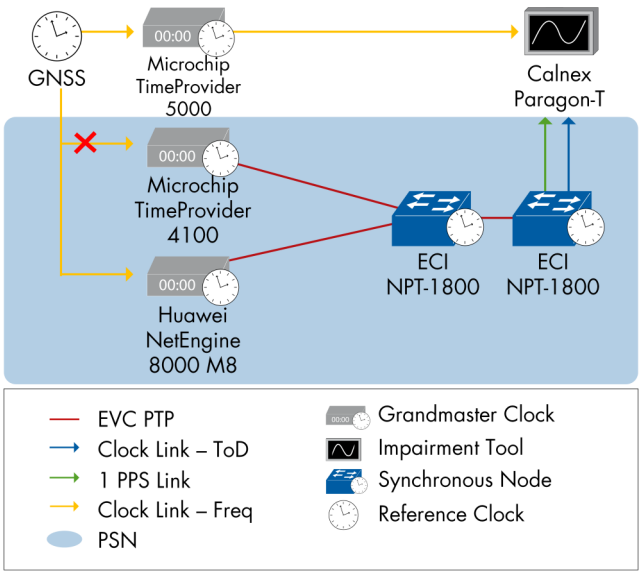

Figure 53: Phase/Time Partial Timing Support over MACsec

The combination passed the G.8271 accuracy level 6.

It is recommended for External Transport Security Mechanisms. The external transport mechanisms include Internet Protocol Security (IPsec) and IEEE 802.1AE MAC Security Standard (MACsec), which provide integrity protection and authentication at the transport (Ethernet) layer.

- Boundary Clock: ECI NPT-1800
- Frequency and Phase Analyzer: Calnex Paragon-T
- Grandmaster: Huawei NetEngine 8000 M8, Microchip TimeProvider 4100
- Reference Clock: Microchip TimeProvider 5000

The following devices successfully demonstrated the MACsec scenario:

- Slave Clock: ECI NPT-1800

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

The goal of the test is to verify that a slave clock can maintain the required synchronization quality when it is using the G.8275.1 profile, with a chain of Class C boundary clocks according to G.8273.2 profile.

The revision of G.8273.2 recommends the performance standards for boundary clocks in the network. This revision adds two new high-accuracy clocks, Classes C and D to the original Classes A and B. The purpose of the new clocks is to be used in the new mobile networks, particularly in the context of 5G, which has some very strict timing requirements.

We built multiple combination of Class C boundary clocks, and added failover scenario to create reliable, and real case scenario.

All the Boundary clocks were locked on GM A, we measured the time accuracy for 1000 seconds, then we degraded the GM A clock class by disconnecting the GNSS link, then we measured accuracy during the failover to GM B.

We tested these testcase against the ITU-T G.8273.2 Maximum Absolute Time Error requirements less than 30 ns for Class C Boundary Clocks, and assuming that the whole chain will achieve the requirement.

All depicted combinations passed the test, and achieved the requirement value of time error.

We faced some interop issues regarding different implementation of Correction Field value in PTP messages. Some vendor sent huge CF values which caused the connected nodes to discard the value and have a huge time error.

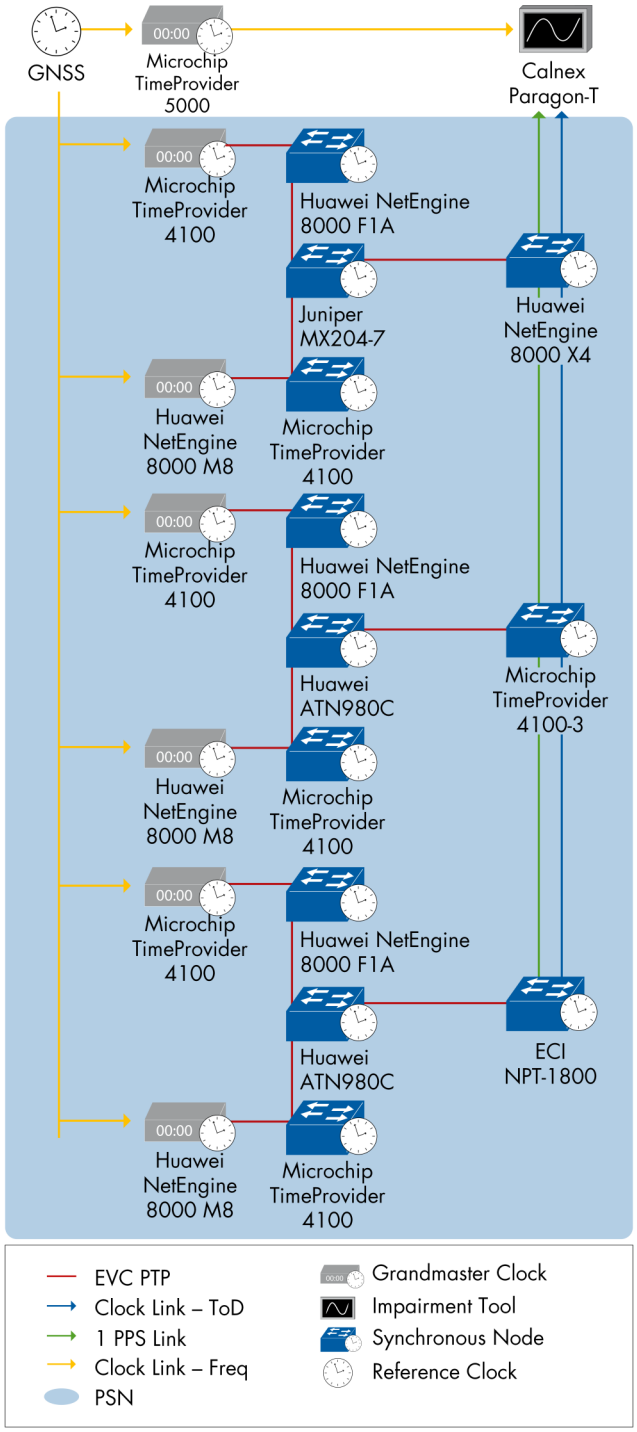

Figure 54: Phase/Time Full Timing Support, Boundary Clocks Class-C Test 1-3 Pairs

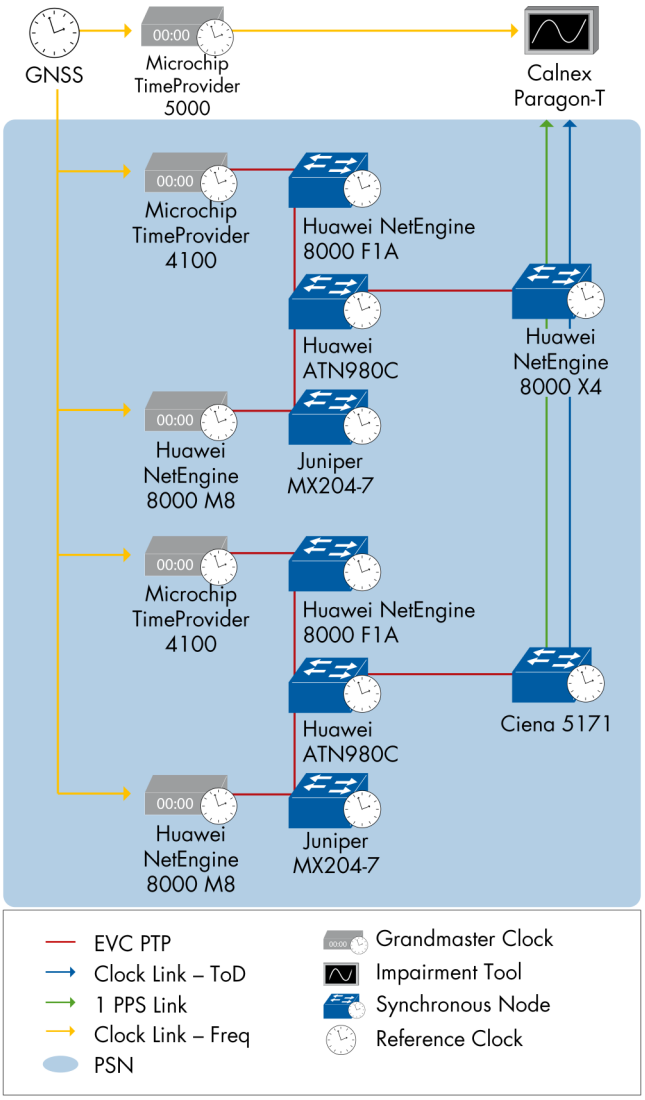

Figure 55: Phase/TIme Full Timing Support, Boundary Clocks Class-C Test 4-5 Pairs

The following devices successfully participated in the test:

- Boundary Clock: Huawei ATN980C, Huawei NetEngine 8000 F1A, Juniper MX204, Microchip TimeProvider 4100
- Frequency and Phase Analyzer: Calnex Paragon-T
- Grandmaster: Huawei NetEngine 8000 M8, Microchip TimeProvider 4100
- Slave Clock: ECI NPT-1800, Huawei NetEngine 8000 X4, Microchip TimeProvider 4100

Phase/Time GNSS Vulnerability Test

Providing the security demand for 5G, GNSS is a reliable system. Professional GNSS receivers require to be aware of all possible vulnerabilities which could be exploited. Spoofing is an intelligent form of interference that makes the receiver believe it is at a false location.

To protect these receivers against spoofing, a GNSS security device can be used, which analyzes the GPS signal. GPS signal data is received and evaluated from each satellite to ensure compliance along with analyzing received signal characteristics.

This test verifies the GNSS Firewall capabilities of detecting spoofed GNSS signal and stop it.

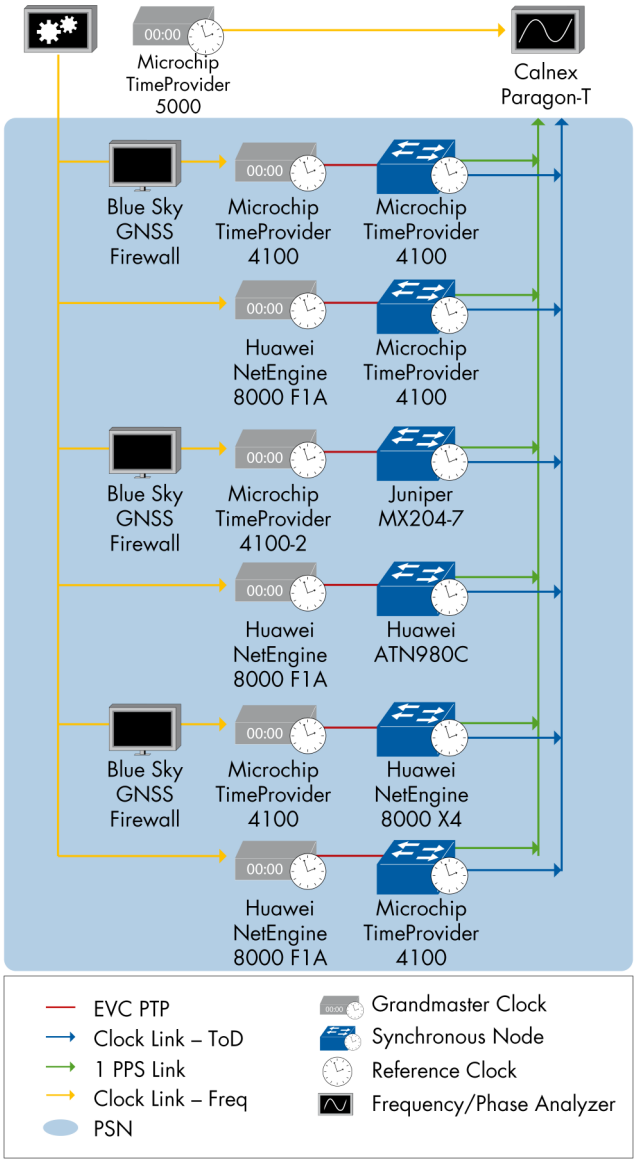

Figure 56: Phase/Time GNSS Vulnerability Test

In this test, we tested the BlueSky firewall from Microchip, some ports can be configured as Verified output ports, which able to detect the spoofed signal and Hardened ports which do not stop the signal.

We connected the BlueSky to the GNSS antenna, and it's outputs to the GMs, we emulated the spoofing by disconnecting the GNSS signal which caused the BlueSky running on its internal oscillator, and measured the time error on the slave clocks.

On the Hardened ports, the slave clocks didn't detect any lose of the GNSS signal and didn't show any transient response, and showed the clock class 6 as the source of the time. Unlike the Verified ports which detected the GNSS spoofing and stopped forwarding the signal, which caused the Slave Clocks to rise an alarm for Master Clock Class degradation, which means the test was successfully passed.

The following devices successfully participated in the test:

- Frequency and Phase Analyzer: Calnex Paragon-T
- GNSS security device: Microchip BlueSky
- Grandmaster: Huawei NetEngine 8000 F1A, Microchip TimeProvider 4100
- Slave Clock: Huawei ATN980C, Huawei NetEngine 8000 X4, Juniper MX204, Microchip TimeProvider 4100

Phase/Time Passive Port Monitoring

This test case aims to verify the Relative Time Error monitoring option as per G.8275.1 Annex G and Verify alternateMasterFlag usage and potential accuracy of observed TE difference.

It is always necessary to check the features of the standards, and that the devices behavior follows it. This test is a two features verification:

- Passive port monitoring
- Delay Asymmetry detection

We connected the setup which is depicted in the diagram, and checked the ports states for the DUT, the port status was passive.

We used the Calnex Paragon-X to capture the PTP packets and checked the alternateMasterFlag through the Calnex PTP Field Verifier, and the value was as per the standard: False.

Then we used the Calnex device to apply 250 µs constant delay in one direction, the DUT detected the delay and showed an alarm as expected.

The test was successfully passed.

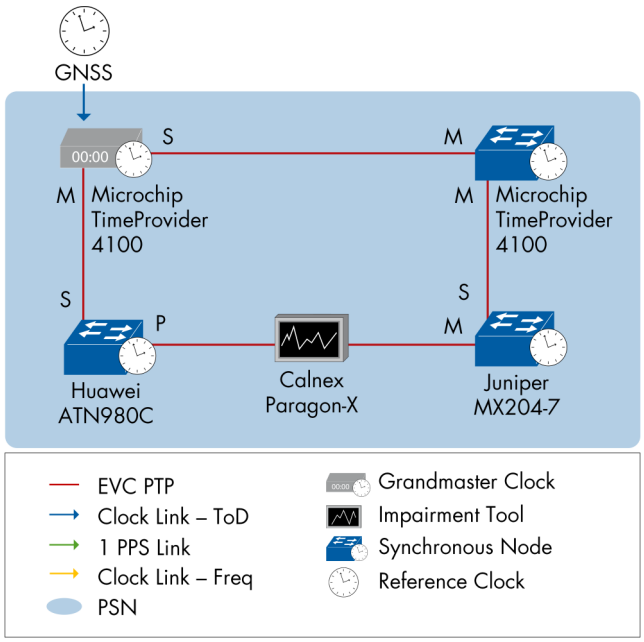

Figure 57: Phase/Time Passive Port Monitoring

The following devices successfully participated in the test:

- Boundary Clock: Huawei ATN980C, Juniper MX204, Microchip TimeProvider 4100
- Frequency and Phase Analyzer: Calnex Paragon-T
- Grandmaster: Microchip TimeProvider 4100

Remote Collaboration Aspects

At EANTC, we have conducted many remote testing programs in the Network Functions Virtualization (NFV) space, for example, the New IP Agency (NIA) test series with more than 80 participants of which only 20 % were present locally, the Intel and Lenovo performance testing program for a substantial number of virtual network functions, or the NetSecOpen next-gen firewall performance certification program.

That said, it was the first time for us to conduct a transport network-focused interoperability test event in a hybrid local and remote fashion. Thanks to the great commitment of all participants whether in Berlin or at many other places worldwide, the collaboration went very well and was quite efficient. All equipment was connected locally, typically installed and cabled by regional vendor representatives or by the EANTC team ahead of the hot-staging. Vendors then accessed the management remotely for configuration, collaborating via permanent Zoom videoconferences set up separately for each test area and by EANTC-hosted Rocketchat messaging. Specifically, the availability of combined video and chats, plus fully transparent remote access to the equipment under test kept remote participants engaged from our point of view. We are really grateful that this ad-hoc experiment forced on us by Covid-2019 has worked out, and will expand it in future test events.

Some additional aspects of workflow pipelines that had been planned way in advance helped to increase the efficiency of testing as well. The EANTC project management scheduled test combinations via our Confluence-based custom-made planning and documentation system. Once a specific test combination had been completed, each involved vendor uploaded the detailed results and confirmed the verdict using a TAN-based authorization scheme. All results were subsequently validated and included in the whitepaper by EANTC employing a standardized workflow. This way, the vendors created more than 500 results in seven testing days, and the EANTC team verified, approved, and distilled the documentation within two weeks after the event.

Summary

We successfully performed a wide range of EVPN tests, among which IRB, all-active multi-homing, MAC Mobility, and loop prevention extensively addressed the EVPN control plane for MAC/route learning intelligence. Carrier Ethernet services tests included EVPNs with different service types and protocols that save resources and time between multicast nodes. Finally, EVPN scenarios for data center interconnection were evaluated. Throughout the event, we observed EVPN-capable devices supporting the entire test scope as well as solutions specializing in some advanced test cases. We observed only a few interoperability issues, helping vendors to further mature their EVPN implementations.

FlexE tests included two specific features: channelization and bandwidth adjustment. They were both successfully tested in this initial round of multi-vendor evaluation. We plan to expand on this area with more test cases and more participating implementations next year.

In the Path Computation Element Protocol (PCEP) test area, we successfully tested the PCE- and PCC-initiated stateful path computation with different vendor combinations. We successfully confirmed traffic rate limitation and traffic drop using BGP Flowspec for IPv6 with various rules. In the Egress Peer Engineering with Segment Routing, the participating vendor implementations successfully redirected the traffic via ingress Autonomous System Boundary Routers (i_ASBR) by modifying the BGP SR color policy in the ingress PE router.

There was quite some progress in the NETCONF/Yang tests this year. Device configuration worked well with multiple NETCONF servers and clients. Furthermore, L2VPN/ L3VPN services were successfully provisioned in multi-vendor, multi-domain environments. This year, multi-vendor EVPN service creation was achieved as well. Generally, EANTC would welcome broader and more detailed interop testing of Yang models in the future.

EANTC witnessed a substantial number of new, successful test combinations in the Segment Routing area. In addition to the frequent testing of SR-MPLS using ISIS, this year we verified the interoperability of nine network vendors to set up an SR-MPLS topology using the OSPFv2 routing protocol for the first time. We tested service protection over the SR-MPLS data plane using TI-LFA. We examined different TI-LFA protection flavors against link or SRLG failure.

Another test case included the seamless BFD interoperability over SR-TE path. We successfully examined the SRv6 data plane for EVPN services, traffic engineering, and TI-LFA. We reintroduced the BGP Segment Routing this year with a spine layer consisted of two nodes and three different vendors took the role of the leaf node.

The Clock Synchronization tests started with classic Partial and Full Timing support implementation, went on with failover resiliency tests and concluded with new tests of Class C Boundary clocks, GNSS security, and PTP over MACsec. This year, the test coverage caught up with recently ratified standards, and precision elevated to meet the demands of new 5G deployment scenarios. We were able to achieve G.8271 level 6 accuracy in most of the tests - and even for level 6C, we tested the passive port feature and asymmetry detection.

Security aspects are not very popular in PTP testing, but they are becoming ever more important with the larger scale field deployments to less secure locations in 5G networks. We tested a GNSS firewall, and encrypting PTP packets with MACsec successfully, which are very difficult tests to pass.

Throughout all test areas, we faced a number of interoperability issues and problems - as expected in this type of events. It was a great experience to see very thoughtful, experienced, and knowledgeable engineers working together, troubleshooting and solving problems towards production-ready interoperability of advanced technology solutions.

EANTC AG
European Advanced Networking Test Center

Salzufer 14
10587 Berlin, Germany

Tel: +49 30 3180595-0
Fax: +49 30 3180595-10
info@eantc.de
http://www.eantc.com

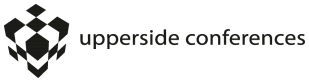

Upperside Conferences

54 rue du Faubourg Saint Antoine
75012 Paris - France

Tel: +33 1 53 46 63 80
Fax: + 33 1 53 46 63 85
info@upperside.fr
http://www.upperside.fr

20200401 v1.2