EANTC

WHITE PAPER 2019

Multi-Vendor Interoperability Test

MPLS+SDN+NFVWORLD
PARIS 2019

Editor's Note

Carsten Rossenhövel
Managing Director, EANTC

EANTC's annual multi-vendor showcase at the MPLS, SDN and NFV World Congress in Paris verifies, documents and helps to improve the interoperability of commercial solutions for communication service provider (CSP) transport networks. This has always been the case – since we started the series of test events in the mid-2000s. In the past few years, we have focused Software-Defined Networks (SDN) a lot. Over time, we have noticed that the telecoms industry follows some kind of "pork cycle": From the outset, standards compliance is rated very important in requests for quotation (RfPs) issued by telecom operators. Two to three years later, CSPs typically start believing that standardized solutions hamper innovation and competitive differentiation – at which time they start to lean towards individual proprietary solutions. After an additional few years, though, CSPs notice the vendor lock-in and the increasing cost of maintenance for such proprietary solutions. Around the same time, typically many vendors understand a technology better and aim to align their solutions; at this time, inferior protocol solutions disappear. Towards the end of a cycle, standards are getting important again which enables broader interoperability than at any other stage of the cycle.

I believe we are nearing the end of such a cycle with regards to SDN. In this year's event, we noticed an increasing observance of standards compliance across all vendors. There was no request for single-vendor demos (of which we have seen many in the past). The number of implementations participating in each test case increased across the board, creating a larger choice of more mature, interoperable implementations for CSPs. Admittedly, the SDN standards are complex and contain many protocol options which is a burden for implementations; however, this is quite normal for established technologies. Participating implementations each supported more of these options this year (as the industry moves from differentiation to integration); consequently there was more overlap and better interoperability.

As a result, this was the most successful EANTC interoperability event so far, judging by the numbers. We had a total of 174 successful test combinations involving 68 device types from 20 vendors. All major U.S., European, and APAC manufacturers were present. Without exception, all participating vendors deployed excellent senior engineers to the two-week hot-stage testing – 75 attendees rated among the highest ever. At EANTC, we are very thankful for this support. As a result, we are able to report a lot of details in this white paper – details that can help CSPs to identify suitable solutions for next-generation SDN deployment.

The industry definitely converges on EVPN and Segment Routing technologies, aiming to straighten network design and to reduce the variety of network protocols required to run a wide-area network. In particular, for EVPN, the industry has made tremendous progress in the last four years. Back then, the technology started as a data center solution: Very few vendors were able to interface on a basic feature set. Today, and based on the testing at EANTC, quite a number of vendors can successfully interoperate using a common and well understood EVPN feature set, not only for VXLAN overlays but also for Segment Routing networks.

Likewise, we are seeing convergence in the NETCONF/YANG area. Industry efforts in the past five years are paying off now – the test execution was smooth and successful in this year's EANTC test, with more common YANG models observed. Interestingly, the mix of participants in this area changed: This time only two incumbents submitted NETCONF/YANG implementations; the area was dominated by smaller vendors and newcomers. The reasons are unclear; we encourage readers to ask their large suppliers about continued commitment to standardized network configuration. We will keep monitoring the space and will aim to encourage more vendors to join this topic next year again. There is still a lot to be done – specifically when it comes to common YANG models for network service-level configuration, aligning with standardization efforts at the IETF.

5G services need to be taken into account in all network aspects these days. We evaluated interoperable slicing concepts, for example. One of the most challenging and critical aspects for 5G support in the backhaul network is clock synchronization, of course. Sync requirements are changing in subtle ways with 5G, and surprisingly the support for higher speed 100 Gbit/s connections creates a specific challenge as well. In this year's event, we developed new test cases to be

Table of Contents

pursuant with new technologies and new requirements. Thanks to an excellent, highly experienced team of recurring vendor participants, the test results were a success across the board and fulfilled all of our expectations.

The whole EANTC team is very thankful of the great spirit of all participating vendors, organizationally and individually. It is an honor to provide the melting pot where the industry can collaborate towards interoperable, common use cases, despite the fact that each participant has their own affiliations and interests. I hope that the white paper conveys this spirit and is an interesting read, explaining the wealth of results in a truly open, fair, and impartial manner.

Technical Summary

The goals of EANTC's multi-vendor test event series are to:

a) Improve the interoperability of different implementations of the same standards, helping to eliminate software development issues;

b) Be a proving ground for Internet drafts, validating how a new standard can be implemented;

c) Provide a platform for fruitful discussions regarding standard and Internet draft interpretations. Normative text is sometimes not precise enough and can be understood in different ways, which would lead to non-interoperable implementations.

Network equipment manufacturers typically deploy senior developers or network architects for the intense two-week hot-staging. A total of 75 engineers from 20 participating companies collaborated at EANTC's lab in March 2019.

This year's interoperability event focused on the full range of next-generation wide-area transport solutions, network management and clock synchronization again:

- Ethernet VPN (EVPN)
- Segment Routing (SPRING)
- SDN Controllers using Path Computation Element Protocol (PCEP)
- SDN Controllers using BGP with Segment Routing programmed SR-TE tunnels
- Precision Time Protocol (PTP) interoperability and performance (ready for 5G)
- Common Yang models
- Integration of Microwave equipment with SDN

In the following, we summarize the main test aspects and some results highlights.

Software-Defined Networks

Similarly to previous events, vendors corrected their implementations on the fly to improve standards compliance, increasing the test case success rate. It is always impressive to see new development builds available for testing in the next morning. We have seen very good progress in the SDN controller interop testing. However, there is still more work to be done to enable controller-based, multi-vendor network designs.

The Ethernet VPN/SPRING area, on the other hand, showed much greater commonalities across vendors. The interop issues found did not affect the main aspects of traffic delivery but were rather related to auxiliary functions such as OAM (specifically MPLS segment routing traceroute) or traffic optimization functions.

SRv6 is an interesting topic. Two major participating network equipment manufacturers strongly believe in the technology; SRv6 eliminates further protocols, thus enabling more elegant network configuration. The tests between these two manufacturers and two test tool vendors were very successful. The interoperability achieved during our test sessions was groundbreaking in any case; time will tell its relevance. SRv6 requires hardware support on each node handling the new SR Header (SRH) field. Some implementations validated this year used merchant-silicon forwarding ASICs.

Network Management and Automation

The NETCONF/YANG standards are all about enabling network automation. The cost and reliability of network automation solutions have long been the Achilles heel for many business cases in enterprises and service providers alike. By providing a solid, high-level interface for management applications to leverage, the cost is slashed while the agility and reliability are dramatically improved.

Systems participate in NETCONF/YANG tests in different roles. A client application acts as a NETCONF/YANG manager towards a single system. An orchestrator provides a service level NETCONF/YANG management interface and manages several controllers and/or devices using network-wide transactions. A controller provides a service level NETCONF/YANG management interface and manages several devices using network-wide transactions (if the devices support this). Finally, a device provides a NETCONF/YANG management interface for the device functionality.

The NETCONF/YANG tests cases are designed as a ladder with three levels. To succeed in the basic level, basic interoperability, a controller must be able to read and write at least something on a third-party device, so that traffic can be enabled and disabled.

To succeed in the intermediate service level, an L2- or L3VPN service must be correctly provisioned and deprovisioned, using network-wide transactions without leaving anything behind. Finally, to master the multi-level service, a service that crosses multiple technology domains must be provisioned and deprovisioned.

In this year's event, we saw progress in vendors migrating from RESTCONF/API to NETCONF with regards to the communication between the Controller and the Orchestrator. Some vendors have not completed this migration yet. A subset of participating controllers already supported Northbound communication between the Controller and the Orchestrator. Finally, we noticed a larger group of vendors supporting device management and configuration of L3VPN, compared with the support of L2VPN service configuration over NETCONF.

Clock Synchronization

This year, 5G was a major focus area for the synchronization test cases. In the high-precision clocking test, participants successfully achieved ITU-T G.8271 Level 6 in two test scenarios with multiple combinations. Additionally, as in previous years a lot of testing was done with new combinations or updated software to ensure interoperability of PTP profiles.

We successfully demonstrated Precision Time Protocol (PTP) support over 100GbE links in two combinations; increasingly, 100GbE is used in the service provider core and aggregation networks. It was deemed critical to test the implications on the time, phase and frequency synchronization.

The security aspect of the time synchronization is still limited, as the next version of the PTP is not published yet. Since there is no standardization for PTP over MACsec or IPsec, we postponed these test cases.

Participants and Devices

Participants	Devices
ADVA Optical Networking	ADVA GO102Pro ADVA OSA 5430 ADVA XG480
Arista Networks	Arista 7050SX3 Arista 7280SR
BISDN GmbH	BISDN Basebox
Calnex Solutions	Calnex Paragon-t Calnex Paragon-X Calnex SNE
Cisco	Cisco ASR 9000 Cisco IOS XRv9000 Cisco NCS540 Cisco NCS 5500 Cisco Network Services Orchestrator (NSO) Cisco Nexus 3100-V Cisco Nexus 3600-R Cisco Nexus 7700 Cisco Nexus 9300-FX Cisco Nexus 9300-FX2
Delta Electronics	Delta AG7648 Delta AGC7648A
ECI Telecom	ECI Neptune 1050 ECI Neptune 1300
Ericsson	Ericsson 6274 Ericsson 6471 Ericsson 6672 Ericsson 6675 Ericsson MINI-LINK 6352 Ericsson MINI-LINK 6363 Ericsson MINI-LINK 6366 Ericsson MINI-LINK 6654 Ericsson MINI-LINK 6691 Ericsson MINI-LINK 6693

Participants	Devices
Huawei Technologies	HUAWEI ATN910C-F HUAWEI ATN950C HUAWEI Network Cloud Engine (NCE) HUAWEI NE40E-F1A HUAWEI NE40E-M2K HUAWEI NE40E-X8A HUAWEI NE9000-8
Intracom Telecom	Intracom Telecom OmniBAS-2W IDU Intracom Telecom OmniBAS ODU
IP Infusion	IP Infusion OcNOS 1.3.5
Ixia, a Keysight business	Ixia IxNetwork Ixia XGS2 Chassis
Juniper Networks	Juniper NorthStar Controller Juniper MX104 Juniper MX204 Juniper MX480 Juniper QFX10002-72Q Juniper QFX5110-48S
Meinberg	Meinberg LANTIME M1000S Meinberg microSync HR
Microsemi, a Microchip company	Microsemi TimeProvider 4100 Microsemi TimeProvider 5000
Nokia	Nokia 7750 SR-7 Nokia Network Services Platform
Seiko Solutions	Seiko TS-2912-22
Spirent Communications	Spirent TestCenter
ZTE Corporation	ZTE Corporation ZENIC ONE ZTE ZXCTN 6180H ZTE ZXCTN 9000-18EA ZTE ZXCTN 9000-8EA

Table 1: Participants and Devices

Interoperability Test Results

As usual, 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

In addition to the EVPN signaling required to scalable data centers, the EVPN tests also focused on EVPN data center interconnect. We tested the various protocols as introduced by the IETF working group that carried scalability properties in mind. The EVPN control plane arose clear interoperability in multiple areas including Carrier Ethernet Services, EVPN enhancement, and EVPN routing and switching. We also successfully tested EVPN maintenance using Y.1731.

Carrier Ethernet Services

EVPN family includes various solutions that address Ethernet point-to-point (E-LINE) and Ethernet rooted-multipoint (E-TREE) service types. This white paper focuses on both solutions.

E-Line Service

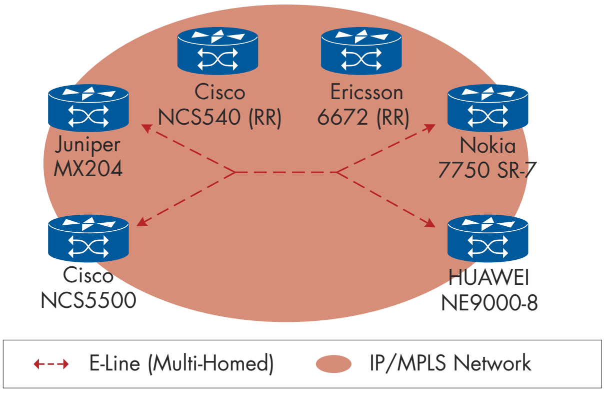

Figure 1: E-Line Service (M-H)

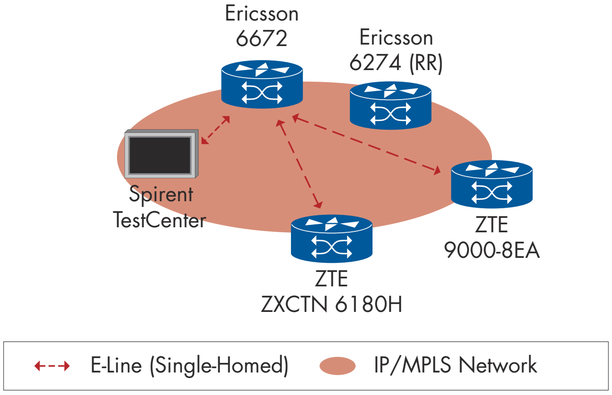

Figure 2: E-Line Service (S-H)

The EVPN family enables all types of Ethernet services under a common architecture. These solutions are currently under standardization by the IETF L2VPN Working Group. This test focused on the E-Line (point-to-point) service implemented in a single-homed as well as a multi-homed scenario within an Autonomous System (AS). In both scenarios, we sent IPv4 traffic to the E-Line service 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 E-Line service. Excluding extreme values in seconds, the general out of service time was good (average 178 ms).

The following devices successfully participated in both single-active and All-Active multi-homed scenario:

- PE: Cisco NCS 5500, HUAWEI NE9000-8, Juniper MX204 and Nokia 7750 SR-7 using Cisco NCS540 (Route Reflector) and Ericsson 6672 (RR), together with HUAWEI NE40E-M2K (CE).

Observed Out of Service Time	Number of Measurements
5 – 9 s	2
17 – 56 ms	2
172 – 490 ms	3

Table 2: E-Line Service Out of Service Time

E-Tree Service Flexible

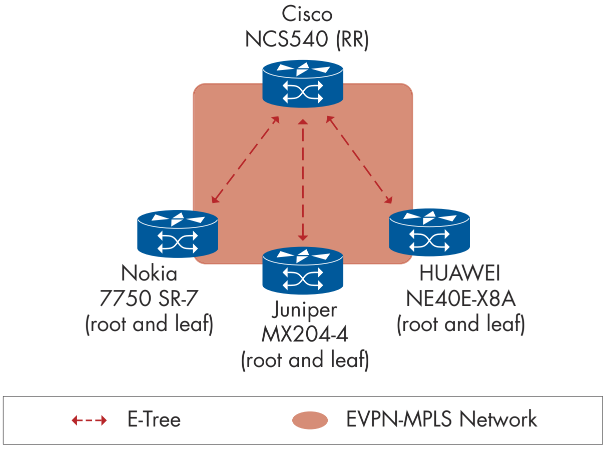
Cisco NCS540 (RR)
Nokia 7750 SR-7 (root and leaf)
Juniper MX204-4 (root and leaf)
HUAWEI NE40E-X8A (root and leaf)
E-Tree
EVPN-MPLS Network

Figure 3: E-Tree Service Flexible

This test focused on E-Tree (rooted-multipoint). In an E-Tree service, endpoints are labeled as either Root or Leaf site. Root sites can communicate with all other sites. Leaf sites can communicate with Root sites but not with other Leaf sites. We sent IPv4 unicast traffic to the established E-Tree service and did not observe any packet loss. We also verified the capability of network elements to support both P2MP tunnel and Ingress Replication tunnel at the same time using the composite tunnel type for broadcast (BUM traffic) and multicast traffic delivery.

The following devices successfully established E-Tree service:

- HUAWEI NE40E-X8A, Juniper MX204 and Nokia 7750 SR-7 using Cisco NCS540 (RR).

EVPN Enhancement

ARP suppression and MAC mobility tests were hallmarks of all EVPN participants. EVPN and extensions in a scalable data center are significant design goals. We were glad to achieve such a great number of interoperability results.

Proxy-ARP

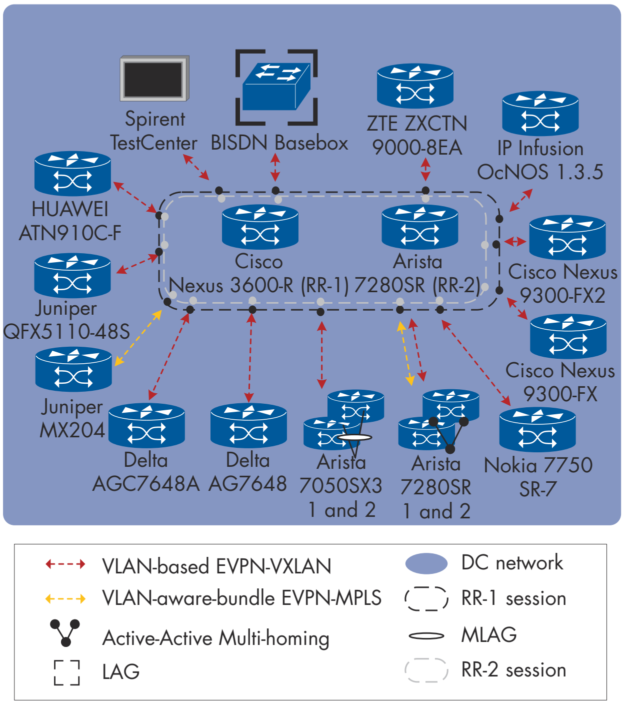

Figure 4: Proxy-ARP

We tested the ARP proxy functionality of EVPN, which eliminates ARP flooding within the transport network between PE nodes. We started this test by ensuring that a PE learns the local MAC/IP address from an emulated CE without any entries in the ARP table. In this process, we expected to observe the route exchange of the PE, since remote PEs learn its ARP entries with BGP updates. As described in the IETF draft, the PE with Proxy-ARP function advertises route type 2 (EVPN MAC/IP Advertisement route) carrying the MAC address along with the IP address in the MAC/IP Advertisement route. As expected, once we generated first ARP requests from the emulated CE, the PE learns the local MAC/IP address into the ARP table. We then observed that the information for these addresses was shown as EVPN route type 2 (RT-2) routes in the EVPN VRF indicating that ARP entries were successfully routed via RT-2 according to the IETF draft. Since the remote PE needed this information to populate the ARP proxy table, we also verified that the ARP proxy table was populated successfully on the remote PE, we sent ARP requests for remote IP addresses by the emulated CE and started capture packets. We expected that the PE intercepts the ARP request and performs a Proxy-ARP lookup for the requested IPs without any flooding to other CEs. As expected the lookup succeeded for the emulated queries, the PE did not flood any ARP Request in the EVPN network and the other local CEs.

- The following devices successfully established VLAN-based EVPN-VXLAN with ARP-proxy: Arista 7050SX3, Arista 7280SR, BISDN Basebox, Cisco Nexus 9300-FX, Cisco Nexus 9300-FX2, Delta AG7648, Delta AGC7648A, HUAWEI ATN910C-F, IP Infusion OcNOS 1.3.5, Juniper QFX5110-48S, Nokia 7750 SR-7, Spirent TestCenter and ZTE ZXCTN 9000-8EA using Arista 7280SR (RR) and Cisco Nexus 3600-R (RR).
- We also tested proxy-APR per VLAN-aware-bundle EVPN-MPLS with both Arista 7280SR and Juniper MX204.

One vendor did not establish the BGP session with all route reflectors since the BGP updates were rejected by one of the RR. The same vendor also failed to establish EVPN with another visualized solution.

MAC Mobility

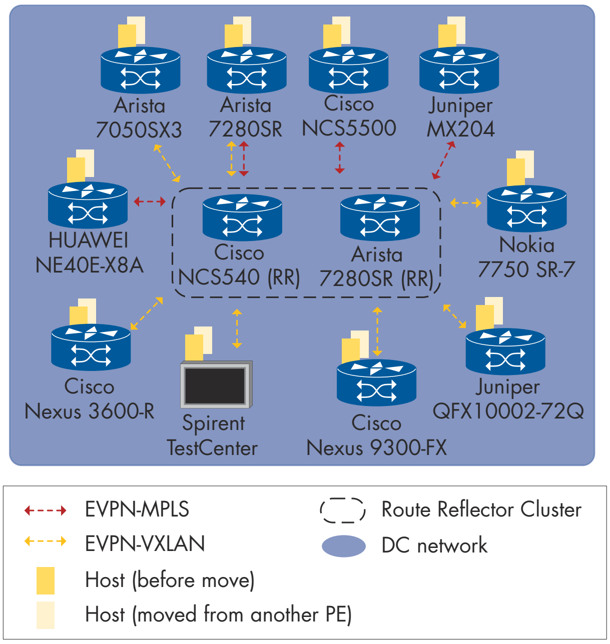

Figure 5: MAC Mobility

The challenge with the MAC mobility is how to do the Layer 2 stretch in a way that ensures that the entities can be reached after the relocation. We successfully tested the MAC mobility of EVPN which allows flexibility for a given host or end-station (as defined by its MAC address) to move from one Ethernet segment to another. EVPN introduces sequence numbering in its type 2 routes which prevent race conditions which might exist with multiple rapid moves.

In this test, we first selected an emulated host that has not been moved before from an Ethernet segment to confirm that within this initial state MAC/IP adver-

tisement 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 Ethernet segment by moving the traffic from previous PE to a new PE. Then the value increased to 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 withdraws its MAC/IP Advertisement route. We sent test traffic and did not observe any frame loss, we also did not receive any flooded traffic.
The following devices successfully participated in the MAC mobility test:

- using EVPN-VXLAN: Arista 7050SX3, Arista 7280SR, Cisco Nexus 3600-R, Cisco Nexus 9300-FX, Juniper QFX10002-72Q, Nokia 7750 SR-7 and Spirent TestCenter using Arista 7280SR (RR) and Cisco NCS540 (RR).
- using EVPN-MPLS: Arista 7280SR, Cisco NCS 5500, Juniper MX204 and HUAWEI NE40E-X8A.

EVPN Routing and Switching

The EVPN control plane with its scalable design goals, resided in the center of the test stage. The tested concepts include integrated routing and switching, IP subnet routing and All-Active multi-homing.

Flexible Cross-Connect Service

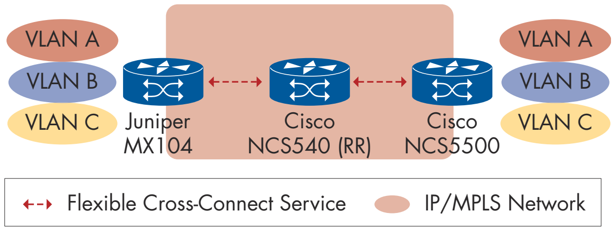

Figure 6: Flexible Cross-Connect Service

The VPWS flexible cross-connect (FCX), a tested draft among others first time in this event, was released by the IETF BGP-enabled services (bess) working group in April 2018. It defines a VID (VLAN IDs) lookup that allows multiplexing of different ACs (Attachment Circuits) to a single point-to-point service. The driving force of this protocol is to reduce the number of services in an operator network and maintenance through ACs' multiplexing, such as carrying multiple ACs through the same VPWS thus saving the number of EVPN service tags associated with it, EVPN-VPWS service tunnels and related OAM monitoring. These all reduce network resource consumption and the work load for operators.

We first determined the routes which carried the AC information in the EVPN routing table because FCX complements the traditional lookup of labels in EVPN and places VLANs into the Ethernet tag field. As all required information was shown in the routing table, we confirmed that the service was established, then we sent IPv4 unicast traffic to the service and did not observe any packet loss.

We successfully established the flexible cross-connect service between the following devices: Juniper MX104 and Cisco NCS 5500 using Cisco NCS540 (RR).

Integrated Routing and Switching (IRB)

EVPN with IRB solutions arose clear interoperability in this event and almost all EVPN vendors in this event participated in this test. Ethernet VPN Integrated Routing and Bridging (IRB) status is currently "work in progress" at the IETF and provides a solution for inter-subnet forwarding in data center environments with EVPN. MP-BGP EVPN enables communication between hosts in different VXLAN overlay networks by distributing Layer 3 reachability information in the form of either a host IP address route (route type-2) or an IP prefix (route type-5). We tested two different modes of IRB depending on the required lookup at the ingress or/and egress Network Virtualization Edge (NVE).

Integrated Routing and Switching (IRB) - Symmetric IRB

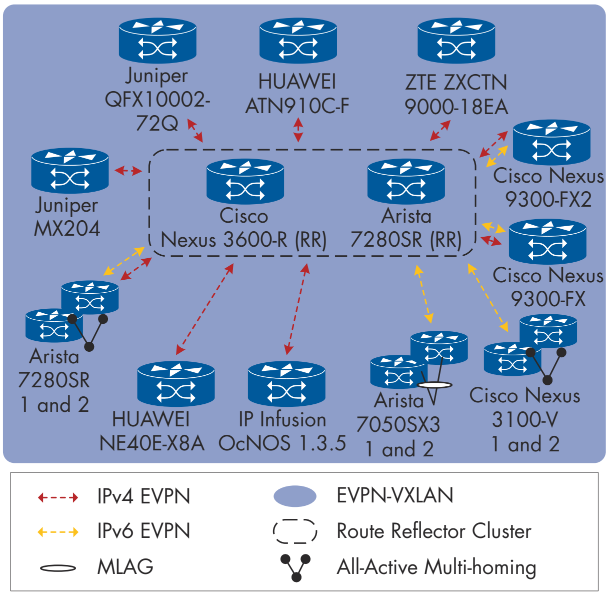

Figure 7: Symmetric IRB-VXLAN

In the symmetric IRB semantic, both IP and MAC lookup are required at both ingress and egress NVEs, their results per EVPN-MPLS or EVPN-VXLAN were:

- The following PE participated successfully in IPv4-based EVPN-VXLAN with IRB: Arista 7280SR, Cisco Nexus 9300-FX, Cisco Nexus 9300-FX-2, IP Infusion OcNOS 1.3.5, HUAWEI ATN910C-F, Juniper QFX10002-72Q, and ZTE ZXCTN 9000-18EA. The PEs per IPv6-based EVPN-VXLAN with IRB were: Arista 7050SX3, Arista 7280SR, Cisco Nexus 3100-V, Cisco Nexus 9300-FX and Cisco Nexus 9300-FX2. The route reflectors of both cases were Arista 7280SR (RR) and Cisco Nexus 3600-R (RR).

- While Arista 7280SR, Cisco NCS 5500, HUAWEI NE40E-X8A and Keysight (Ixia) IxNetwork successfully established VLAN-based EVPN-MPLS, we also established VLAN-aware-bundle EVPN-MPLS between Arista 7280SR and Juniper MX204. The route reflectors were Arista 7280SR (RR), Cisco NCS 5500 (RR) and Cisco NCS540 (RR).

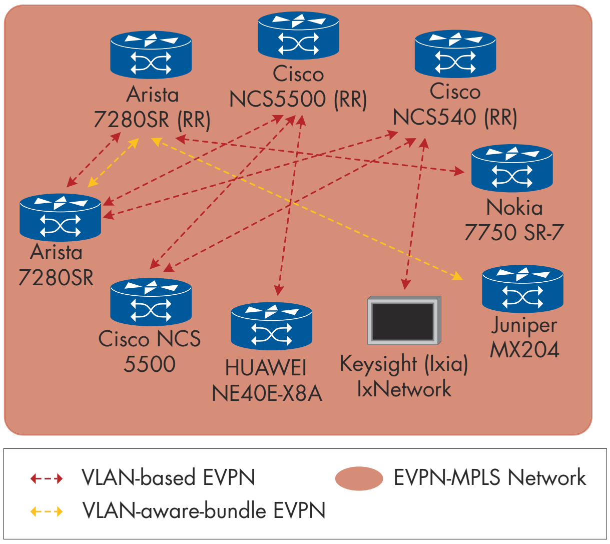

Figure 8: Symmetric IRB-MPLS

Integrated Routing and Switching (IRB) - Asymmetric IRB

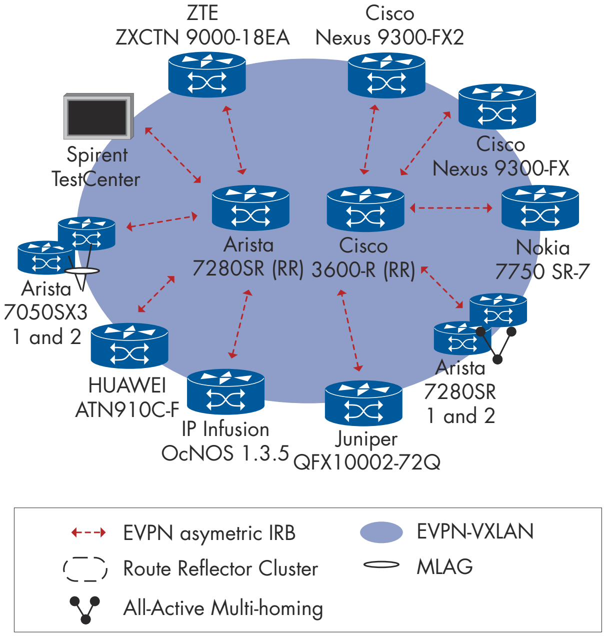

Figure 9: Asymmetric IRB-VXLAN

The asymmetric IRB semantic requires both IP and MAC lookups at the ingress NVE with only MAC lookup at the egress NVE, their results per EVPN-MPLS or EVPN-VXLAN were:

- PEs per EVPN-VXLAN with IRB: Arista 7050SX3, Arista 7280SR, Cisco Nexus 9300-FX, Cisco Nexus 9300-FX2, HUAWEI ATN910C-F, IP Infusion OcNOS 1.3.5, Juniper QFX10002-72Q, Nokia 7750 SR-7, Spirent TestCenter, and ZTE ZXCTN 9000-18EA. The route reflectors were Arista 7728SR and Cisco 3600-R.
- PEs with VLAN-aware-bundle EVPN-MPLS: Arista 7280SR and Juniper MX204.

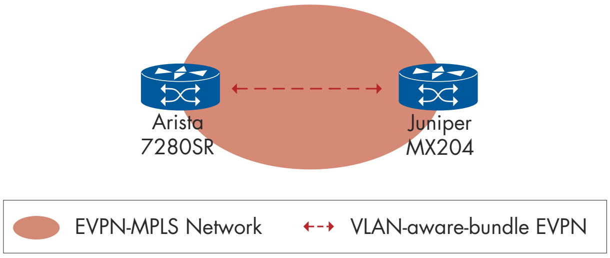

Figure 10: Asymmetric IRB-MPLS

IP Subnet Routing

In an EVPN network environment, there is a requirement for IP prefix advertisement for subnets and IPs residing behind an IRB interface. This scenario is referred to as EVPN IP-VRF-to-IP-VRF. The EVPN prefix advertisement draft provides different implementation options for the IP-VRF-to-IP-VRF model:

1. Interface-less model, where no Supplementary Broadcast Domain (SBD) and overlay index are required
2. Interface-full with unnumbered SBD IRB model, where SBD is required as well as MAC addresses as overlay indexes
3. Interface-full with SBD IRB model, where SBD is required as well as Gateway IP addresses as overlay indexes

This resulted in the following tested combinations:

- PEs of EVPN-VXLAN with the interface-less model were: Arista 7280SR, Cisco Nexus 3600-R, Cisco Nexus 9300-FX, Cisco Nexus 9300-FX2, HUAWEI ATN910C-F, Juniper QFX10002-72Q, Nokia 7750 SR-7, Spirent TestCenter, and ZTE ZXCTN 9000-18EA. We tested the interface-full model per EVPN-VXLAN with IP Infusion OcNOS 1.3.5, Nokia 7750 SR-7 and Spirent TestCenter. The route reflectors were Arista 7280SR (RR) and Cisco 3600-R (RR).

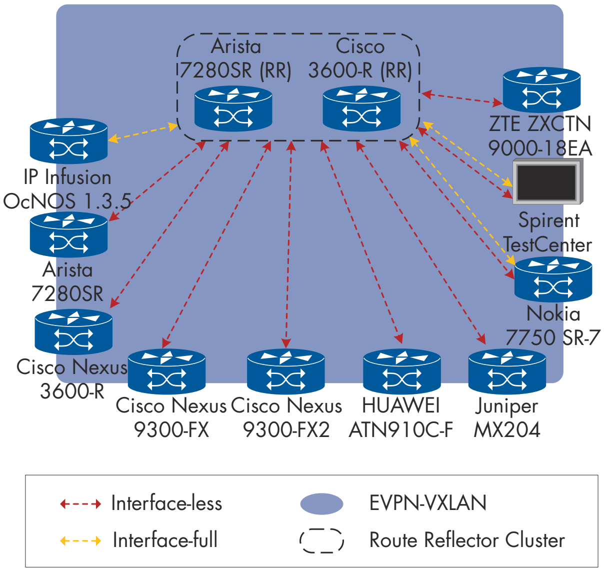

Figure 11: IP Subnet Routing-VXLAN

- PEs of EVPN-MPLS were: Arista 7280SR, Cisco NCS 5500, Cisco Nexus 9300-FX2, HUAWEI ATN910C-F, and HUAWEI NE40E-X8A. Both Arista 7280SR and Juniper MX204 tested the interface-less model using VLAN-aware-bundle EVPN-MPLS. The route reflector was Cisco NCS540 (RR)

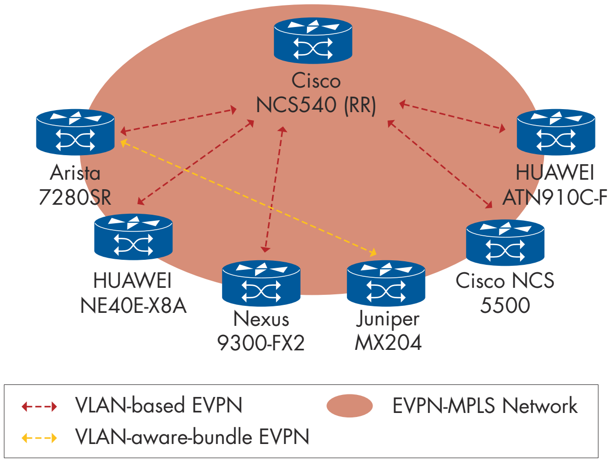

Figure 12: IP Subnet Routing-MPLS

All-Active Multi-homing

We tested the All-Active multi-homing on EVPN and verified that the emulated CE which was connected to two or more PEs via two Ethernet links (referred to as an Ethernet segment) and all multi-homed PEs forwarded known unicast traffic to/from that Ethernet segment for a given VLAN. In addition, we tested that BUM traffic stopped on the PE and did not cause any loop. Finally, we verified IRB attached to the All-Active multi-homed EVPN.

- The following devices acted as All-Active multi-homed PEs, per VLAN-based EVPN-VXLAN: and Nokia 7750 SR-7, also with IRB was: Arista 7280SR and Cisco Nexus 3100-V-2; per VLAN-aware-bundle with IRB: Arista 7280SR and Juniper QFX5110-48S. The route reflectors were Cisco 3600-R (RR) and Arista 7280SR (RR).

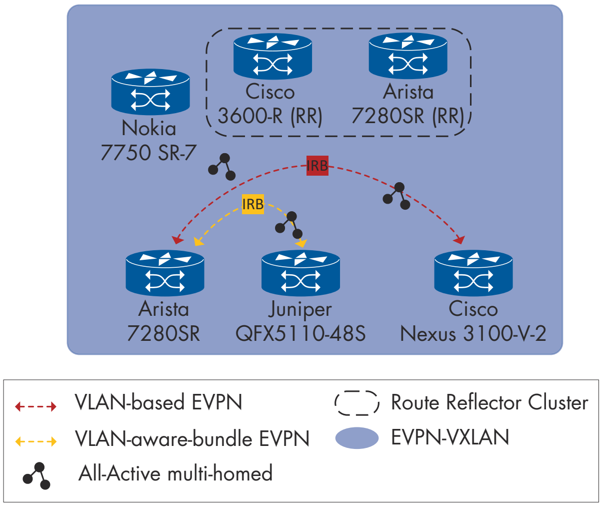

Figure 13: All-Active Multi-homing-VXLAN

- We tested two types of EVPN-MPLS services: per VLAN-based consisting of All-Active multi-homed PEs: Nokia 7750 SR-7 and NCS 5500, also with IRB were: Arista 7280SR, Cisco NCS 5500 and HUAWEI NE40E-X8A using route reflectors Arista 7280SR (RR) and Cisco NCS540 (RR). Single-homed PEs: Cisco NCS 5500 and Cisco Nexus 3600-R; Per VLAN-aware-bundle consisting of All-Active multi-homed PEs: HUAWEI NE40E-X8A with IRB were Arista 7280SR and Juniper MX204.

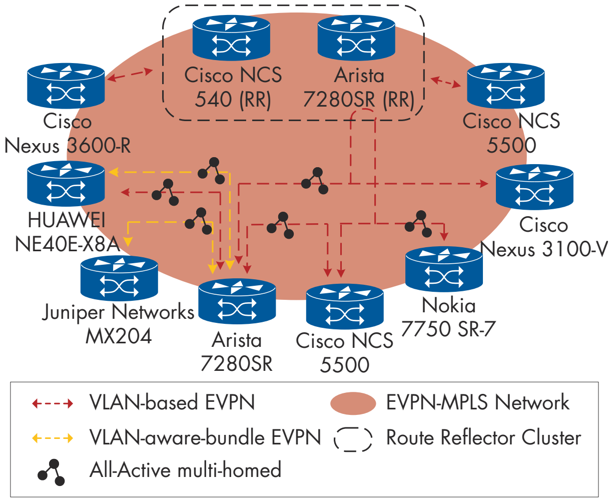

Figure 14: All-Active Multi-homing-MPLS

EVPN Interconnect

The role of the data center gateway was also verified in the EVPN interoperability test. Together with VXLAN and MPLS participants we built two different network typologies, to verify the use of Ethernet VPN (EVPN) by different vendors (multi-vendor environment) could extend data center business services over a MPLS network in the lab, which represents, a geographically dispersed campus or corporate network in the real world.

EVPN-VXLAN and EVPN-MPLS Interworking

This test focused on the EVPN-VXLAN extension over an EVPN-MPLS network.

Three vendors successfully participated in the gateway role: Cisco ASR9000, HUAWEI NE40E-X8A, and Nokia 7750 SR-7. The following devices acted as PE of EVPN-VXLAN: Arista 7050SX3, Arista 7280SR and HUAWEI NE40E-F1A. The PE of EVPN-MPLS was Cisco NCS 5500.

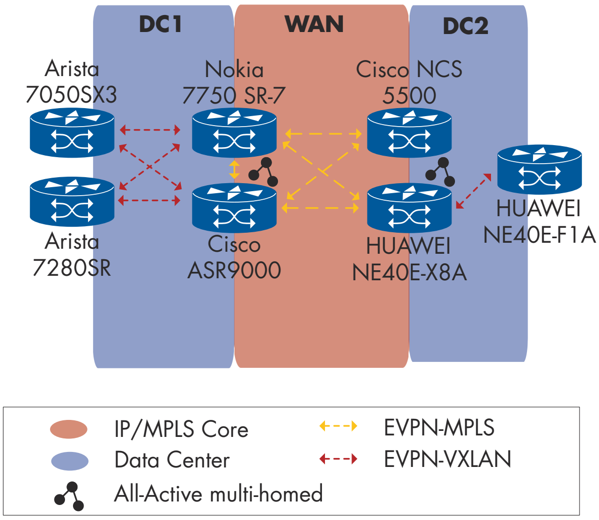

Figure 15: EVPN-VXLAN and EVPN-MPLS Interworking

EVPN Interworking with IPVPN

This test focused on interconnect of EVPN-VXLAN and EVPN-MPLS over a IP/MPLS VPN. We successfully tested the following devices:

- EVPN-MPLS PE: Arista 7280SR and Arista 7050SX3
- EVPN-MPLS/IP-VPN Gateway: Arista 7280SR and Nokia 7750 SR-7
- IP-VPN/EVPN-VXLAN Gateway: Arista 7280SR, Cisco ASR9000 and HUAWEI NE40E-X8A
- EVPN-VXLAN PE: HUAWEI ATN950C and Cisco Nexus 3100-V
- Route reflector in EVPN-VXLAN: Cisco 3600-R (RR)

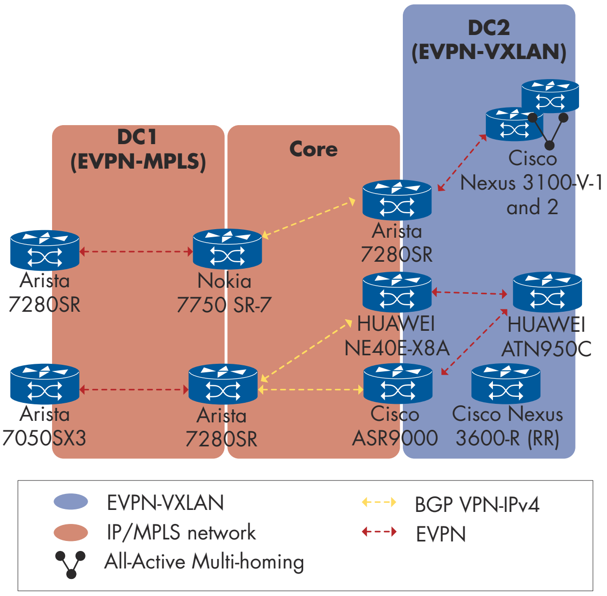

Figure 16: EVPN Interworking with IPVPN

EVPN Maintenance

The IETF EVPN OAM draft "salam-bess-evpn-oam-req-frmwk" defines the operations, administration and maintenance requirements and framework for EVPN. We tested the fault management and performance monitoring. In addition, we also tested the loop prevention of EVPN.

EVPN Loop Protection

This test verified that the EVPN PEs were able to automatically resolve a loop over different Attached Circuits in the same broadcast domain. RFC7432 defines the mechanism via RT-2 route to prevent loop in the EVPN forwarding plane risked by a backdoor connection (global loop) between multiple ACs. EVPN Loop Protection as described in "draft-snr-bess-evpn-loop-protect" extends this to the data plane so the PEs shall stop BUM traffic from being forwarded by disabling the AC or by dropping the looped frames.

We first observed the EVPN under normal condition on the PE to ensure that the loop can be added without any interruption. As expected, once the loop was added (as shown in the Figure 17), the link on the PE was active, and the PE installed the duplicate MAC address. The loop did not cause any congestion. The loop detection log was visible in the log. The remote PE removed the blackhole MAC from the MAC-VRF table received from the RT-2 route. There was no packet loss during this process. Similarly, after removing the loop, the blackhole MAC route was revoked without any packet loss.

- The following devices successfully participated in the test: HUAWEI NE40E-X8A and Nokia 7750 SR-7

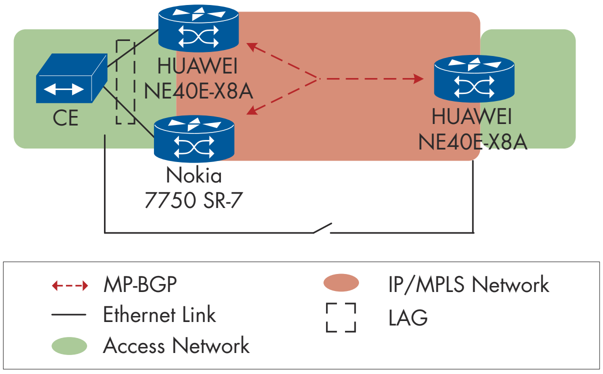

Figure 17: EVPN Loop Protection

EVPN Performance Management

We verified the interoperability of the loss measurement as well as delay measurement as defined in Y.1731.

We started with the loss measurement, first by observing 0 packet loss via MEP counter on the PE, while sending the baseline traffic at 100 Mbit/s. Then we repeated the step while introducing 30% packet loss, implemented on a third-party switch that was connected between the PEs. By changing the bandwidth of the switch, we received 70 Mbit/s at the traffic generator and observed 30% packet drop in the MEP counter on the PE as expected. Finally, we removed the impairment from the switch and observed no packet loss as shown in initial state on the PE.

During the delay measurement, we first recorded the baseline delay value and compared it between the traffic generator and PE. Then we added delay over the third-party switch connected between the PEs. As expected, we observed that the PE showed the increased delay value.

- The following devices provided the Maintenance End Points (MEPs): HUAWEI NE40E-X8A, Juniper MX204, and Nokia 7750 SR-7.

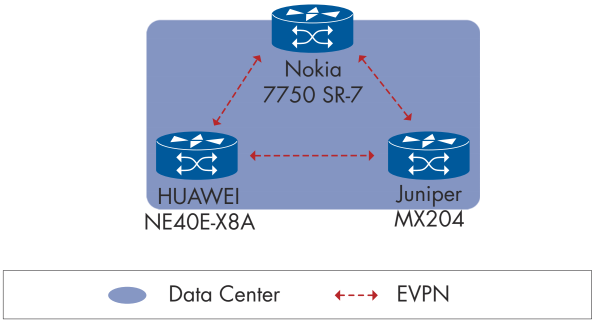

Figure 18: EVPN Performance Management

EVPN Fault Management

This test focused on Connectivity Fault Management (CFM) functionality as defined in Y.1731 to detect link failure for a network service. In this test, we first observed that no packet loss was shown in the EVPN service to ensure a baseline setup. Then we introduced a 100% packet loss while traffic was running. This was performed on a third-party switch that was connected between the PEs, we removed the configuration on the switch without losing any connection to the PEs. As expected, the PE detected the link failure via CFM packet drop. After removing the failure on the switch, the PE also detected the link was up as shown in the baseline setup.

- The following devices provided the Maintenance End Points (MEPs): Ericsson 6672, HUAWEI NE40E-X8A, Juniper MX204, and Nokia 7750 SR-7.

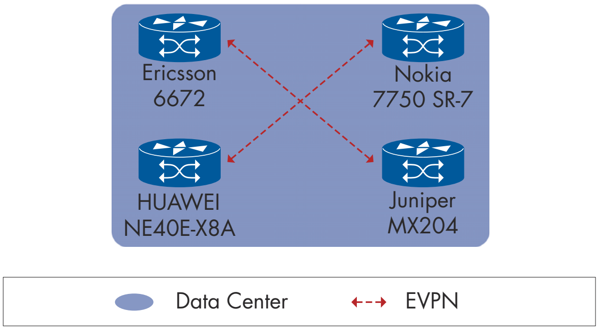

Figure 19: EVPN Fault Management

Segment Routing

Segment Routing provides complete control over the forwarding paths by using Source Packet Routing in Networking (SPRING). Based on the traditional MPLS or IPv6 forwarding plane, it simplifies the network protocols and provides end-to-end traffic engineering without any additional signaling or midpoint fabric-state.

The source routing architecture allows the use of different control-plane models. The centralized model, using external controllers for path computation, is tested in the following SDN sections of this white paper. The distributed model, Network Elements (NEs) using dynamic routing protocols, is tested in this Segment Routing section.

In this section, we will present the SR tests based on MPLS or SRv6 data plane, using IGP or BGP with SR extensions to allocate and distribute Segment Identifiers (SIDs). We will also demonstrate some use cases of Segment Routing to perform failure protection, Multi-Plane network slicing, network OAM, etc.

IPv4/IPv6 VPN over SRv6

The draft “dawra-bess-srv6-services” defines procedures and messages for BGP SRv6-based EVPNs and L3 VPNs in order to provide a migration path from MPLS-based VPNs to SRv6 based VPNs.

In order to provide SRv6-VPN service with best-effort connectivity, the egress PE signals an SRv6-VPN SID with the VPN route. The ingress PE encapsulates the VPN packet in an outer IPv6 header where the destination address is the SRv6-VPN SID provided by the egress PE. The underlay between the PEs only need to support plain IPv6 forwarding.

In our test, vendors configured IPv4 L3VPN and IPv6 EVPN L3VPN over SRv6, the egress node performed the END.DT4/END.DT6 function. BGP was used to advertise the reachability of prefixes in a particular VPN from an egress Provider Edge (egress-PE) to ingress Provider Edge (ingress-PE) nodes. EVPN VPWS over SRv6 was not tested, due to unavailability of vendor support.

We sent bidirectional traffic between the ingress PE and egress PE, therefore in each group, both PEs were performing the encapsulation and END.DT4/END.DT6 functions.

Additionally, we added P node without SRv6 functions enabled to perform IPv6 forwarding only, and verify that plain IPv6 forwarding is enough for the transit node, no need for supporting SRv6 capability.

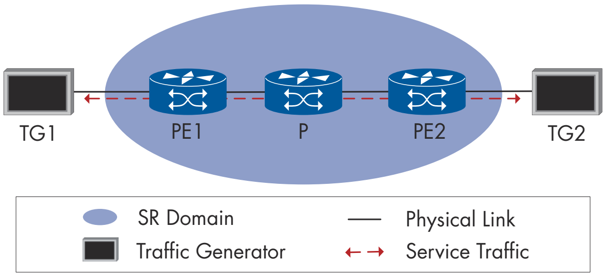

Figure 20: SRv6 Test Scenario

Scenario	PE1	P	PE2
IPv4 L3VPN over SRv6	HUAWEI NE9000-8	Cisco NCS 540, HUAWEI NE40E-X8A	Cisco NCS 5500
	HUAWEI NE9000-8	Cisco NCS 540, HUAWEI NE40E-X8A	Keysight (Ixia) IxNetwork
	Cisco NCS 5500	HUAWEI NE9000-8, HUAWEI NE40E-X8A	Keysight (Ixia) IxNetwork
	HUAWEI NE9000-8	Cisco NCS 540, HUAWEI NE40E-X8A	Spirent TestCenter
	Cisco NCS 5500	HUAWEI NE9000-8, HUAWEI NE40E-X8A	Spirent TestCenter
IPv6 EVPN L3VPN over SRv6	HUAWEI NE40E-F1A	HUAWEI ATN950C	Keysight (Ixia) IxNetwork
	HUAWEI NE40E-F1A	HUAWEI ATN950C	Spirent TestCenter

Table 3: SRv6 Test Pairs

Segment Routing TI-LFA

Loop-free alternate (LFA) and remote LFA (RLFA) have been used to provide fast-reroute protection. However, depending on the network topology, the percentage of destinations protected by LFA and remote LFA is usually less than 100 percent.

Topology Independent Loop-free Alternate (TI-LFA) extends the concept of LFA and remote LFA by allowing the Point of Local Repair (PLR) to use deeper label stacks to construct backup paths. In addition, the TI-LFA imposes the constraint that the backup path used by the PLR be the same path that a packet takes once the interior gateway protocol (IGP) has converged for a given failure scenario. This path is referred to as the post-convergence path. TI-LFA relies on segment routing to build a protection mechanism based on proven IP-FRR concepts.

TI-LFA provides protection against link failure, node failure, and local Shared Risk Link Group (SRLG) failure. In Segment Routing protection test, we tested TI-LFA for link failure and SRLG failure on MPLS data plane and SRv6 data plane. Since not all vendors support TI-LFA with SRLG, we split the MPLS data plane protection test into two scenarios: TI-LFA with SRLG and TI-LFA without SRLG. TI-LFA for SRv6 is tested in separate scenario without SRLG.

To observe the TI-LFA protection, we agreed on below setups and to observe only unidirectional traffic from the PLR (ingress PE) to the Egress PE, in this case only the PLR (ingress PE) was performing the TI-LFA protection. Since there are vendors supporting only TI-LFA as P node, we separated the PLR roles to PLR(PE) and PLR(P). While vendor was acting as PLR(P) role, the Traffic Generator was acting as PE and sending labeled traffic to the PLR(P) node, and PLR(P) node performed the TI-LFA protection.

For the TI-LFA with Local SRLG Protection test, vendors configured TI-LFA based on MPLS. SRLG was enabled on the PLR Node and included two links (Link1 and Link2) in the SRLG. We checked the TI-LFA calculation result, it showed that the backup path was calculated but it's not using the link in same SRLG (Link2). We disconnected the primary link while traffic was running through the primary path and measured the service interruption time based on the packet loss. Figure 21 and Table 4 is the setup for the first scenario.

For the TI-LFA with Link Protection test, vendors created a similar setup and configuration but without the link for SRLG. Figure 22 and Table 5 is the setup for the second scenario.

For the TI-LFA for SRv6, vendors created a setup without SRLG. Figure 23 and Table 6 is the setup for the third scenario.

The results showed that in all scenarios, all vendors performing the TI-LFA protection can switchover to the backup path in less than 50ms.

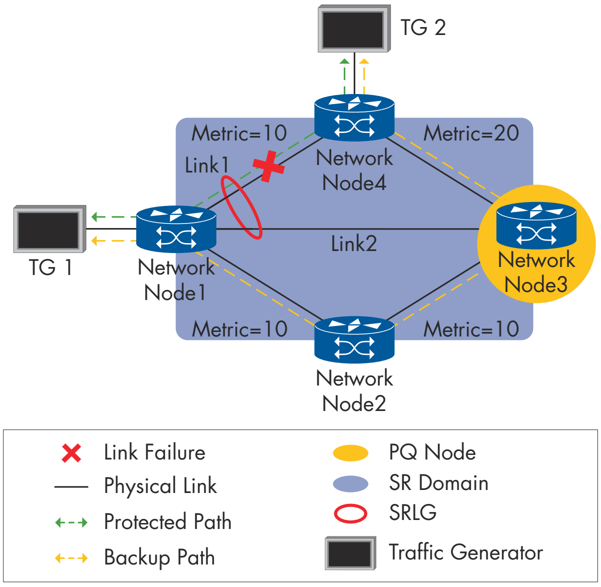

Figure 21: TI-LFA with Local SRLG Protection

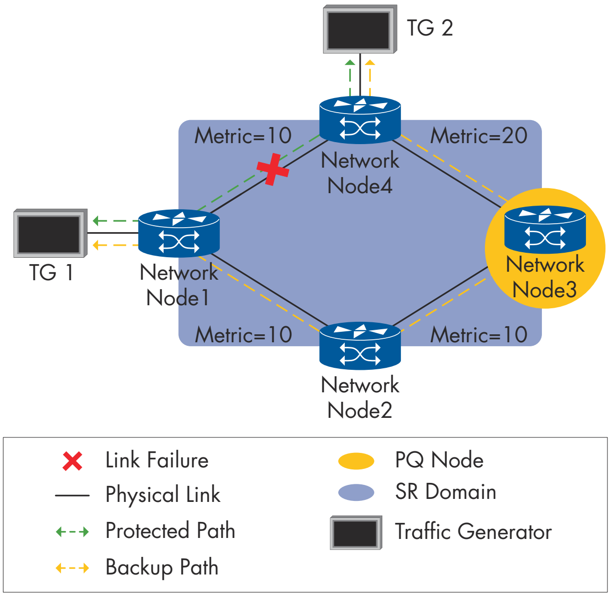

Figure 22: TI-LFA with Link Protection

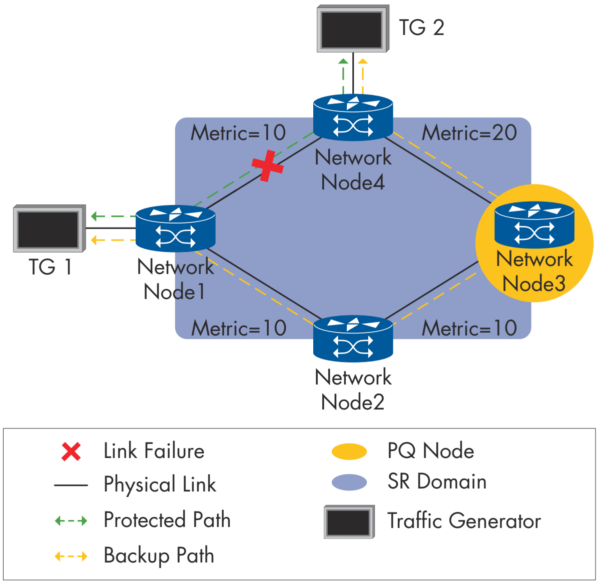

Figure 23: TI-LFA for SRv6

Test Setup	PLR (PE) (Network Node 1)	P Node (Network Node 2)	PQ Node (Network Node 3)	Egress PE (Network Node 4)
1	Juniper MX204	Cisco ASR 9000	Nokia 7750 SR-7	Ericsson 6675
2	Ericsson 6675	Juniper MX204	Cisco ASR 9000	Nokia 7750 SR-7
3	Cisco ASR 9000	Nokia 7750 SR-7	Ericsson 6675	Juniper MX204

Table 4: TI-LFA with SRLG Setups

Test Setup	PLR (PE) (Network Node 1)	P Node (Network Node 2)	PQ Node (Network Node 3)	Egress PE (Network Node 4)
1	ECI Neptune 1300	Arista 7280SR	ZTE 6180H	ZTE 9000-18EA
2	ZTE 6180H	ZTE 9000-18EA	Arista 7280SR	ECI Neptune 1300
3	Arista 7280SR*	Cisco NCS 5500	Cisco ASR 9000	Nokia 7750 SR-7

*only acting as PLR (P), PE was simulated by TG

Table 5: TI-LFA without SRLG Setups

Test Setup	PLR (PE) (Network Node 1)	P Node (Network Node 2)	PQ Node (Network Node 3)	Egress PE (Network Node 4)
1	Cisco NCS 5500	Cisco NCS 540	HUAWEI NE40E-X8A	HUAWEI NE9000-8
2	HUAWEI NE40E-X8A	HUAWEI NE9000-8	Cisco NCS 540	Cisco NCS 5500

Table 6: TI-LFA for SRv6 Setups

Segment Routing Label Switched Path Ping/Traceroute

The RFC 8287 defines the LSP ping and traceroute method for segment routing (SR) 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.

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 a 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. After ping/traceroute test between all vendors in the test network, we gathered all results and come up to below table including all of the successful results. Each successful ping/traceroute result pair is marked with 'Y', not tested combinations are marked with '/'.

In this year, vendors demonstrated good interoperability results, most vendors can ping/traceroute each other via Segment Routing LSP but still there are some interoperability issues found during the test. Some vendors had a different understanding and implementation of 'RFC 8287 - Label Switched Path (LSP) Ping/Traceroute for Segment Routing (SR)', causing SR ping/traceroute failure or unexpected value in MPLS echo packets. Fortunately, some vendors fixed the issues during the test, leading to a better result for SR ping/traceroute test.

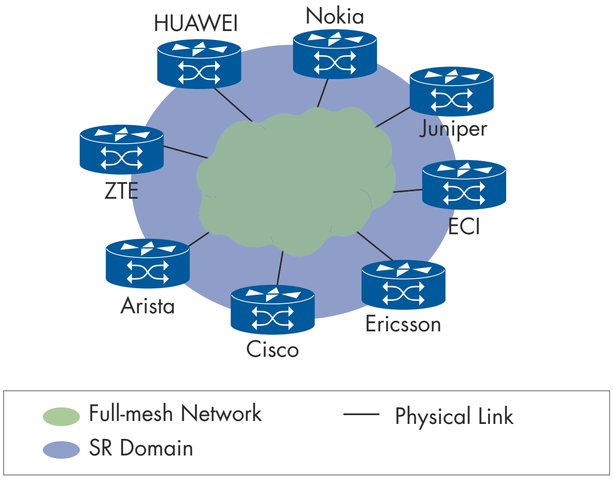

Figure 24: Full-mesh Network for Ping/Traceroute

	A	C	NX	ECI	E	H	J	N	ZTE
A									
C	Y								
NX	/	/							
ECI	Y	Y	Y						
E	Y	Y	Y	/					
H	Y	Y	Y	/	Y				
J	Y	Y	Y	Y	Y	Y			
N	Y	Y	/	Y	Y	Y	Y		
ZTE	/	Y	Y	Y	/	/	Y	Y	

Table 7: SR LSP Ping Results[1]

	A	C	NX	ECI	E	H	J	N	ZTE
A									
C	Y								
NX	/	/							
ECI	Y	Y	Y						
E	Y	Y	Y	/					
H	Y	Y	Y	/	Y				
J	Y	/	/	/	/	/			
N	Y	Y	/	Y	Y	Y	/		
ZTE	/	/	Y	Y	/	/	Y	Y	

Table 8: SR LSP Traceroute Results

	Cisco NCS 5500	Cisco NCS 540
HUAWEI NE9000-8	Y	Y
HUAWEI NE40E-X8A	Y	Y

Table 9: SRv6 LSP Ping/Traceroute Results

Segment Routing Anycast

In this section, we verified that the Segment Routing Anycast Segment could be used to disjoint traffic forwarding paths within dual plane networks.

Vendors configured Anycast Segment Identifier (Anycast SID) for all DUTs, DUTs in same Anycast group are configured with same Anycast SID. We sent three service traffic from the traffic generator to PE1, PE1 disjoints the traffic to three data planes according to the Anycast SID and encapsulates the Anycast SID into the packets.

Initially, all links were configured with the default metric and all traffic was forwarded along the shortest path. Then we increased the metric of the shortest path and observed the traffic was switched to the other anycast node within the same anycast group.

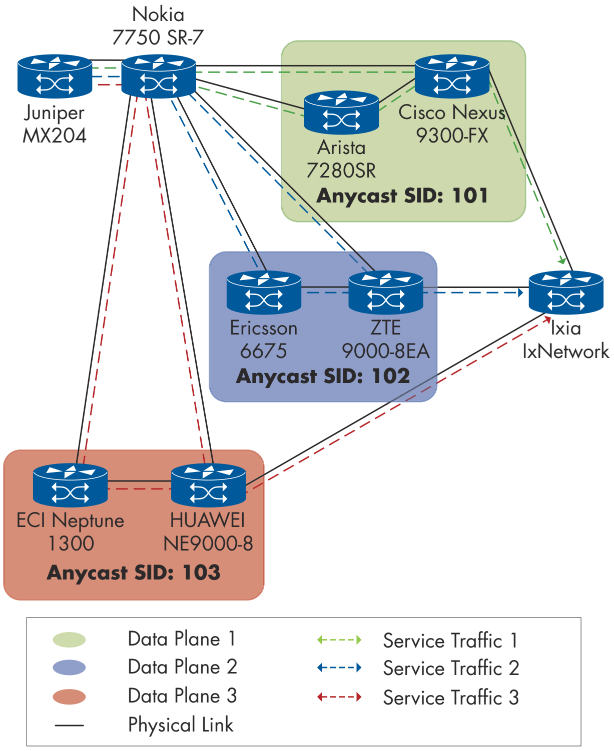

Figure 25: Segment Routing Anycast

Additionally, we tested to cut the shortest path link to one of the anycast node in each data plane. Traffic could be switched over to another anycast node in the same anycast group, showing the traffic protection of anycast nodes within the same data plane.

Vendors participating in the Segment Routing Anycast Segment tests were:

- PE1: Juniper MX204
- P: Nokia 7750 SR-7
- DUT: Arista 7280SR, Cisco Nexus 9300-FX, ECI Neptune 1300, Ericsson 6675, HUAWEI NE9000-8, ZTE ZXCTN 9000-8EA
- PE2: Keysight (Ixia) IxNetwork

[1] A: Arista 7280SR, C: Cisco ASR 9000, NX: Cisco Nexus 9300-FX, ECI: ECI Neptune 1300, E: Ericsson 6675, H: HUAWEI NE9000-8, J: Juniper MX204, N: Nokia 7750 SR-7, ZTE: ZTE ZXCTN 9000-8EA/ZTE ZXCTN 6180H

BGP Segment Routing: BGP-LU

Segment Routing can be used in large scale networks as a simple solution to provide traffic engineering and fast re-route capabilities. 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.

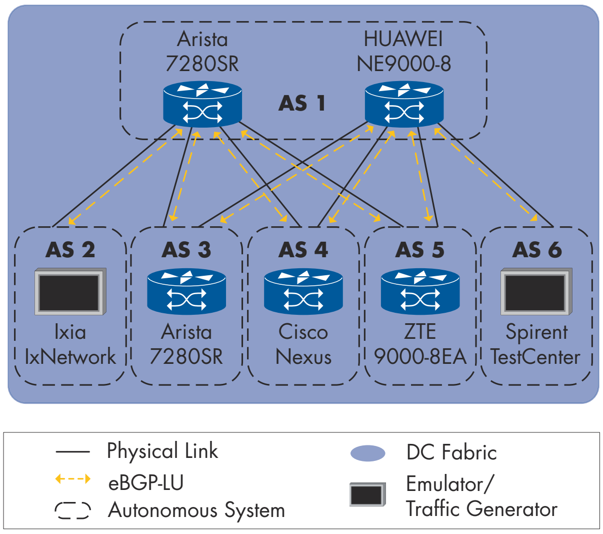

Figure 26: Segment Routing BGP-Label Unicast

We tested BGP Segment Routing using BGP Labeled Unicast (BGP-LU) NLRI in a typical Clos topology with two Spines and five Leaves. Vendors configured the Leaf Nodes (DUTs) to advertise BGP 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, Arista Spine Node enabled BGP Segment Routing capability to generate MPLS labels for the BGP updates received from Leaf Nodes. Full-mesh traffic was tested between all Leaf Nodes and the Node Emulators – Ixia and Spirent.

In this case, we were focusing on the BGP SR capability using BGP-LU on Leaf Switches. Both Ixia and Spirent were used as Traffic Generator and Emulated Leaves.

Vendors participating in the BGP-LU tests were:

- Spine: Arista 7280SR, HUAWEI NE9000-8
- Leaf: Arista 7280SR, Cisco Nexus 9300-FX, ZTE ZXCTN 9000-8EA, Keysight (Ixia) IxNetwork, Spirent TestCenter

Segment Routing TWAMP

TWAMP uses the methodology and architecture of OWAMP [RFC4656] to define an open protocol for measurement of two-way or round-trip metric, in addition to the one-way metric of OWAMP.

TWAMP employs time stamps applied at the echo destination (reflector) to enable greater accuracy (processing delays can be accounted for). The TWAMP measurement architecture is usually comprised of only two hosts with specific roles, and this allows for some protocol simplifications, making it an attractive alternative to OWAMP in some circumstances.

In this test, we tested the Segment Routing TWAMP, in both light and full mode. Vendors configured their devices as TWAMP Sender and Reflector to test the interoperability of TWAMP. We sent bidirectional traffic between the TWAMP sender and reflector, while a Calnex SNE was used to apply delay in one direction.

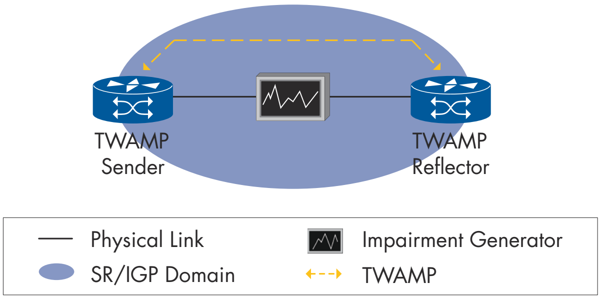

Figure 27: Segment Routing TWAMP

At first, we sent traffic with no delay applied, we observed the TWAMP measured traffic delay matched the statistics shown on the traffic generator. Then we applied delay on one direction of the traffic, we observed the delay on the traffic generator, and TWAMP also showed the increase of delay in the two-way delay measurement result.

As some vendors only supported light mode or full mode, and some vendor only support to work as the reflector, we executed the test in below combinations:

Mode	Sender	Reflector
Light	Ericsson 6675	HUAWEI NE9000-8
	HUAWEI NE9000-8	Ericsson 6675
	Ericsson 6675	Cisco ASR 9000
Full	Juniper MX204	Arista 7280SR
	Juniper MX204	Cisco ASR 9000

Table 10: TWAMP Test Pairs

Seamless BFD

Bidirectional Forwarding Detection (BFD) is a widely used failure detection mechanism for multiple protocols and applications. It can be improved to expand failure detection coverage and to allow BFD usage for wider scenarios. Seamless Bidirectional Forwarding Detection (S-BFD) is a simplified mechanism to improve the efficiency of Bidirectional Forwarding Detection. It can provide applications a smooth and continuous operational experience.

In this case, we tested the Seamless BFD for Segment Routing Traffic Engineering (SR-TE) Tunnel. Vendors configured SR-TE tunnel with a primary path and a backup path. Seamless BFD runs on both paths. We sent unidirectional traffic from PE1 (S-BFD initiator) to PE2 (S-BFD reflector). In a normal situation, the traffic was forwarded along the primary path. Then Calnex dropped all MPLS packets in the primary path, causing the Seamless BFD session over the primary path to fail. Upon failure of primary Seamless BFD, PE1 (S-BFD initiator) redirected the traffic over the backup path. The result showed that Seamless BFD can be used to detect the SR-TE LSP failure and trigger the traffic to be switched to the backup path.

Vendors participating in the S-BFD tests were:

- PE1, S-BFD initiator: Juniper MX204
- PE2, S-BFD reflector: Keysight (Ixia) IxNetwork
- P Nodes: Cisco ASR 9000, Nokia 7750 SR-7
- Impairment tool: Calnex SNE

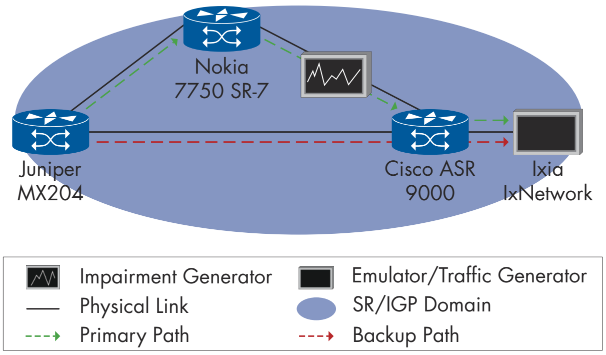

Figure 28: Seamless BFD

Flexible Algorithms

In a Multi-Plane network, traffic must be forwarded over a specific path. It is usually different from the shortest IGP path. Traffic engineering is used to compute the optimal path based on constraints.

The draft “draft-ietf-lsr-flex-algo” defines a solution that allows IGPs themselves to compute constraint-based paths over the network. It also specifies a way of using Segment Routing Prefix-SIDs to steer packets along the constraint-based paths. ISIS Flexible Algorithm Definition Sub-TLV (FAD Sub-TLV) is used to advertise the definition of the Flexible Algorithm.

Segment Routing Flexible Algorithms enriches the SR-TE solution by adding additional segments having different properties than the IGP Prefix segments.

In this test, we tested the multi-plane network slicing with Segment Routing Flexible Algorithms based on new ISIS FAD Sub-TLV extensions. Two Flexible Algorithms were configured on each PE. Each P node was configured with only one of the Flexible Algorithms. Two bidirectional service traffic flows were sent between PEs, each service was forwarded within the specific network plane.

Vendors participating in the Flexible Algorithms tests were:

- PE: Cisco NCS 5500, Juniper MX480
- P: Cisco ASR 9000, Juniper MX-204

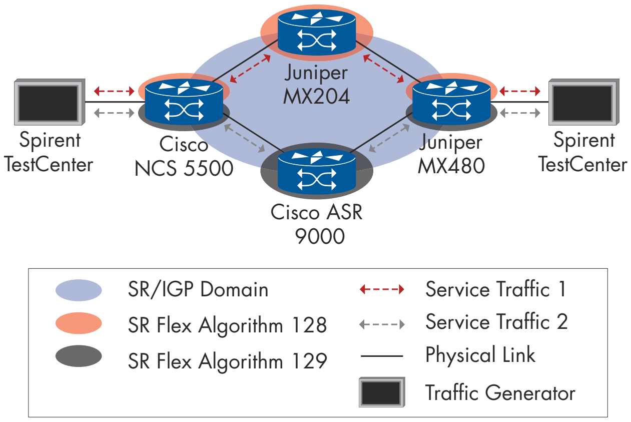

Figure 29: Segment Routing Flexible Algorithms

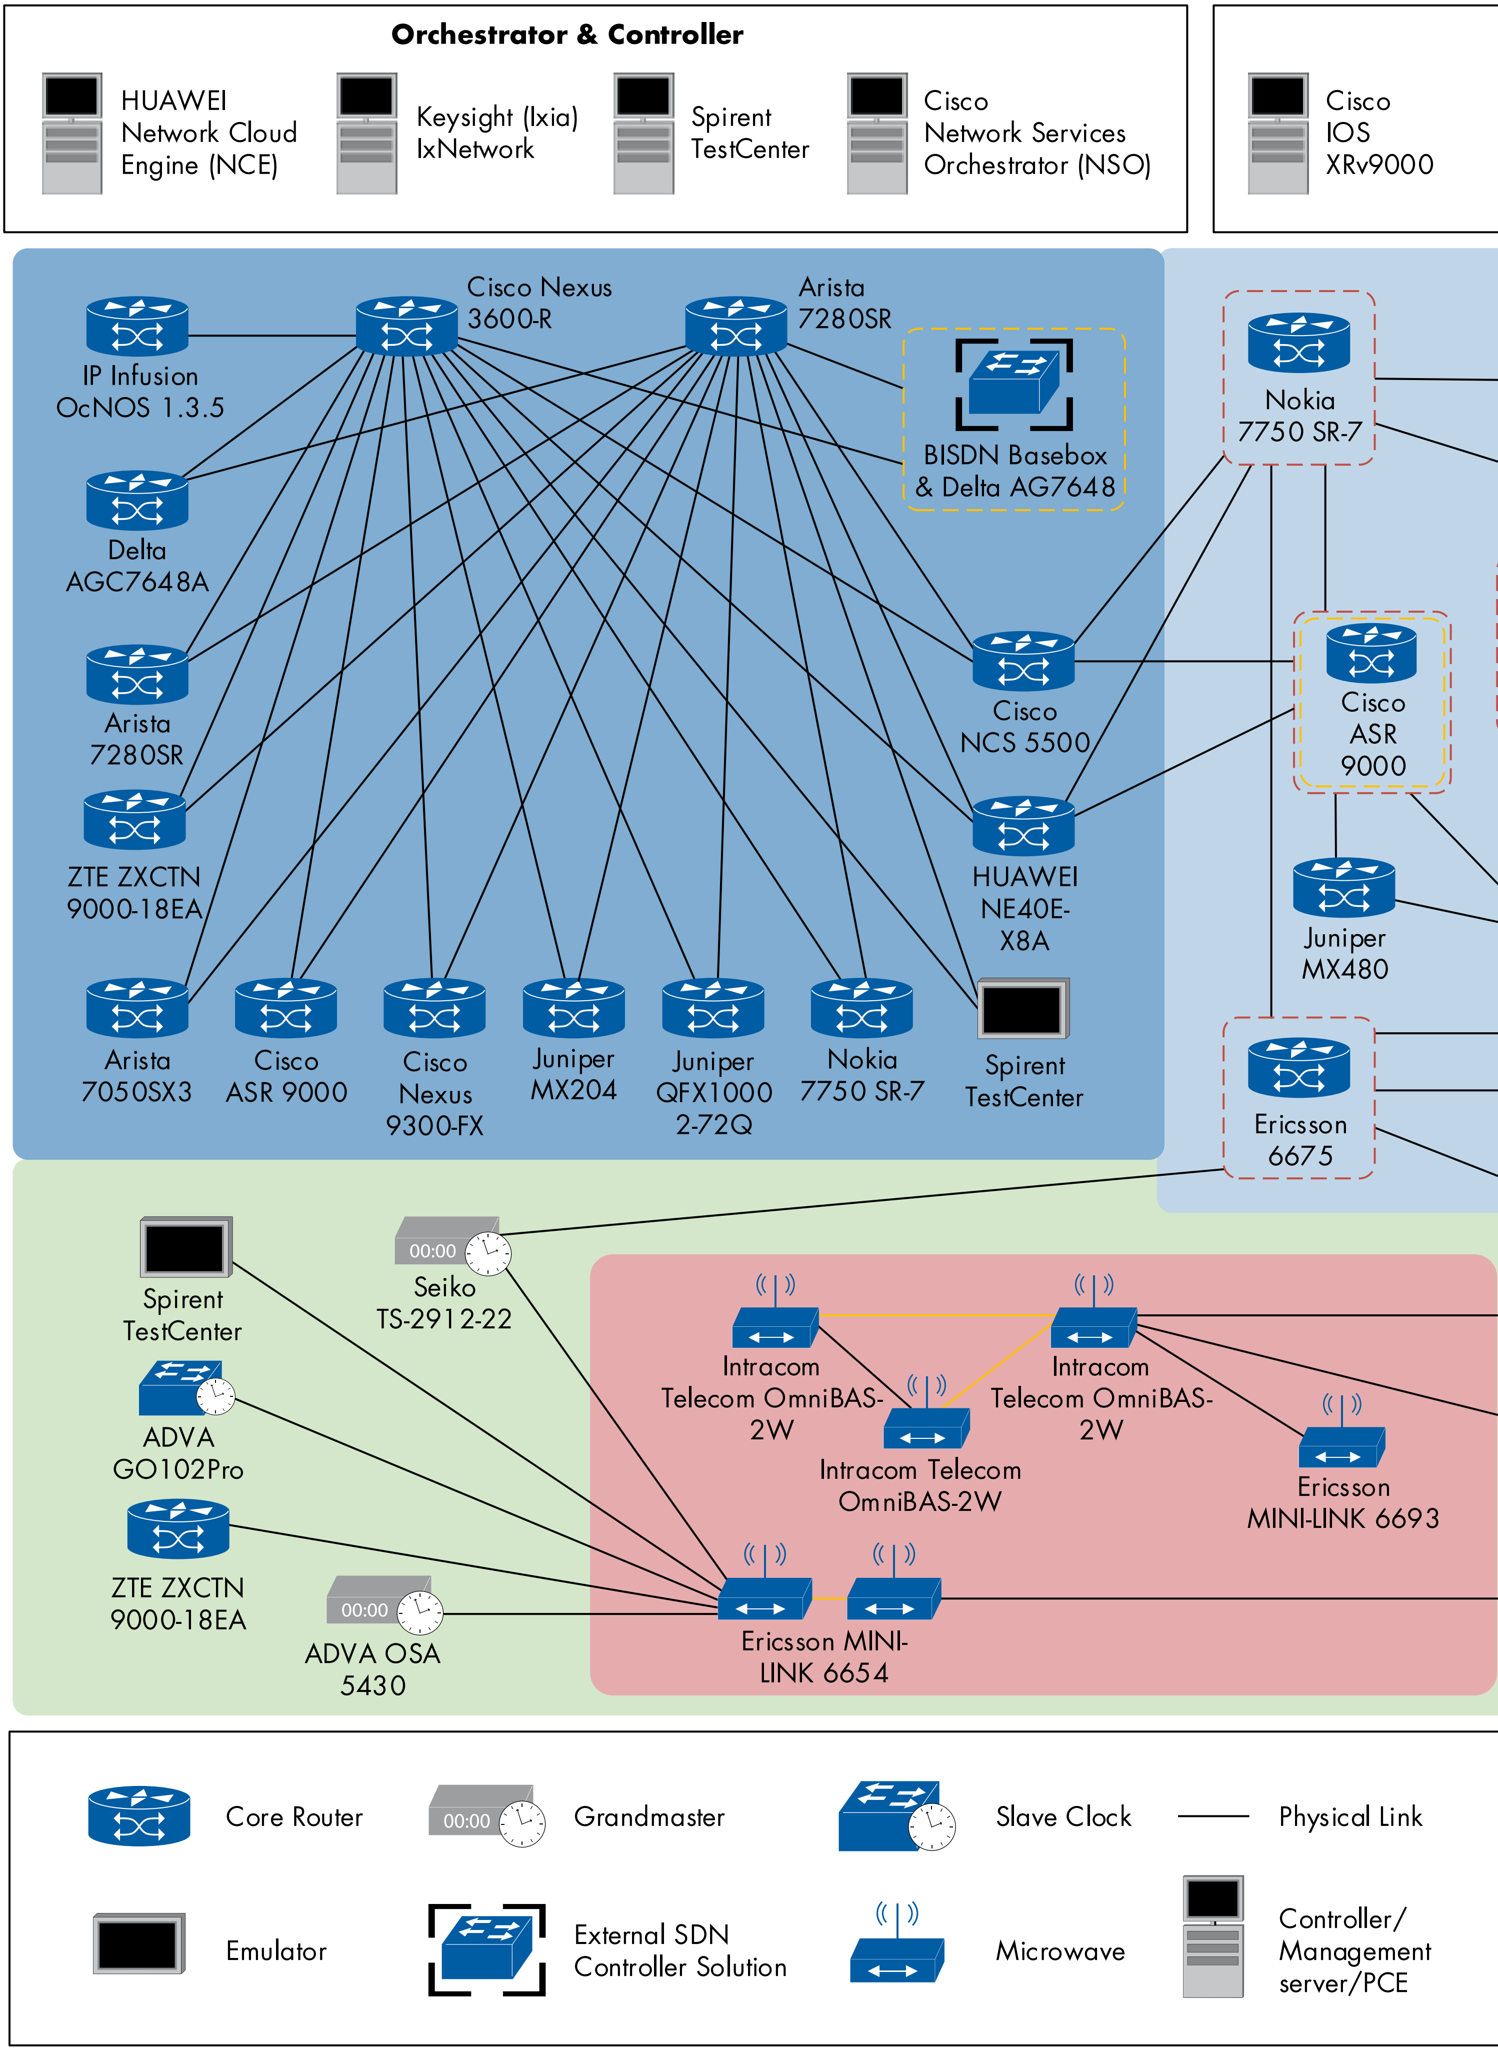
Orchestrator & Controller
HUAWEI Network Cloud Engine (NCE)
Keysight (Ixia) IxNetwork
Spirent TestCenter
Cisco Network Services Orchestrator (NSO)
Cisco IOS XRv9000
IP Infusion OcNOS 1.3.5
Cisco Nexus 3600-R
Arista 7280SR
BISDN Basebox & Delta AG7648
Delta AGC7648A
Arista 7280SR
ZTE ZXCTN 9000-18EA
Cisco NCS 5500
HUAWEI NE40E-X8A
Arista 7050SX3
Cisco ASR 9000
Cisco Nexus 9300-FX
Juniper MX204
Juniper QFX10002-72Q
Nokia 7750 SR-7
Spirent TestCenter
Nokia 7750 SR-7
Cisco ASR 9000
Juniper MX480
Ericsson 6675
Spirent TestCenter
Seiko TS-2912-22
ADVA GO102Pro
ZTE ZXCTN 9000-18EA
ADVA OSA 5430
00:00
00:00
Intracom Telecom OmniBAS-2W
Intracom Telecom OmniBAS-2W
Intracom Telecom OmniBAS-2W
Ericsson MINI-LINK 6693
Ericsson MINI-LINK 6654
Core Router
00:00
Grandmaster
Slave Clock
Physical Link
Emulator
External SDN Controller Solution
Microwave
Controller/ Management server/PCE

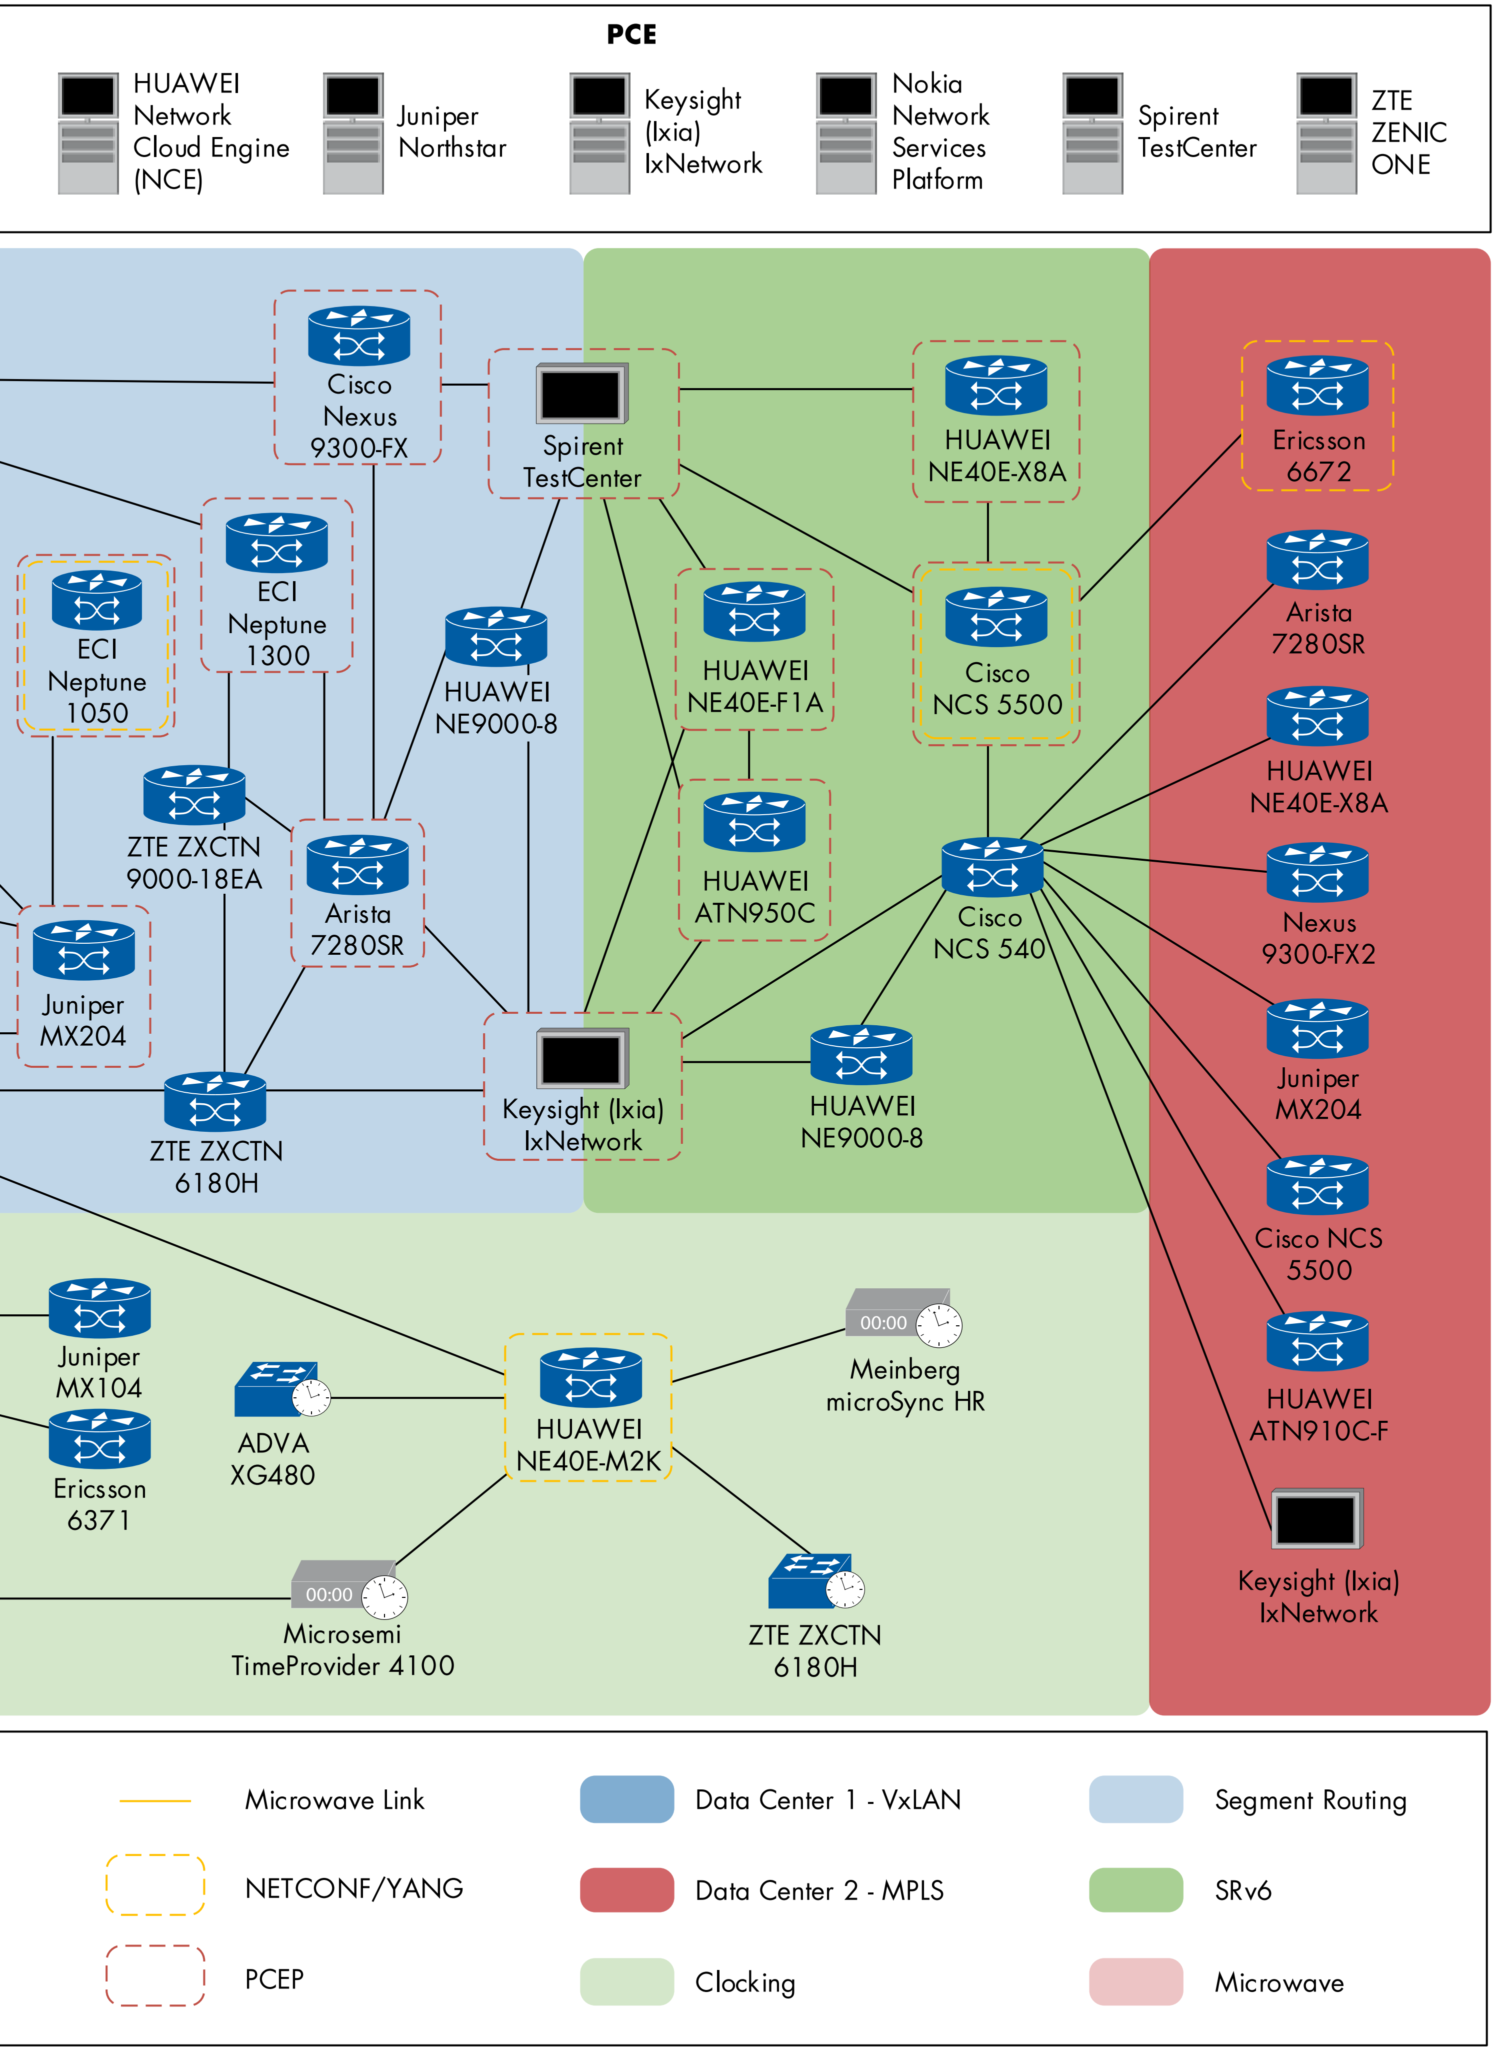
PCE
HUAWEI Network Cloud Engine (NCE)
Juniper Northstar
Keysight (Ixia) IxNetwork
Nokia Network Services Platform
Spirent TestCenter
ZTE ZENIC ONE
Cisco Nexus 9300-FX
Spirent TestCenter
HUAWEI NE40E-X8A
Ericsson 6672
ECI Neptune 1300
ECI Neptune 1050
HUAWEI NE9000-8
HUAWEI NE40E-F1A
Cisco NCS 5500
Arista 7280SR
HUAWEI NE40E-X8A
ZTE ZXCTN 9000-18EA
Arista 7280SR
HUAWEI ATN950C
Cisco NCS 540
Nexus 9300-FX2
Juniper MX204
Juniper MX204
ZTE ZXCTN 6180H
Keysight (Ixia) IxNetwork
HUAWEI NE9000-8
Cisco NCS 5500
Juniper MX104
ADVA XG480
HUAWEI NE40E-M2K
00:00
Meinberg microSync HR
HUAWEI ATN910C-F
Ericsson 6371
Keysight (Ixia) IxNetwork
00:00
Microsemi TimeProvider 4100
ZTE ZXCTN 6180H
Microwave Link
Data Center 1 - VxLAN
Segment Routing
NETCONF/YANG
Data Center 2 - MPLS
SRv6
PCEP
Clocking
Microwave

SDN & NFV

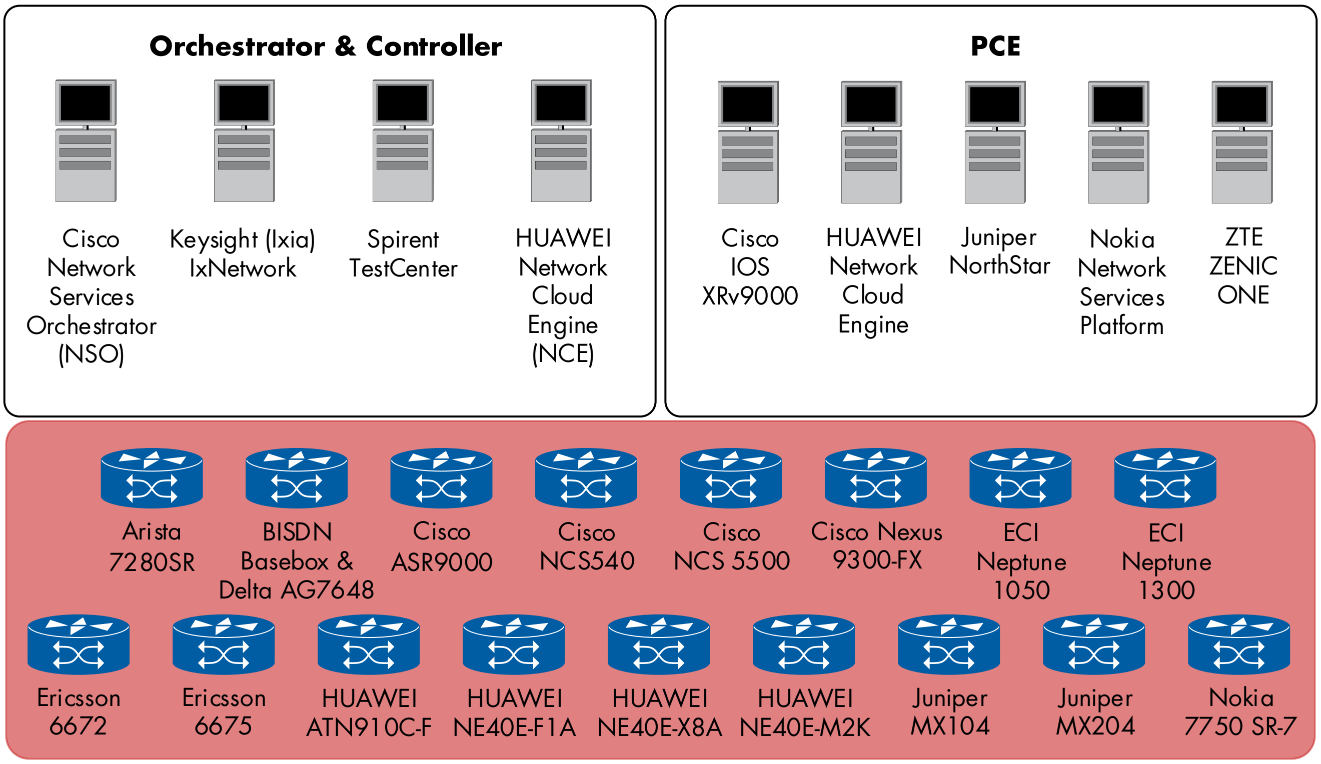

Figure 30: SDN Topology

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 shows a wide range of interoperability between vendors in multiple scenarios including some advanced cases.

Path Computation Element Protocol (PCEP)

The mechanism to communicate between a Path Computation Element (PCE) and a Path Computation Component (PCC) is described in RFC5440. All messages between the PCE and PCC run over TCP. Label Switched Paths are computed by the PCE and can be initiated by either the PCC or PCE.

Similar to last year, vendors have shown interest in using PCEP session to deploy SR-TE paths. This is mainly due to the shift in the market towards Segment Routing.

PCE-initiated Paths in a Stateful PCE Model

In an MPLS environment, a path between two nodes is called Labeled Switched Path (LSP). In some applications, it's important to be able to create or delete an LSP dynamically as a response to any change in the environment. In a case of a network failure that damages a certain LSP, it's crucial to be able to tear down the damaged LSP and create another one.

In this test, we verified that the PCE is capable of creating an SR-TE path in a single IGP domain. Furthermore, the test checks LSP deletion, path recomputation, PCEP session re-verification, and state synchronization.

The test topology included a PCE and three network nodes. One of the nodes acted as a PCC. Initially, ISIS-TE was used to synchronize the TED information. The LSP was set initially to follow a sub-optimal path, i.e., the path with not the lowest IGP cost.

The test started by establishing a PCEP session and verifying it by checking TED information on the network nodes. After an LSP is initiated by the PCE, we checked that the LSP database of the PCC included one LSP. We then started an L3VPN service to ensure that the traffic was flowing through the path we chose. Next step was to verify the synchronization. We asked the PCE vendors to terminate the PCEP session and delete the LSP in its LSP database. We then asked the PCE vendors to re-establish the session. We check the synchronization of the LSPs on the PCC with the PCE. As expected the traffic was not interrupted. After the PCE deleted the LSP and terminated the PCEP session, the IGP path with the lowest cost was expected to be followed. This was tested by checking the traffic on this path.

Some test case combinations faced initial issues due to some signaling message incompatibility caused by the different implementation of RFC8231. Other issues appeared caused by some vendors being incapable of handling certain SRGB base. This lead to the inability to create an LSP by the PCE. A certain vendor didn't support PCE initiated SR-TE paths. Some vendors faced issues when PCEP messages included unexpected BW objects.

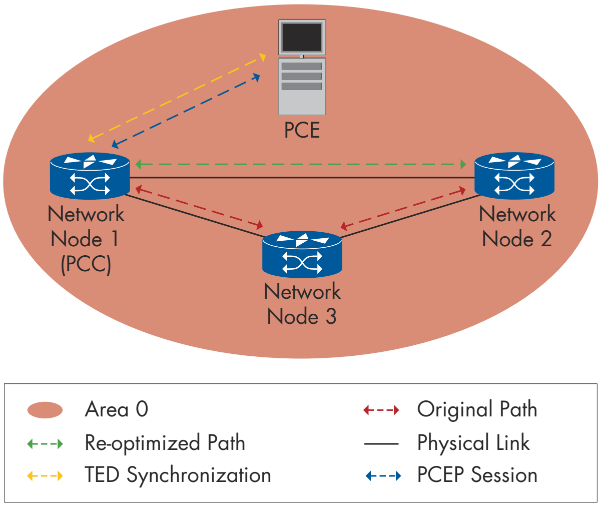

Figure 31: PCE-initiated Paths in a Stateful PCE Model

PCC-initiated Paths in a Stateful PCE Model

In some useful scenarios, the PCC can request an LSP from the PCE in case a PCEP session is established between them. 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.

In this test, we verified that the PCC is capable of requesting an SR-TE path in a single IGP domain. Furthermore, the test checked LSP delegation, LSP re-delegation path re-computation and LSP revocation.

As in the previous case, the test topology included a PCE and three network nodes. One of the nodes acted as a PCC. Initially, ISIS-TE was used to synchronize the TED information. The constraints LSP where set initially to follow a sub-optimal path, i.e., the path with not the lowest IGP cost.

The test started by establishing a PCEP session and verifying it by checking TED information on the network nodes. After an LSP was initiated by the PCC, we checked that the LSP database of the PCC included one LSP. We then started an L3VPN service to ensure that the traffic was flowing through the path we chose. In the next step, the PCC delegated the LSP to the PCE. Finally, we issued the termination of the LSP. LSP database was checked to verify that the LSP was deleted. As expected, data was following the IGP shortest path.

PCE	PCC	Network Node 2	Network Node 3
Cisco IOS XRv9000	Nokia 7750 SR-7	Cisco ASR 9000	HUAWEI NE40E-X8A
Cisco IOS XRv9000	Ericsson 6675	Juniper Networks MX204	ECI Neptune 1050
Cisco IOS XRv9000	Juniper Networks MX204	Cisco ASR 9000	Nokia 7750 SR-7
Juniper NorthStar	Ericsson 6675	Nokia 7750 SR-7	Juniper Networks MX204
Juniper NorthStar	Nokia 7750 SR-7	Juniper Networks MX204	Cisco ASR 9000
Juniper NorthStar	Cisco ASR 9000	Juniper Networks MX204	Nokia 7750 SR-7
Nokia Network Services Platform	Cisco ASR 9000	Nokia 7750 SR-7	HUAWEI NE40E-X8A
Nokia Network Services Platform	Juniper Networks MX204	Nokia 7750 SR-7	Cisco ASR 9000
Nokia Network Services Platform	ECI Neptune 1050	ECI Neptune 1300	Nokia 7750 SR-7
ZTE Corporation ZENIC ONE	Nokia 7750 SR-7	Cisco ASR 9000	Juniper Networks MX204
ZTE Corporation ZENIC ONE	Cisco ASR 9000	Nokia 7750 SR-7	Juniper Networks MX204
ZTE Corporation ZENIC ONE	Ericsson 6675	Nokia 7750 SR-7	Juniper Networks MX204

Table 11: PCE-initiated Paths in a Stateful PCE Model - Successful Combinations

Some vendor combinations faced initial issues due to some signaling message incompatibility caused by the different implementation of RFC8231. Section 5.8.2 specifies two LSP operation states which are active and passive states. Due to different implementations of the states, the communication between the PCE and PCC ended with an error message. Also, some PCCs delegated the LSP to the PCE with no option to revoke the delegation.

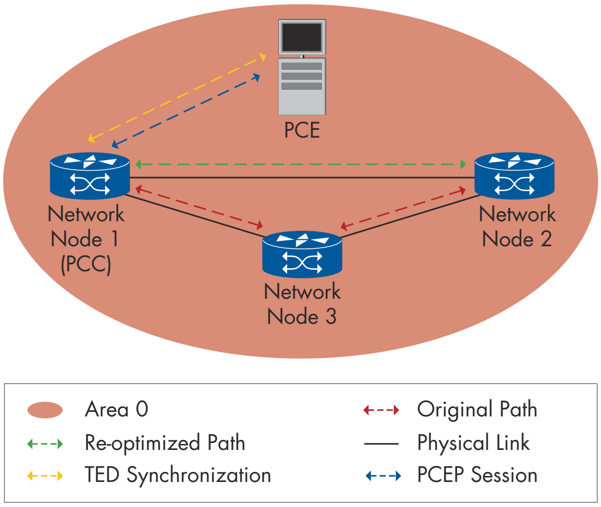

Figure 32: PCC-initiated Paths in a Stateful PCE Model

Path Re-optimization in a PCEP Network

Nowadays networks are dynamic. The centralized SDN architecture should allow the controller to make service updates as a response to a variety of network changes.

In this test, we verify that the PCE can trigger the recalculation and re-optimization of the transport path as a response to a network change. Since there are many events that can trigger a change on the network, we limited this event to one of the following two options:

- Tear down the primary link
- Increase the cost on the primary link

A simple topology for this case includes a PCE and three network nodes. One of the nodes acts as a PCC. Initially, ISIS-TE was used to synchronize the TED information. The constraints LSP was set initially to follow a sub-optimal path, i.e., the path with not the lowest IGP cost.

The initial steps of establishing a PCEP session and assigning an LSP are the same as the previous two scenarios PCE-initiated Paths in a Stateful PCE Model and PCC-initiated Paths in a Stateful PCE Model depending on the vendors' choice. When one of the events trigger a change on the network, LSP is being updated to show the new path and traffic is forwarded without loss over the new path.

Vendors had also the choice to run this case with:

- PCE initiated scenario
- PCC initiated scenario

PCE	PCC	Network Node 2	Network Node 3
Cisco IOS XRv9000	Nokia 7750 SR-7	Cisco ASR 9000	Juniper Networks MX104
Cisco IOS XRv9000	HUAWEI NE40E-M2K	Cisco ASR 9000	Juniper Networks MX204
HUAWEI Network Cloud Engine (NCE)	Nokia 7750 SR-7	HUAWEI NE40E-M2K	Juniper Networks MX104
HUAWEI Network Cloud Engine (NCE)	Cisco ASR 9000	HUAWEI NE40E-M2K	Juniper Networks MX204
Nokia Network Services Platform	Cisco ASR 9000	Nokia 7750 SR-7	Ericsson 6675
Nokia Network Services Platform	HUAWEI NE40E-M2K	Juniper Networks MX204	Nokia 7750 SR-7
Nokia Network Services Platform	HUAWEI ATN910C-F	Nokia 7750 SR-7	Juniper Networks MX204
ZTE Corporation ZENIC ONE	Cisco ASR 9000	Juniper Networks MX104	Nokia 7750 SR-7

Table 12: PCC-initiated Paths in a Stateful PCE Model - Successful Combinations

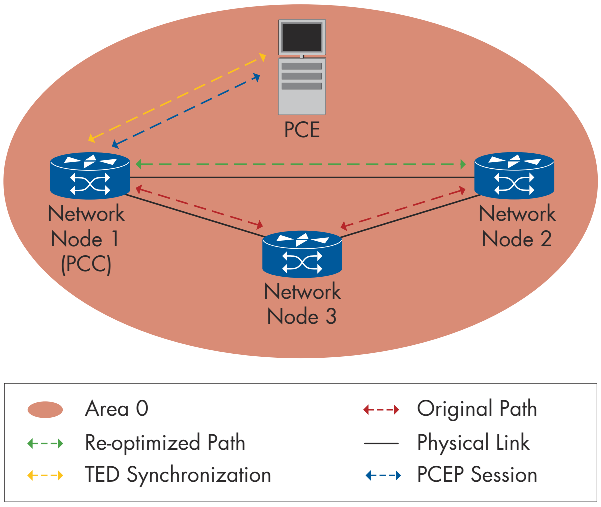

Figure 33: Path Re-optimization in a PCEP Network

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.

PCE	PCC	Network Node 2	Network Node 3
Cisco IOS XRv9000	Ericsson 6675	Juniper Networks MX204	ECI Neptune 1050
Cisco IOS XRv9000	HUAWEI NE40E-M2K	Cisco ASR 9000	Juniper Networks MX204
Cisco IOS XRv9000	Juniper Networks MX204	Cisco ASR 9000	Nokia 7750 SR-7
Cisco IOS XRv9000	Nokia 7750 SR-7	Cisco ASR 9000	Juniper Networks MX204
HUAWEI Network Cloud Engine (NCE)	Nokia 7750 SR-7	Juniper Networks MX204	HUAWEI NE40E-M2K
HUAWEI Network Cloud Engine (NCE)	Cisco ASR 9000	HUAWEI NE40E-M2K	Juniper Networks MX204
HUAWEI Network Cloud Engine (NCE)	Nokia 7750 SR-7	HUAWEI NE40E-M2K	Juniper Networks MX204
Juniper NorthStar	Ericsson 6675	Nokia 7750 SR-7	Juniper Networks MX204
Juniper NorthStar	Nokia 7750 SR-7	Juniper Networks MX204	Cisco ASR 9000
Juniper NorthStar	Cisco ASR 9000	Juniper Networks MX204	Nokia 7750 SR-7
Nokia Network Services Platform	HUAWEI ATN910C-F	Nokia 7750 SR-7	Juniper Networks MX204
Nokia Network Services Platform	Juniper Networks MX204	Nokia 7750 SR-7	Cisco ASR 9000
Nokia Network Services Platform	Cisco ASR 9000	Nokia 7750 SR-7	Juniper Networks MX204
ZTE Corporation ZENIC ONE	Nokia 7750 SR-7	Cisco ASR 9000	Juniper Networks MX204
ZTE Corporation ZENIC ONE	Cisco ASR 9000	Nokia 7750 SR-7	Juniper Networks MX204

Table 13: Path Re-optimization in a PCEP Network - Successful Combinations

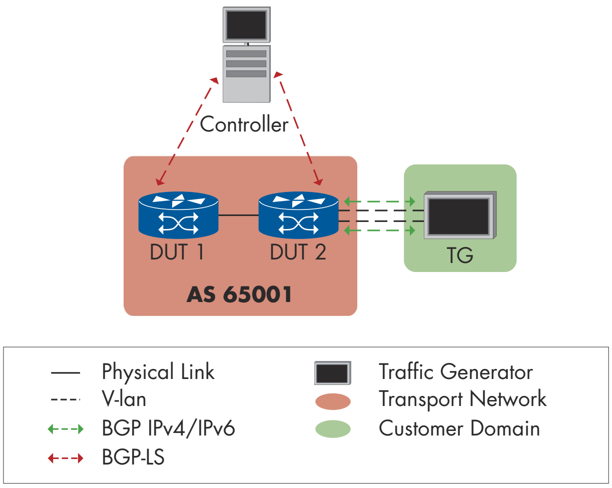

Figure 34: Egress Peer Engineering with Segment Routing

In this test, we verified that EPE Segment Routing could be used to allocate MPLS label for each engineered peer and use a label stack to steer traffic to a specific destination. The test topology includes an SDN controller and two network nodes. Initially, we enabled BGP on all network nodes. Then we enabled BGP-LS between the controller and the network nodes. After that, controller vendors used either PCEP or BGP to advertise SRTE path to the first network node.

We checked that the BGP-EPE controller collected the internal topology and maintained an accurate description of the egress topology of both network nodes. The initial path was set to go through DUT1-DUT2-vlan1. We tested that by running traffic. We configured the controller to push the policy for test traffic to select vlan2. We then verified that traffic followed DUT1-DUT2-vlan2 path.

Flowspec for IPv4/IPv6

BGP Flowspec defines a new BGP Network Layer Reachability Information (NLRI) encoding format that can be used to distribute traffic flow specifications. This allows the routing system to propagate information regarding more specific components of the traffic aggregate defined by an IP destination prefix.

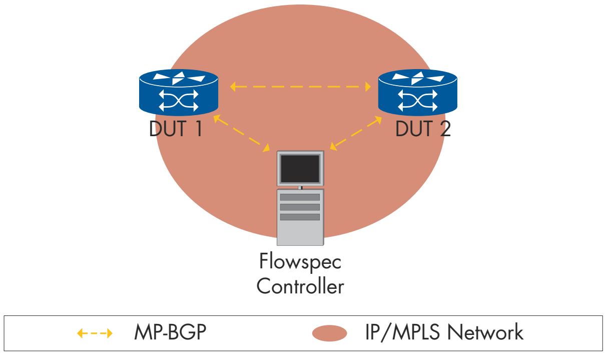

Figure 35: Flowspec for IPv4/IPv6

Flowspec for IPv4/IPv6 proposes a subset of new encoding formats to enable Dissemination of Flow Specification Rules in RFC5575.

The test topology includes a Flowspec controller and two network nodes. The test starts with establishing a BGP session between the two network nodes. After that, bidirectional traffic was generated between the two nodes. After verifying the traffic, BGP sessions were established between the Flowspec controller and the two network nodes. The Flowspec controller sends Flowspec policies to cap the data rate to 128 kbit/sec. Finally, we can see the data rate limit is applied.

Vendors had the choice to choose between IPv4 or IPv6 traffic. Furthermore, Arista tested drop profile with redirect-to-VRF action for policy-based forwarding in the second combination presented in the table below.

Flowspec Controller	DUT 1	DUT 2
Keysight (Ixia) IxNetwork	Nokia 7750 SR-7	HUAWEI NE40E-X8A
Keysight (Ixia) IxNetwork	Arista 7280SR	Nokia 7750 SR-7

Table 14: Flowspec for IPv4/IPv6 - Successful Combinations

Controller	DUT 1	DUT 2	Traffic Generator
Cisco IOS XRv9000	HUAWEI NE40E-F1A	Cisco Nexus 9300-FX	Spirent TestCenter
Cisco IOS XRv9000	Nokia 7750 SR-7	Cisco Nexus 9300-FX	Keysight (Ixia) IxNetwork
Cisco IOS XRv9000	Arista 7280SR	HUAWEI NE40E-F1A	Spirent TestCenter
Keysight (Ixia) IxNetwork	Arista 7280SR	Cisco Nexus 9300-FX	Keysight (Ixia) IxNetwork
Keysight (Ixia) IxNetwork	Nokia 7750 SR-7	Cisco Nexus 9300-FX	Keysight (Ixia) IxNetwork

Table 15: Egress Peer Engineering with Segment Routing - Successful Combinations

BGP-signaled Segment Routing Policies

An SR policy is a set of candidate paths consisting of one or more segment lists. The headend of an SR Policy may learn candidate paths via some different mechanisms, e.g., CLI, NetConf, PCEP, or BGP. In this test case, we verified how BGP could be used to add, update and remove candidate paths of an SR policy. A new BGP SAFI with a new NLRI alongside new sub-TLVs for the Tunnel Encapsulation Attribute was defined to signal an SR policy candidate path. A Cisco route reflector was used for all the cases below.

Three different topologies were tested. The first included an SDN Controller and four network nodes where DUT 1 acts as the ingress node and DUT 4 as the egress node. The second included an SDN controller and three network nodes where DUT 1 acted as the ingress node and DUT 3 as the egress node. Finally, a third topology included an SDN controller and four network nodes where DUT 1 acted as the ingress node and DUT 2 as the egress node.

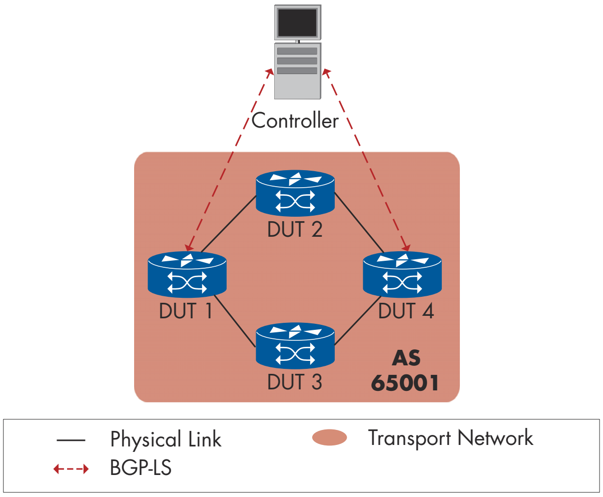

Figure 36: BGP-signaled Segment Routing Policies - Scenario 1

Two different SR policies were tested which are:

- IPv4 EP
- IPv6 EP

Four different traffic types were tested:

- Unlabeled IPv4
- Unlabeled IPv6
- Labeled (BSID+IPv4)
- Labeled (BSID+IPv6)

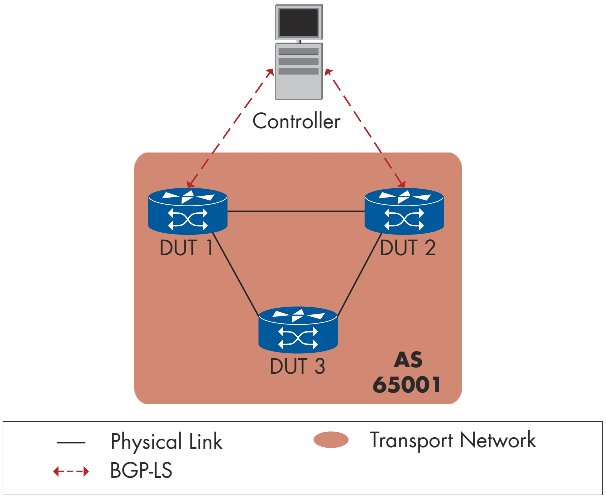

Figure 37: BGP-signaled Segment Routing Policies - Scenario 2

Controller	DUT 1	DUT 2	DUT 3
Keysight (Ixia) IxNetwork	Cisco NCS 5500	Arista 7280SR	HUAWEI NE40E-X8A
Spirent TestCenter	Cisco NCS 5500	Arista 7280SR	HUAWEI NE40E-X8A

Table 16: BGP-signaled Segment Routing Policies - Successful Combinations - Scenario 2

Controller	DUT 1	DUT 2	DUT 3	DUT 4
Cisco IOS XRv9000	Arista 7280SR	Cisco NCS 5500	Nokia 7750 SR-7	HUAWEI NE40E-X8A
Cisco IOS XRv9000	Cisco NCS 5500	Cisco ASR 9000	Arista 7280SR	Nokia 7750 SR-7
Cisco IOS XRv9000	Arista 7280SR	Cisco ASR 9000	Cisco NCS 5500	Nokia 7750 SR-7
Keysight (Ixia) IxNetwork	Arista 7280SR	Cisco NCS 5500	Nokia 7750 SR-7	HUAWEI NE40E-X8A
Nokia Network Services Platform	Arista 7280SR	Cisco NCS 5500	Cisco ASR9901	Nokia 7750 SR-7
Spirent TestCenter	HUAWEI NE40E-F1A	Nokia 7750 SR-7	Cisco ASR 9000	HUAWEI NE40E-M2K

Table 17: BGP-signaled Segment Routing Policies - Successful Combinations - Scenario 1

Furthermore, Cisco and Arista tested automated steering of IP traffic into SR policies based on “color-only” matching.

We first verified that the egress node could be configured to add a BGP color extended community to the prefixes learned from a traffic generator. We then verified that the controller can advertise an SR policy to the ingress node.

We ran one of the mentioned traffic types to prove that traffic flow through the path specified by the policy. Finally, the controller was asked to withdraw the SR policy advertised previously.

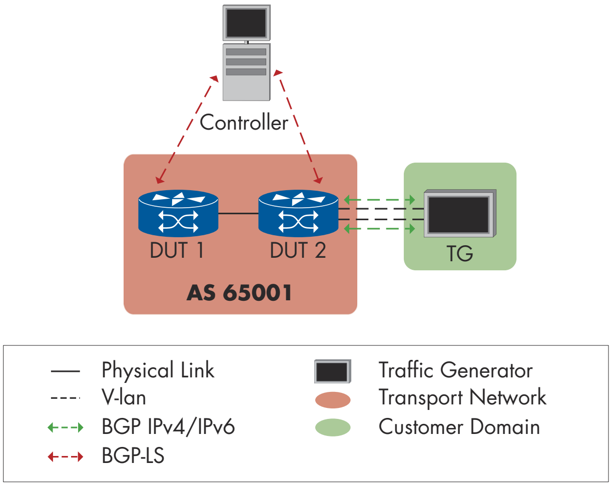

Figure 38: BGP-signaled Segment Routing Policies - Scenario 3

Controller	DUT 1	DUT 2
Cisco IOS XRv9000	Nokia 7750 SR-7	Cisco Nexus 9300-FX
Cisco IOS XRv9000	HUAWEI NE40E-F1A	Cisco Nexus 9300-FX
HUAWEI Network Cloud Engine (NCE)	Nokia 7750 SR-7	Cisco ASR 9000

Table 18: BGP-signaled Segment Routing Policies - Successful Combinations - Scenario 3

Multi-domain Segment Routing Traffic Engineering

An inter-AS TE LSP is an LSP that is using at least two Autonomous Systems (AS) in the path. Topology visibility remains local to a given AS and a head-end LSR cannot compute an inter-AS shortest constrained path. One key application of the PCE based architecture is the computation of inter-AS TE LSP.

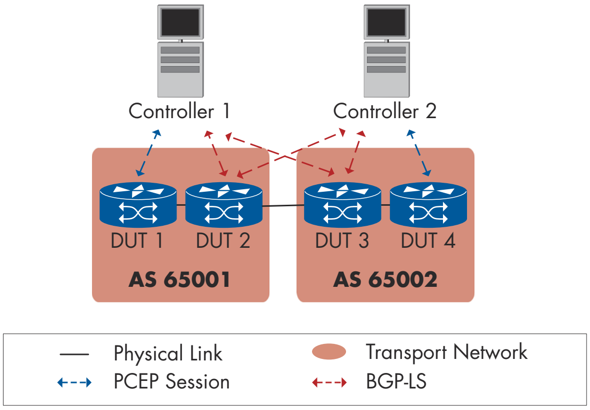

Figure 39: Multi-domain Segment Routing Traffic Engineering - Scenario 1

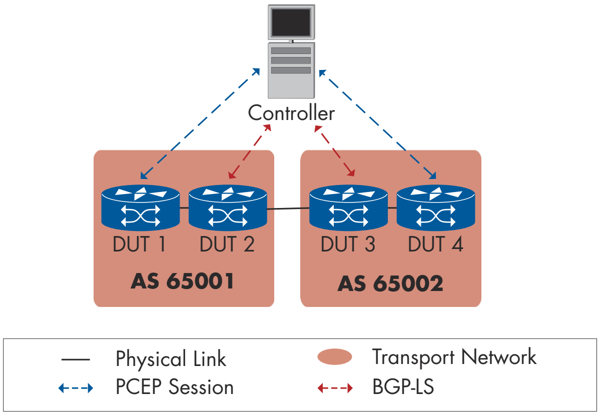

Figure 40: Multi-domain Segment Routing Traffic Engineering - Scenario 2

In this test case, we verified the scenario where we used of a Path Computation Element (PCE) to compute such inter-AS TE LSPs across a predetermined sequence of domains, using a backward-recursive path computation technique.

We also verified the traffic steering capabilities in an inter-domain scenario by pushing a Segment Routing routing policy through PCEP to the ingress PE’s FIB.

The test topology includes two controllers, one for each AS. Each AS contains two network nodes. The two AS are connected via an EPE link. We started the test by establishing PCEP sessions with the ingress nodes. We checked that the controllers had retrieved the TED information related to each AS domain. BGP-LS was established with network nodes in both AS domain. We triggered an L3VPN service between DUT 1 and DUT 4. We checked that the controllers have computed the inter-AS shortest path and sent it to the network nodes. Two transport paths were installed. The path originated at DUT 1 and terminated at DUT 4 and vice versa. We made sure that no traffic loss was seen.

Controller 1	Controller 2	DUT 1	DUT 2	DUT 3	DUT 4
Cisco IOS XRv9000	HUAWEI Network Cloud Engine (NCE)	Cisco ASR 9000	Cisco NCS540	HUAWEI ATN910C-F	HUAWEI NE40E-M2K

Table 19: Multi-domain Segment Routing Traffic Engineering - Successful Combinations - Scenario 1

Controller	DUT 1	DUT 2	DUT 3	DUT 4
Nokia Network Services Platform	Cisco ASR 9000	Cisco NCS540	HUAWEI ATN910C-F	HUAWEI NE40E-M2K

Table 20: Multi-domain Segment Routing Traffic Engineering - Successful Combinations - Scenario 2

NETCONF/YANG

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

In this test section, we provided a combination of NETCONF protocol and different YANG modules. Our main focus was to test L2VPN and L3VPN 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 multi-domain network regardless of the control plane technology of each vendor. In this test, we provisioned end-to-end service using multi-domain's controller. NETCONF was used as a management protocol between domain controllers and between the controllers and the Orchestrator.

The test topology included one Multi Domain Orchestrator, two controllers one for each domain. Each domain has two network nodes.

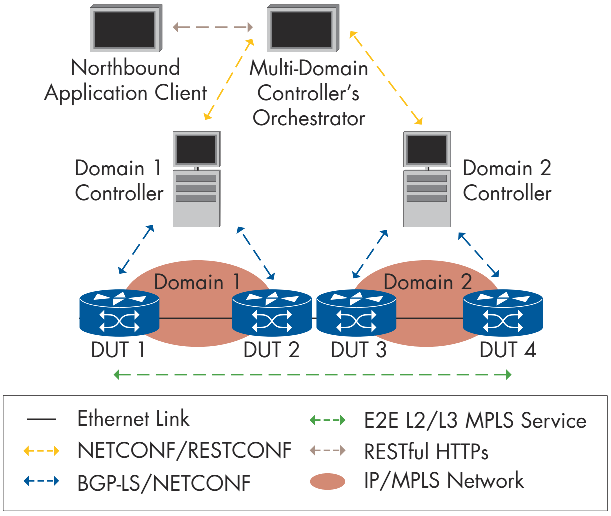

Figure 41: Multi-Vendor/Multi-Domain Controllers Orchestration

First, we verified the NETCONF session between controller and DUTs. Next, we verified the NETCONF session between the Orchestrator and the Domain Controllers. We performed end-to-end service provision using the multi-controller's orchestrator. No traffic loss was observed. Finally, we asked the Orchestrator vendor to delete the provisioned service. We verified that no traffic was flowing.

Multi-Domain Controller's Orchestrator	Domain 1 Controller	Domain 2 Controller	DUT 1	DUT 2	DUT 3	DUT 4
Cisco Network Services Orchestrator (NSO)	Cisco Network Services Orchestrator (NSO)	Cisco Network Services Orchestrator (NSO)	Cisco NCS 5500	Ericsson 6672	Cisco ASR 9000	Ericsson 6672

Table 21: Multi-Vendor/Multi-Domain Controllers Orchestration - Successful Combinations

L2VPN Service Creation Using NETCONF/YANG

YANG is a data modelling 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 a YANG model can be used to configure and manage L2VPNs. It verified VPLS 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 a L2VPN services.

The test topology included two provider edge nodes, a controller, and an orchestrator. NETCONF management protocol was used between the controller, the orchestrator and the provider edges.

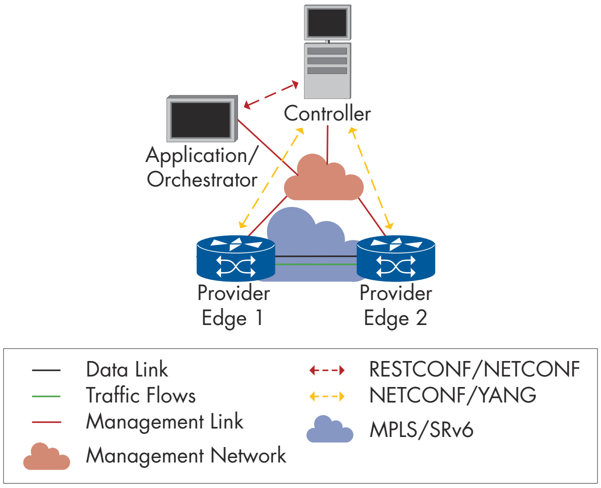

Figure 42: L2VPN Service Creation Using NETCONF/YANG

First, we verified that the complete configuration from the device was retrieved and synchronized with controller configuration data base. We then asked the orchestrator to initiate a L2VPN service and we verified that the status of the service is up on both the controller and the orchestrator. We checked that no traffic loss was seen. We asked the orchestrator to delete the previously configured service and we verified that the current configuration is identical to the initial configuration. Finally we confirmed that none of the traffic was forwarded over the MPLS network.

L3VPN Service Creation Using NETCONF/YANG

This case is very similar to the previous case. The only difference was that the orchestrator and the controller must provision a L3VPN service. This included initiating the service, provisioning it and then delete the service while verifying that the required traffic was flowing as expected.

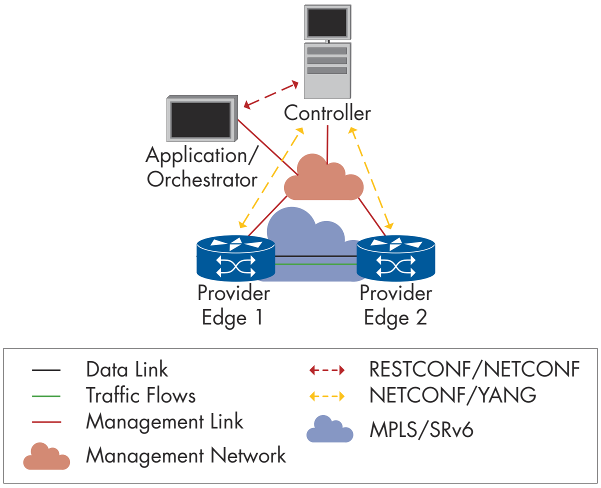

Figure 43: L3VPN Service Creation Using NETCONF/YANG

Controller	Application/ Orchestrator	Provider Edge 1	Provider Edge 2
Cisco Network Services Orchestrator (NSO)	Keysight (Ixia) IxNetwork	Ericsson 6672	Ericsson 6672

Table 22: L2VPN Service Creation Using NETCONF/YANG - Successful Combinations

Controller	Application/ Orchestrator	Provider Edge 1	Provider Edge 2
HUAWEI Network Cloud Engine (NCE)	HUAWEI Network Cloud Engine (NCE)	Cisco NCS 5500	HUAWEI NE40E-M2K
Cisco Network Services Orchestrator (NSO)	Keysight (Ixia) IxNetwork	Cisco ASR 9000	Ericsson 6672

Table 23: L3VPN Service Creation Using NETCONF/YANG - 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 topology includes two DUTs and one NETCONF client. First, we verified that a NETCONF session between a NETCONF client and the NETCONF server is up. We then verified a configuration change against a supported YANG model on the NETCONF server. Finally, we deleted the configurations and terminated the NETCONF session.

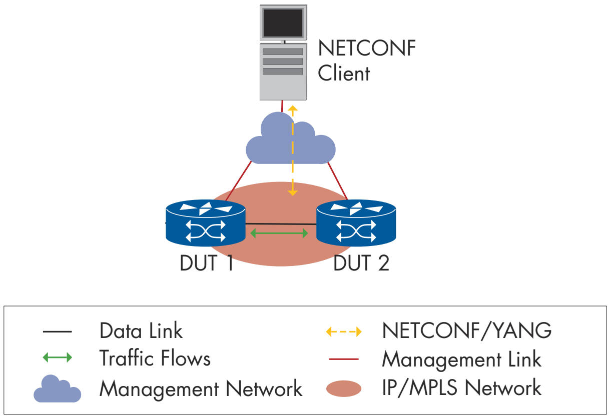

Figure 44: Device Configuration Using NETCONF/YANG

NETCONF Client	DUT 1	DUT 2
HUAWEI Network Cloud Engine (NCE)	HUAWEI NE40E-M2K	Cisco ASR 9000
Keysight (Ixia) IxNetwork	BISDN Basebox & DELTA AG7648	Ericsson 6672
Cisco Network Services Orchestrator (NSO)	Ericsson 6672	Cisco ASR 9000
Cisco Network Services Orchestrator (NSO)	ECI Neptune 1050	Ericsson 6672

Table 24: Device Configuration Using NETCONF/YANG - Successful Combinations

Microwave

As the conversion to 5G is happening now and all the operators are getting ready for it, we tried to answer the question about the role and the capability of microwave transport.

We see a trend of integrating the more specialized microwave devices into the standard IP/MPLS router domain, and thus we looked into two particular aspects of this in the following tests.

Bandwidth Notification

Today, mobile backhaul networks are often built as an overlay with routers setting on top of microwave devices. In the past, there was limited communication between these two domains, but with the bandwidth notification messages (ETH-BN) defined by ITU-T Y.1731, it is now possible for the microwave systems to signal a change in bandwidth to the routers. This enables a router to apply service policies to the traffic it sends on to the microwave system based on the bandwidth information within the ETH-BN packets.

At the beginning of this test, the microwave nodes were using the maximum modulation possible (4096 QAM at 56 MHz), and sent end-to-end traffic.

In the next step, we emulated severe weather conditions in the link between the microwave nodes by using an RF attenuator and verified that the microwave nodes generated and signaled the ETH-BN message to the aggregation router, which subsequently could process the bandwidth notification messages (ETH-BN) and accordingly apply service policies to the traffic sent to the microwave system.

We successfully tested the following combination:

- Ericsson Router 6371 acted as the aggregation router. The microwave link was established between two Intracom Telecom OmniBAS-2W devices.

Layer 3 Microwave MPLS-based Services

This test aimed to confirm the capability to establish IP/MPLS service on a microwave platform crossing or terminating on existing infrastructure.

We tested four different combinations relying on different transport profiles, and verified that an L2VPN VPWS/VPLS service (in first and second combinations)/L3VPN service (in first, third, and fourth combinations) can be set up between IP/MPLS capable microwave systems and IP/MPLS aggregation routers in multi-vendor scenario. In the first and second scenario, we used OSPF as the IGP protocol and LDP for the MPLS label allocation/distribution. In the third scenario, we changed the IGP to IS-IS with LDP.

We created both end-to-end services between the microwave system operating at maximum modulation (4096 QAM at 56 MHz) and the standalone routers participating as aggregation router as well as directly between two microwave vendors.

In the tests we used the following combinations:

- Microwave System: Intracom Telecom OmniBAS-2W, Aggregation Router: Juniper Networks MX104
- Microwave System: Intracom Telecom OmniBAS-2W, Aggregation Router: Ericsson 6371
- Microwave System: Intracom Telecom OmniBAS-2W, Aggregation Router: Juniper Networks MX104, Microwave System: Ericsson 6693, Ericsson MINI-LINK 6691
- Microwave System: Intracom Telecom OmniBAS-2W. Microwave System: Ericsson 6693, Ericsson MINI-LINK 6691

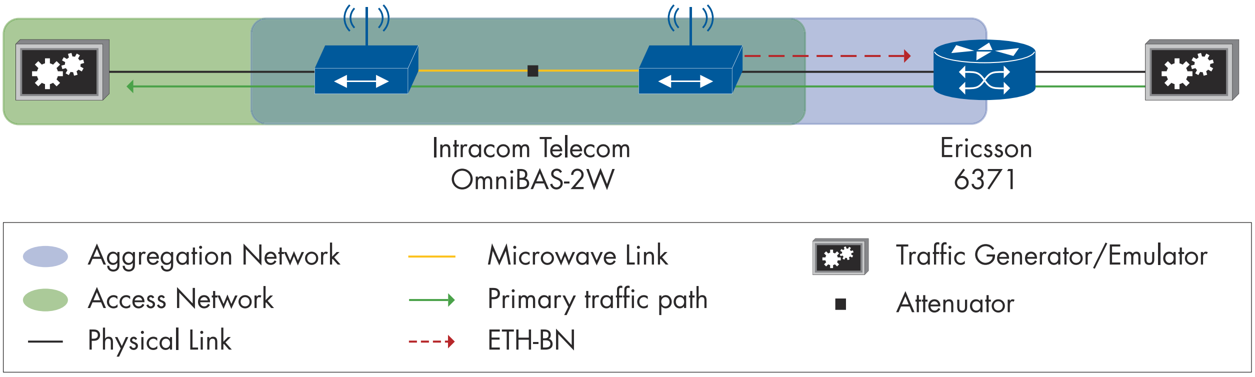

Figure 45: Bandwidth Notification

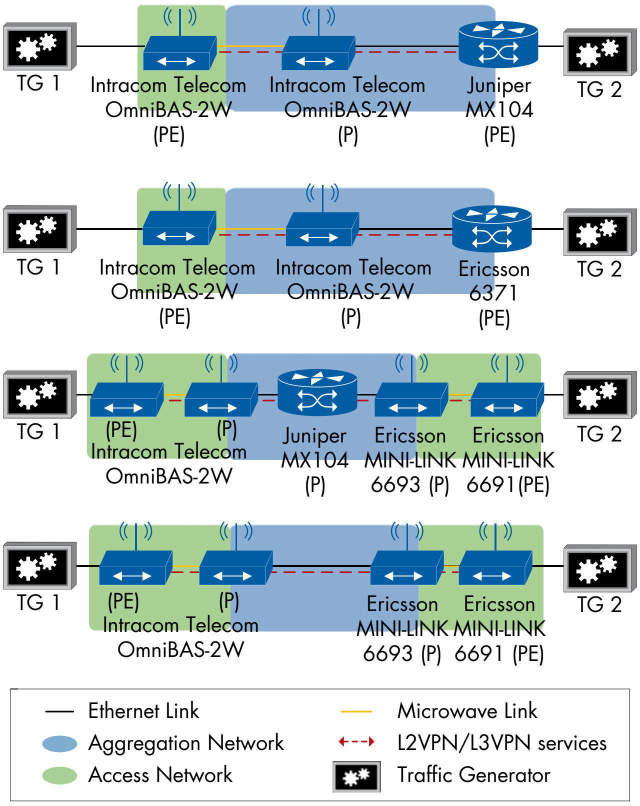
TG 1
Intracom Telecom OmniBAS-2W (PE)
Intracom Telecom OmniBAS-2W (P)
Juniper MX104 (PE)
TG 2
TG 1
Intracom Telecom OmniBAS-2W (PE)
Intracom Telecom OmniBAS-2W (P)
Ericsson 6371 (PE)
TG 2
TG 1
(PE) (P) Intracom Telecom OmniBAS-2W
Juniper MX104 (P)
Ericsson MINI-LINK 6693 (P)
Ericsson MINI-LINK 6691(PE)
TG 2
TG 1
(PE) (P) Intracom Telecom OmniBAS-2W
Ericsson MINI-LINK 6693 (P)
Ericsson MINI-LINK 6691 (PE)
TG 2
Ethernet Link
Aggregation Network
Access Network
Microwave Link
L2VPN/L3VPN services
Traffic Generator

Figure 46: Layer 3 Microwave MPLS-based Services

Layer 3 Microwave Transport Resiliency

Bringing IP/MPLS to the access provides additional resiliency options in the access network and increases the end-to-end service availability.

The goal of the test was to verify interoperability on IP/MPLS transport and service control plane and data plane between IP/MPLS capable microwave systems and/or IP/MPLS core routers in a multi-vendor scenario.

Additionally, in this test, we measured service interruption time for multiple profiles on the MPLS control plane. To emulate severe weather conditions, we reduced the bandwidth between the nodes of the microwave system using an RF attenuator.

We used Spirent TestCenter to act as CE and sent bidirectional traffic across the network. We verified that the microwave nodes were using the main path with the maximum modulation scheme available (4096 QAM at 56 MHz) and that no packets were lost.

We then emulated severe weather conditions by reducing the available bandwidth of the channel.

Two L2VPN VPWS services were established (between Intracom Telecom PE microwave nodes and Juniper PE aggregation router). On link failure, traffic in first VPWS was switched to the backup traffic path, while traffic in second VPWS remained unaffected.

In this test, three Intracom Telecom OmniBAS-2W microwave devices were used, two acted as PE nodes and one as P node. Juniper Networks MX104 acted as a PE node.

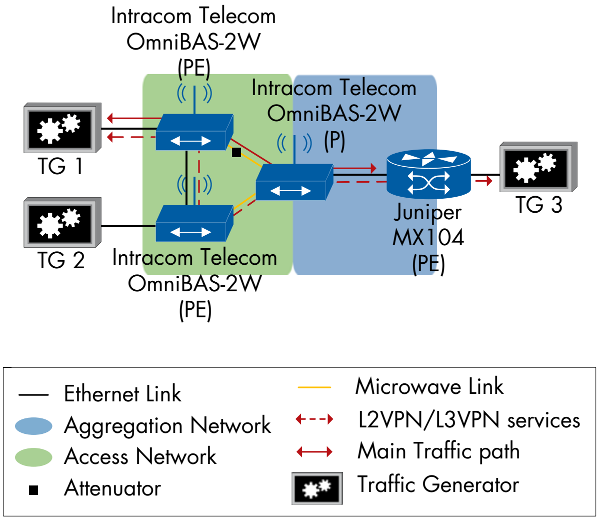

Figure 47: Layer 3 Microwave Transport Resiliency

Clock Synchronization

In this year's event, we focused on time/phase delivery trying to achieve the requirements for 5G networks, including a lot of resiliency scenarios, using full and assisted partial timing support setups.

We tested the behavior of the time signal delivery in optimal and suboptimal conditions: network delay asymmetry, hold-over performances, source failover between two Grandmaster Clocks with high-precision clocking and we reached 45 successful combinations.

For the high-precision clocking 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.

Again the Calnex Paragon suite of products proved invaluable in both generating the network impairment characteristics (G.8261 Test case 12), in providing accurate measurement, in reporting against the 5G network limits and clock mask performance.

The primary reference time clock (PRTC) was using an GNSS 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. The R&D departments of the vendors reacted quickly with providing patches and troubleshooting support.

Phase/Time Partial Timing Support

This test was performed using only the ITU-T G.8275.2 profile (PTP telecom profile for Phase/Time-of-day synchronization with partial timing support from the network), without any physical frequency reference – such as SyncE.

In this setup, the Grandmaster Clock was provided with GPS input, while the slave and Boundary Clock started from a free running condition.

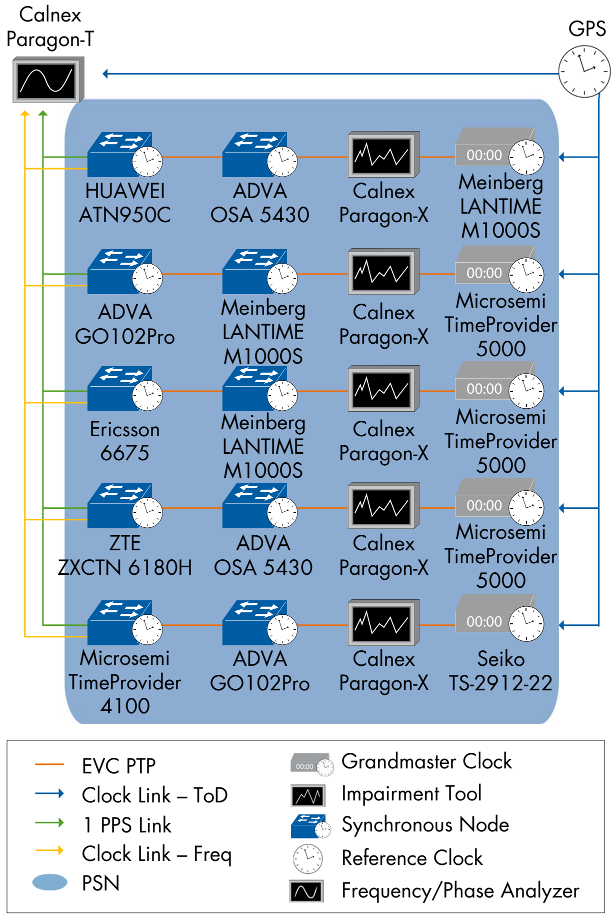

Figure 48: Phase/Time Partial Timing Support

We performed the following combinations with Boundary Clocks:

Grand-master	Boundary Clock	Slave Clock
Meinberg LANTIME M1000S	ADVA OSA 5430	HUAWEI ATN950C
Microsemi TimeProvider 5000	Meinberg LANTIME M1000S	ADVA GO102Pro
Microsemi TimeProvider 5000	Meinberg LANTIME M1000S	Ericsson 6675
Microsemi TimeProvider 5000	ADVA OSA 5430	ZTE Corporation ZXCTN 6180H
Seiko TS-2912-22	ADVA GO102Pro	Microsemi TimeProvider 4100

Table 25: Successful Combinations

Phase/Time Assisted Partial Timing Support: Delay Asymmetry

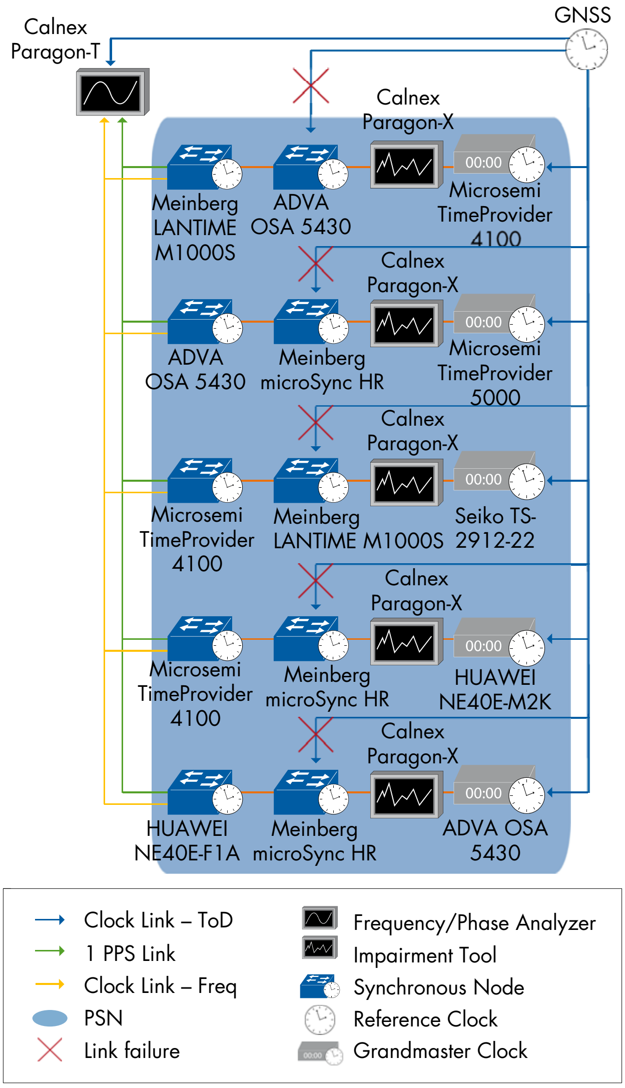

Figure 49: Phase/Time Assisted Partial Timing Support: Delay Asymmetry

This test was performed using the ITU-T G.8275.2 profile between the Grandmaster and Boundary Clock, with the participants having the choice of running G.8275.1 or G.8275.2 between the boundary and Slave Clocks.

After disconnecting the GPS from the Boundary Clock, we used the Calnex Paragon-X to introduce an additional delay asymmetry of 250 μs and verified that the boundary could calculate and compensate the asymmetry introduced.

High-Precision Clocking: Source Failover

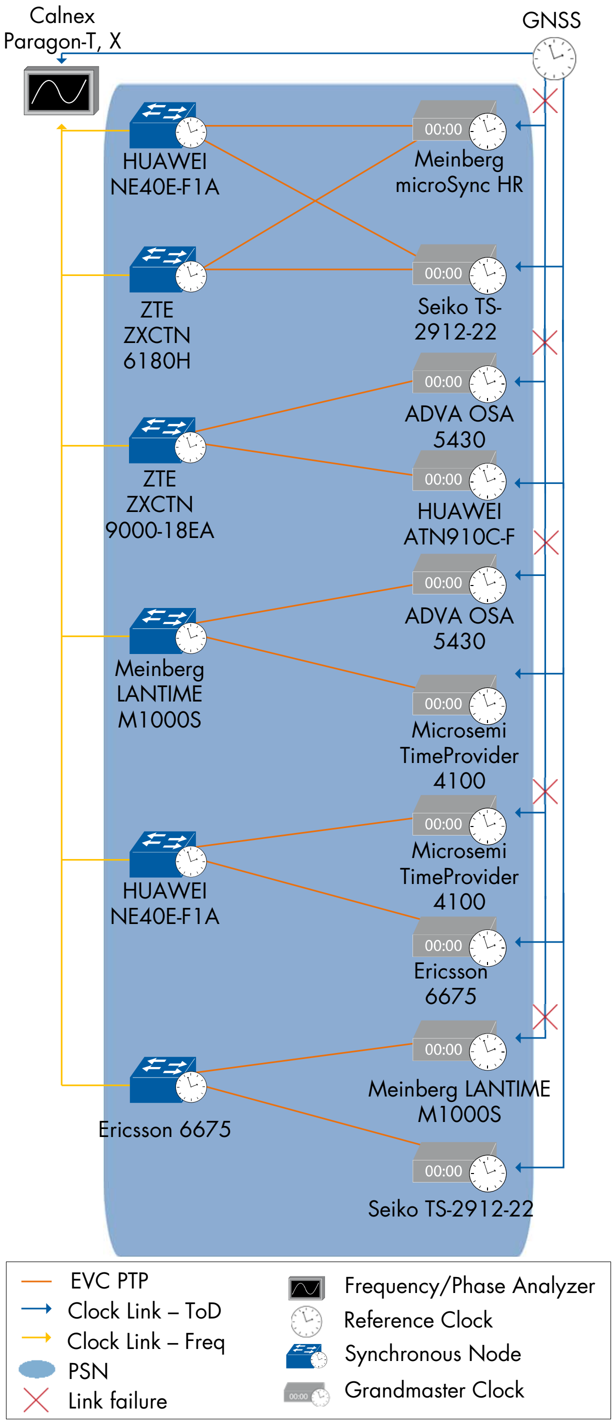

Figure 50: High-Precision Clocking: Source Failover

The goal from this test is to test a resiliency scenario in which both Grandmaster devices were provided with a GPS signal from a common GNSS antenna. We allowed the Boundary Clock to lock to the primary Grandmaster and then degraded the primary Grandmaster's quality by disconnecting its GNSS input. We verified that the Boundary Clock switched over to the secondary Grandmaster and measured the Boundary Clock's transient response. We also tested if the Grandmaster devices are signaling the correct ClockClass values according to the telecom profiles, which allows the alternate best master clock algorithm on the Boundary Clock to correctly select the best Grandmaster during each step of the tests. We used the priority2 field as tie-break parameter.

Grand-master A	Grand-master B	Boundary Clock
ADVA OSA 5430	HUAWEI ATN910C-F	ZTE Corporation ZXCTN 9000-18EA
ADVA OSA 5430	Microsemi TimeProvider 4100	Meinberg LANTIME M1000S
Meinberg microSync HR	Seiko TS-2912-22	HUAWEI NE40E-F1A
Meinberg microSync HR	Seiko TS-2912-22	ZTE Corporation ZXCTN 6180H
Meinberg LANTIME M1000S	Seiko TS-2912-22	Ericsson 6675
Microsemi TimeProvider 4100	Ericsson 6675	HUAWEI NE40E-F1A

Table 26: Successful Combinations

The test was performed using the ITU-T G.8275.1 between the Grandmaster devices and the Boundary Clock.

It is critical to calibrate the Grandmaster devices and to compensate the cable delays between the Grandmaster devices and the GNSS antenna to guarantee that these delays do not affect the test results.

This test is designed to achieve the accuracy requirements ITU-T G.8271 Level 6. The following combinations passed 260 ns (ITU-T G.8271 Level 6A) and some of them passed 130 ns (ITU-T G.8271 Level 6B).

Phase/Time Synchronization: Source Failover

In this setup, we tested a real-life resiliency with two Grandmaster devices, Boundary Clock, and a Slave Clock. The Boundary Clock was locked on the primary Grandmaster, and then we degraded the Grandmaster A quality by disconnecting 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 Grandmaster devices, Boundary Clock and Slave Clock.

It is critical to calibrate the Grandmaster devices and to compensate the cable delays between the Grandmaster devices and the GNSS antenna to guarantee that these delays do not affect the test results.

The goal was to achieve the accuracy of G.8271 accuracy level 4, although some combinations achieved the high-precision clocking ITU-T G.8271 Level 6.

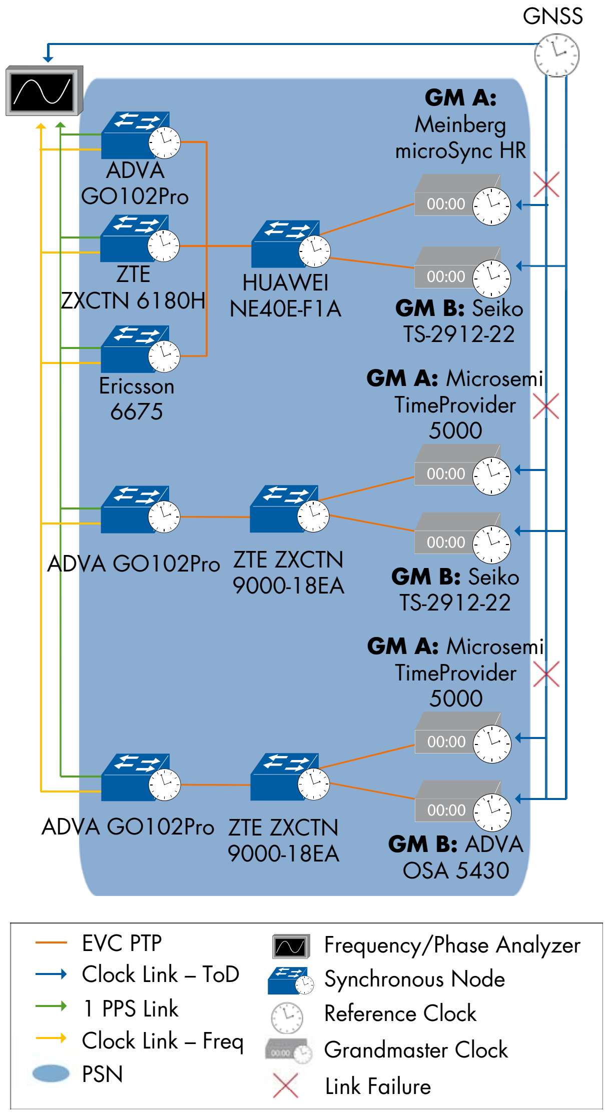

Figure 51: Phase/Time Synchronization: Source Failover - Level 6

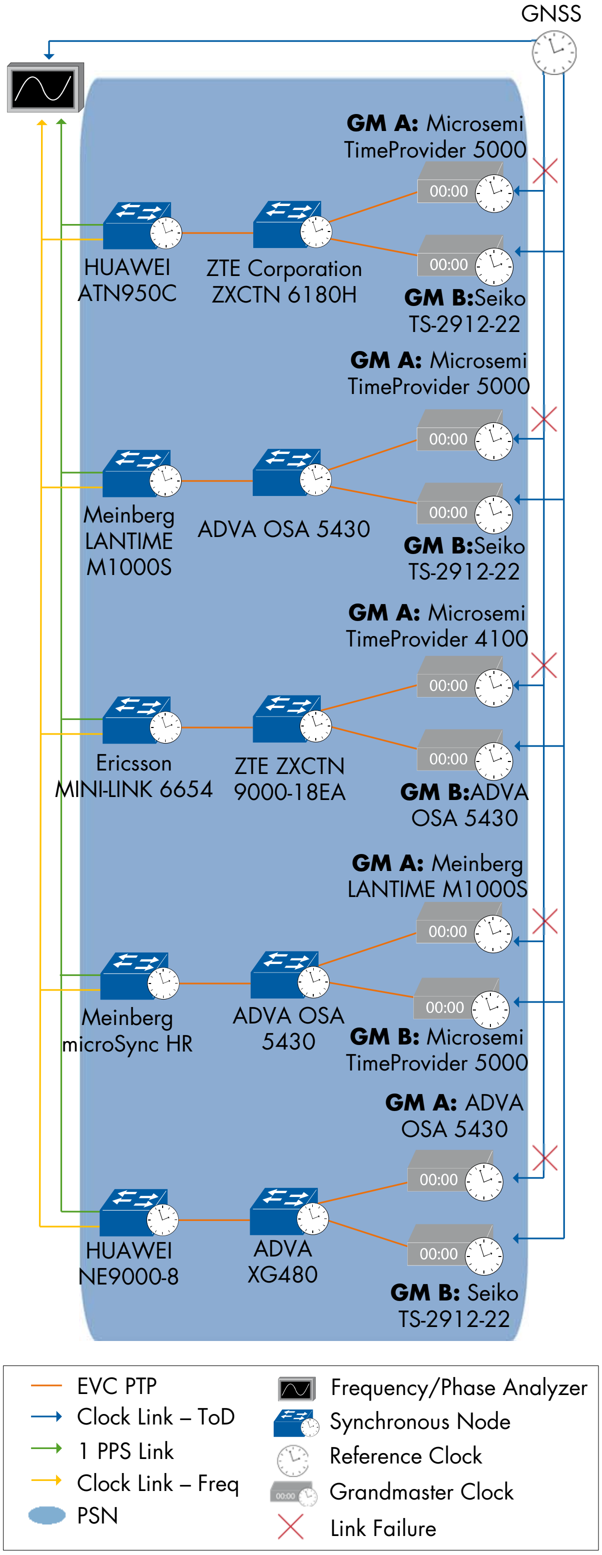

Figure 52: Phase/Time Synchronization: Source Failover - Level 4

The following combinations achieved the ITU-T G.8271 Level 6:

Grandmaster A	Grandmaster B	Boundary Clock	Slave Clock
Microsemi TimeProvider 5000	Seiko TS-2912-22	ZTE Corporation ZXCTN 9000-18EA	ADVA GO102 Pro
Microsemi TimeProvider 5000	ADVA OSA 5430	ZTE Corporation ZXCTN 9000-18EA	ADVA GO102 Pro
Meinberg microSync HR	Seiko TS-2912-22	HUAWEI NE40E-F1A	ADVA GO102 Pro
Meinberg microSync HR	Seiko TS-2912-22	HUAWEI NE40E-F1A	ZTE Corporation ZXCTN 6180H
Meinberg microSync HR	Seiko TS-2912-22	HUAWEI NE40E-F1A	Ericsson 6675

Table 27: Successful Combinations - Level 6

The following combinations achieved the ITU-T G.8271 Level 4:

Grandmaster A	Grandmaster B	Boundary Clock	Slave Clock
Microsemi TimeProvider 5000	Seiko TS-2912-22	ZTE Corporation ZXCTN 6180H	HUAWEI ATN950C
Microsemi TimeProvider 5000	Seiko TS-2912-22	ADVA OSA 5430	Meinberg LANTIME M1000S
Microsemi TimeProvider 4100	ADVA OSA 5430	ZTE ZXCTN 9000-18EA	Ericsson MINI-LINK 6654
Meinberg LANTIME M1000S	Microsemi TimeProvider 5000	ADVA OSA 5430	Meinberg microSync HR
ADVA OSA 5430	Seiko TS-2912-22	ADVA XG480	HUAWEI NE9000-8

Table 28: Successful Combinations - Level 4

Phase/Time Synchronization with Full Timing Support: Microwave Transport

The goal of this test was to verify that a microwave transport in a network providing full timing support maintains the phase accuracy requirements. A microwave system may undergo conditions that cannot be controlled by the network operator, such as severe weather conditions. The goal of this test case is to verify the synchronization functions of IEEE 1588-2008 located at the Grandmaster and Slave Clocks when the Boundary Clock is a microwave system.

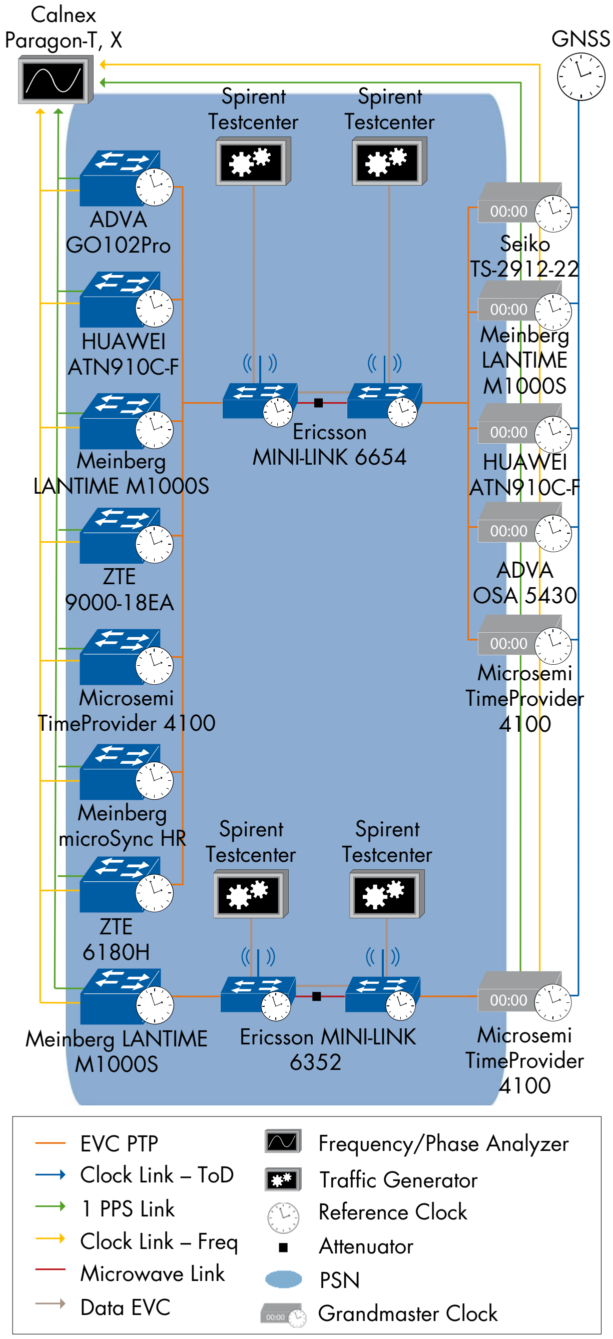

Figure 53: Phase/Time Synchronization with Full Timing Support: Microwave Transport

To emulate severe weather conditions, we reduced the bandwidth between the two nodes of the microwave using an RF attenuator. We then verified that the PTP traffic is prioritized over other data traffic and the Slave Clock output retained the required quality level.

We started the test with the Slave Clock in free running mode and generated a constant bit rate at the maximum line rate for the maximum modulation scheme (10% of 576-byte packets, 30% of 64-byte packets, 60% of 1518 byte packets) and expected no traffic loss. After the Slave Clock locked, we performed baseline measurements using the Calnex Paragon-T device. To emulate severe weather conditions, we reduced the bandwidth between the two nodes of the microwave network using an RF attenuator. As expected the nodes reacted by changing the modulation used. We then verified that the PTP traffic was unaffected by the change of modulation, as it was prioritized over other data traffic and the Slave Clock output retained the required quality level. Since the bandwidth decreased accordingly, we saw that data packets were dropped according to the available bandwidth.

In one combination, the Microwave device was not able to handle a mix between bursts and constant traffic, although it can handle each one separately, so we used the constant traffic for maximum line rate.

All the combinations we tested for this testcase:

Grandmaster	Boundary Clock (Microwave System)	Slave Clock
Seiko TS-2912-22	Ericsson MINI-LINK 6654	ADVA GO102Pro
		HUAWEI ATN910C-F
		Meinberg LANTIME M1000S
Meinberg LANTIME M1000S	Ericsson MINI-LINK 6654	ADVA GO102Pro
		ZTE Corporation ZXCTN 9000-18EA
		Microsemi TimeProvider 4100
HUAWEI ATN910C-F	Ericsson MINI-LINK 6654	ADVA GO102Pro
		ZTE Corporation ZXCTN 9000-18EA
		Microsemi TimeProvider 4100

ADVA OSA 5430	Ericsson MINI-LINK 6654	Meinberg microSync HR
		ADVA GO102Pro
		ZTE Corporation ZXCTN 6180H
Microsemi TimeProvider 4100	Ericsson MINI-LINK 6654	Meinberg microSync HR
		ADVA GO102Pro
		ZTE Corporation ZXCTN 6180H
Seiko TS-2912-22	Ericsson MINI-LINK 6654	ZTE Corporation ZXCTN 9000-18EA
Meinberg LANTIME M1000S	Ericsson MINI-LINK 6352	Microsemi TimeProvider 4100

Table 29: Successful Combinations

Phase/Time Synchronization: Degradation of Primary Source

According to the architecture defined in ITU-T G.8275 a Boundary Clock can become a Grandmaster and can also be slaved to another PTP clock. The goal of this test was to check the capability to swap the role of a Boundary Clock's port from master to slave and vice-versa and also to test the 100 GBE cables in the core network between the Boundary Clocks. This test was performed using the ITU-T G.8275.1 profile. Both the Grandmaster and one of the Boundary Clocks (BC-A) were provided with a GNSS signal. We allowed the Grandmaster and the Boundary Clock A to lock to GPS input. The Boundary Clock A acted as primary Grandmaster for the upstream Boundary Clock (BC-B). We then disconnected the antenna of the Boundary Clock A to emulate a GNSS failure and verified that both Boundary Clocks locked via PTP to the central Grandmaster. In the last step, we recovered the GNSS of the Boundary Clock A and verified that the Boundary Clock B locked again to the downstream Boundary Clock A.

Many devices were not able to act as a Boundary Clock A because they were not able to provide the measurement interfaces while they are using the GNSS as a reference and only a few BCs have 100GbE interfaces.

One combination – Ericsson/ADVA/Huawei – provided a full 100GbE chain from GM to BC-B and from BC-B to BC-A.

One combination was failed because the SyncE was locked via the GNSS reference in the BC A not on BC B, also we faced some physical layer issues during the setup cabling regarding the SFPs or using the Copper cables to carry the frequency between the GM and BC B.

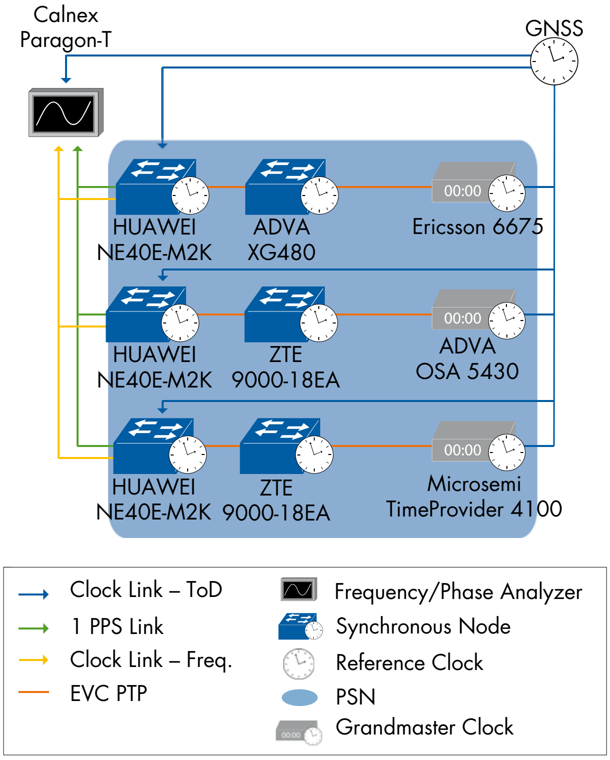

Figure 54: Phase/Time Synchronization: Degradation of Primary Source

Grandmaster	Boundary Clock A	Boundary Clock B
Ericsson 6675	HUAWEI NE40E-M2K	ADVA XG480
ADVA OSA 5430	HUAWEI NE40E-M2K	ZTE Corporation ZXCTN 9000-18EA
Microsemi TimeProvider 4100	HUAWEI NE40E-M2K	ZTE Corporation ZXCTN 9000-18EA

Table 30: Successful Combinations

Summary

The EANTC team thanks all 20 participating vendors for joining us for this fantastic event in Berlin. In the two weeks we were able to complete a wide range of tests with a total of 174 interop combinations. Most combinations worked as expected and some interop issues between different vendor implementations could be seen.

It was a pleasure to meet and work with so many great people and we are looking forward to the next event in 2020!

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

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

While every reasonable effort has been made to ensure accuracy and completeness of this publication,
the authors assume no responsibility for the use of any information contained herein.
All brand names and logos mentioned here are registered trademarks of their respective companies
in the United States and other countries.

20190506 v1.3