Interoperability Showcase 2017

White Paper

EANTC

MPLS + SDN + NFV WORLD → PARIS 2017

Editor's Note

Carsten Rossenhövel
Managing Director, EANTC

Since we started the EANTC multi-vendor interoperability tests in Paris in 2003, Upperside's MPLS + SDN + NFV World Congress has taken place within weeks of the Mobile World Congress in Barcelona. 15 years ago, these two conferences were very different. Mobile operators did not care much about the fixed network aspects (and vice versa). News spread in Barcelona dealt with radio aspects or cell phones, which Paris conference attendees would not care much about.

This year, walking through the halls in Barcelona has been a very different experience. **5G** has brought a whole set of new challenges: Mobile networks need to be ubiquitously available and densely serviced without obstructive bandwidth limits. Customers want to use these networks for many tasks previously assigned to fixed networks. A huge flood of Internet of Things (IOT) devices is expected soon, requiring connectivity to scale. Low-latency services build on Mobile Edge Computing (MEC) which requires distributed data centers in the aggregation and access network.

Each of these areas depends on the innovation of the transport network and related virtualized network services: They need to become more flexible, manageable, and scaling on the service level. Due to the diverse functions involved, service providers will be even more inclined to deploy multi-vendor solutions than before.

Our interoperability test and showcase focused on the innovations in the Software-Defined Network (SDN) areas enabling these services. The marketing hype around SDN has subsided which is a good sign that vendors are now working hard to deliver mature, standardized and interoperable solutions. A convincing level of interoperability has been reached for many SDN aspects (wide area and data center) this year. It has been very encouraging to see that more vendor solutions interoperate on the management and orchestration level as well meanwhile.

At the same time, giving up infrastructure in favor of new solutions (the infamous "forklift upgrade") has mostly been dismissed as a bad idea in the telecoms industry. Integrating existing infrastructure into an improved scenario is vital. This is specifically true for well-working MPLS technology which still (and for many years to come) runs most of the world's packet networks. We have focused the integration of MPLS into SDN specifically on the management and orchestration level. Noticeably, MPLS is embracing microwave transmission components that used to be pure switching devices in the past.

One very important aspect of mobile network innovation cannot be underestimated: Synchronizing clocks with high precision and scale. We have tested many aspects of Precision Time Protocol (PTP) interoperability in the past years; this time we focused on reliability and newly standardized ITU profiles, with rewarding very positive results.

In two weeks of testing, we achieved more results than ever thanks to outstanding, dedicated vendor teams participating with great, innovative products – 32 pages of detailed, first-hand information about the state of MPLS and SDN interoperability are ahead!

Table of Contents

Participants and Devices 3
MPLS, Ethernet & Data Center Interconnect 4
Software Defined Networking 15
Topology 16
Clock Synchronization 23

Introduction

Since 2013, we actively follow the progress of Ethernet VPN (EVPN) standardization and vendors' implementation through our multi-vendor interoperability tests. Despite the positive results over the years and increasing numbers of implementations that support EVPN, we still see the need for intensive EVPN interoperability testing.

This year, in addition to the classical EVPN tests, we verified multiple new EVPN features:

- EVPN route types for the advertisement of IP Prefixes
- Two interworking use cases between the data center network and WAN
- Two new drafts for EVPN support of MEF Ethernet services

We also extended the scope of Segment Routing testing with more advanced Traffic Engineering scenarios. Two new use cases were included: one involving a disjointness enforcement within a so-called dual-plane network, and the other one involving CoS-based policies. Also for the first time, we verified scenarios where the Segment Routing was used as transport for EVPN services.

The microwave vendors showed great interest in advanced transport test cases. Two participants tested adaptive modulation with the capability to propagate the bandwidth information further in the network. We also verified the ability to support IP/MPLS via microwave.

Four vendors tested the interoperability of their SDN controllers with multiple network forwarding devices in the SDN area. The vendors were very keen on NETCONF/YANG and PCEP. During the preparation phase for the tests, some participants also showed interest in advanced SDN use cases such as multi-domain and multi-layer service orchestration. However, in the end there was insufficient

support for multi-vendor interoperability and the tests were not feasible this time.

In the packet clock synchronization area we challenged the participating vendors with new test cases based on the recently realised PTP profile for time/phase synchronization with partial timing support from the network.

We also observed an increased involvement in timing support from microwave vendors with three of them taking part in these test cases.

Participants and Devices

Vendor	Devices
Arista Networks	7280R 7050SX 7050QX
Aviat Networks	CTR 8540 WTM4000 ODU600
Calnex	Calnex Paragon-X Calnex Paragon-T Calnex Rb/GPS Frequency Source Frequency Converter
Cisco	Nexus 9372[a] Nexus 93180 Nexus 7702 Cisco Network Service Orchestrator (NSO) Cisco Aggregation Services Router 9000v (ASR9kv) Cisco Cloud Services Router 1000V (CSR1kv)
ECI	NPT-1800
Ericsson	MINI-LINK 6352 MINI-LINK 6691 Router 6672 eNodeB Baseband T605 Baseband 5212
Huawei	NE40E-X8A NE40E-X3A NE40E-X2-M8A NE40E-X2-M16A CloudOpera IES Agile Controller ATN950C ATN980B ATN910C
Ixia	IxNetwork Traffic Generator NOVUS ONE Application Server
Juniper Networks	MX80 MX104 MX240 NorthStar Controller QFX10002 Virtual MX
Meinberg	LANTIME M4000 LANTIME M1000S
Microsemi	TimeProvider 5000 TimeProvider 2700 TimeProvider 2300 IGM1100
Nokia	Nokia 7750SR Nokia NSP
RAD	ETX MiCLK
SIAE MICRO-ELETTRONICA	SIAE MICROELETTRONICA AGS20_HW version 1 & version 2
Spirent Communications	Spirent TestCenter Spirent TestCenter Virtual

a. Nexus 9372 and 93180 will be called alternately Nexus 9300 Series in the white paper

Interoperability Test Results

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

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

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

MPLS, Ethernet & Data Center Interconnect

E-Line and E-Tree are two standardized end-to-end Ethernet services defined by the MEF (Metro Ethernet Forum). The IETF's (Internet Engineering Task Force) EVPN technology now offers EVPN solutions that meet the requirements of E-Line and E-Tree for data center interconnect. We therefore tested both E-Line and E-Tree services with EVPN in IP/MPLS-based networks.

Next generation data center interconnect can be achieved via various approaches as Software-Defined Networking (SDN) evolves over an IP/MPLS core network. We verified interconnection between network virtualization overlay (NVO) running on EVPN and common IP/MPLS-based technologies used in the WAN, such as IP-VPN, EVPN-MPLS and Segment Routing.

The IETF defined an integrated routing and bridging (IRB) draft to enable inter-subnet forwarding between tenants across different IP subnets in EVPN. We verified two IP-VRF-to-IP-VRF scenarios based on the new Route type 5 (IP Prefix route) as defined in the IP Prefix Advertisement in the EVPN draft: Interface-less model and Interface-full with a core-facing IRB model.

EVPN-VXLAN

Virtual overlay networks separate the virtual network from the underlying physical network devices. EVPN-VXLAN is often a common choice for standardized overlay solutions in next generation data center fabric. The EVPN-VXLAN control plane uses the as standard defined BGP protocol and extends communities to learn tenants' routes and distribute prefix lists. By standardizing five different BGP route types (as defined via the IETF in the integrated routing and bridging EVPN draft), EVPN-VXLAN can achieve forwarding optimization. The forwarding not only supports intra-subnet connectivity among hosts, but handles also the inter-subnet forwarding among hosts/VMs across different IP subnets.

Inter-subnet forwarding in EVPN. EVPN supports inter-subnet forwarding when IP routing is required. This approach ensures optimal forwarding, overcomes the shortcomings of the traditional centralized method, eliminates the need for traffic between subnets all the way from the PE to a centralized gateway node, and then returns to the PE for greater efficiency. In the latest design, the traffic is handled directly and locally from the PE, resulting in an optimal forwarding path. We tested two IETF drafts dedicated to inter-subnet forwarding approaches: the Integrated Routing and Bridging in EVPN draft (IRB) and the IP Prefix Advertisement draft. The EVPN instances used for both methods were based on RFC 7432.

Table 1: Inter-subnet Forwarding: Related Tested Features

IETF Draft and RFC	Tested Features
Integrated Routing and Bridging in EVPN draft	Asymmetric
	Symmetric
IP Prefix Advertisement draft	Interface-less use case
	Interface-full use case
RFC 7432	VLAN-based
	VLAN-Aware Bundle

- The "Integrated Routing and Bridging in EVPN" draft defines two different methods for routing traffic between VXLAN overlays. In the asymmetric model, both MAC-VRF and IP-VRF lookups are performed in ingress PE, only the MAC-VRF lookup is performed in the egress PE. Whereas in the symmetric mode, both MAC-VRF and IP-VRF lookup, are performed in ingress PE and egress PE.
- The VLAN-based and VLAN-Aware Bundle are two service interfaces as defined in RFC 7432. Only the VLAN-based service type has a single broadcast domain per EVI. The VLAN aware bundling service interface bundles different VLANs over a single EVPN instance. It can be used to multiplex several VLANs over the same EVI.
- The interface-less and interface-full scenarios are two use cases as discussed in the IP Prefix Advertisement draft. The interface-less option operates without an overlay next hop or a type-2 route in core network. The IRB is required in network virtual overlay components. In comparison to this, the IRB interface within an interface-full scenario is required in the core. Unnumbered in this case means that there is no IP address configured on the IRB interface (only MAC addresses are configured) for saving IPv4 space purpose.

We tested five different profiles using all features:

Table 2: Inter-subnet: Tested Profiles

No.	Tested Profile	Details
1	Asymmetric	With VLAN-based mode
2		With VLAN-Aware Bundle mode
3	Symmetric	With VLAN-based mode
4	IP-VRF-to-IP-VRF	Interface-less use case
5		Interface-full with unnumbered core-facing IRB

We chose a two-stage Clos topology for all profiles, also referred to as a "Leaf and Spine" network (as discussed in RFC 7938). The route server acted as spine and aggregated a set of horizontal EVPN PE devices for the leafs: the same number of IPv4 subnets was connected to every leaf device. Then we configured eBGP in the

physical network between spine and leaf devices (as underlay). The EVPN-VXLAN was used as the overlay between spine and leaf devices. In this setup, the EVPN-VXLAN was accessed by both IPv4 subnets connected to the leaf device whereas the subnets were emulated by the traffic generator.

Asymmetric IRB. In VLAN-based mode, the VXLAN virtual network identifier (VNI) directly maps to the EVPN EVI. Here we confirmed that RT-2 (MAC/IP advertisement) for the remote MAC addresses and the IPv4 routes were learned on the peer leaf devices via CLI. Then we sent IPv4 test traffic from all IPv4 subnets to any other IPv4 subnets, and expected to receive traffic on all IPv4 subnets without any packet loss.

In VLAN-Aware Bundle mode, we verified that the VNI was mapped to the Ethernet Tag field of the RT-2 route by capturing the BGP Updates. Afterwards, we confirmed that the RT-2 routes were learned on the EVPN-VXLAN PE devices, and sent full-meshed IPv4 traffic without any packet loss.

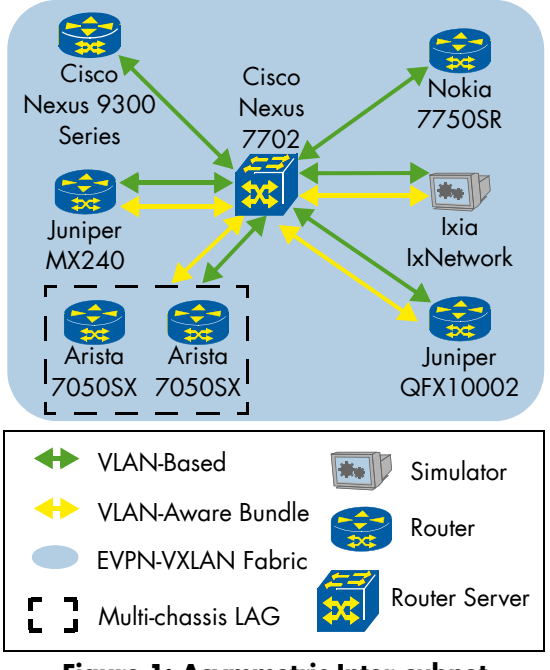

Figure 1: Asymmetric Inter-subnet Forwarding

There were some intensive discussions at the beginning of the test as the participants had different interpretations of the draft, but we were happy to find common ground in the end. We observed that all vendors could learn the routes without any packet lost in the service traffic.

Five vendors participated in the asymmetric VLAN-Based EVPN tests with the following DUTs: Arista 7050SX, Cisco Nexus 9300 Series, Ixia IxNetwork, Juniper MX240, Juniper QFX10002 and Nokia 7750SR.

Three vendors participated in asymmetric VLAN-Aware Bundle EVPN setup with the following DUTs: Arista 7050SX, Ixia IxNetwork, Juniper QFX10002 and Juniper MX240.

Spirent TestCenter and Ixia IxNetwork joined as the traffic generator. The Cisco Nexus 7702 acted as the route server.

Symmetric IRB. VLAN-based mode was tested in symmetric IRB only. We verified that the VXLAN virtual network identifier (VNI) directly mapped to the EVPN EVI via CLI, indicating that EVPN service is established. We then observed that the RT-2 for the remote MAC addresses and the IPv4 host routes were learned through the EVPN service on the peer leaf device via CLI. Finally, we sent IPv4 test traffic from all IPv4 subnets to any other IPv4 subnet, and expected to receive traffic on all IPv4 subnets without any packet loss.

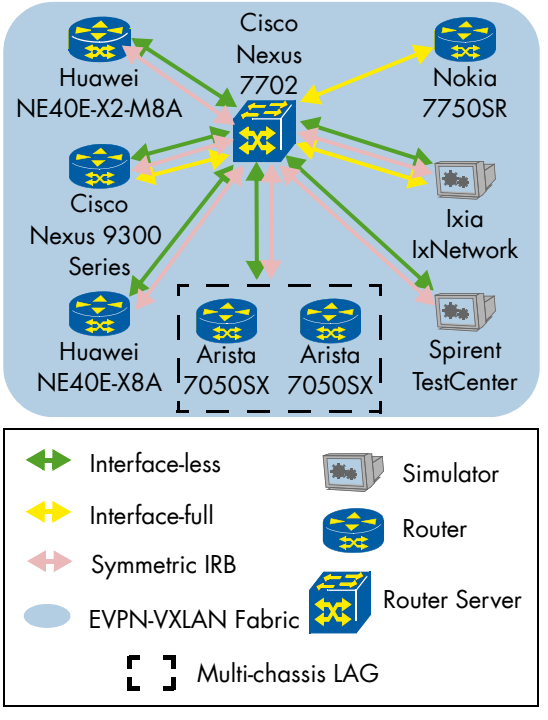

Figure 2: Symmetric IRB versus IP-VRF-to-IP-VRF (interface-less and interface-full)

Five vendors participated in the symmetric EVPN scenario with the following DUTs: Arista 7050SX, Cisco Nexus 9300 Series, Huawei NE40E-X2-M8A, Huawei NE40E-X8A, Ixia IxNetwork and Spirent TestCenter.

Spirent TestCenter and Ixia IxNetwork joined as the traffic generator.

Cisco Nexus 7702 participated as the router server.

IP-VRF-to-IP-VRF. The IETF specified a new route type RT-5 to separate the advertisement of IP prefixes from the advertisement of any MAC address related to it. We tested two selected use cases for the RT-5 based on the IP Prefix Advertisement draft as described below. In both cases we used the same setup as described in the symmetric IRB test (see Figure 2).

At first, we configured the interface-less mode on the EVPN-VXLAN PE devices. We verified that the VXLAN virtual network identifier (VNI) directly maps to the EVPN EVI. We then confirmed that the RT-5 (IP Prefix advertisement route) carried the correct IP length, the Route Target Gateway is set to 0 and was learned on the peer leaf device's IRB interface via CLI. Afterwards, we sent IPv4 test traffic from all IPv4 subnets to any other IPv4 subnet, and expected to receive traffic on all IPv4 subnets without any packet loss.

We configured the interface-full mode on the EVPN-VXLAN PE devices. We verified that the VXLAN virtual network identifier (VNI) directly mapped to the EVPN EVI. We confirmed that the RT-5 (IP Prefix advertisement route) carried the correct IP length, the Route Target Gateway is set to 0 and was learned on the peer leaf device's IRB interface via CLI. Additionally, a RT-2 was used in interface-full mode. It carried MAC address length and MAC address. The IP length was set to 0. Following, we sent IPv4 test traffic from all IPv4 subnets to any other IPv4 subnets, and expected to receive traffic on all IPv4 subnets without any packet loss.

Five vendors participated in the symmetric interface-less EVPN setup with the following DUTs (see Figure 2): Arista 7050SX, Cisco Nexus 9300 Series, Huawei NE40E-X2-M8A, Huawei NE40E-X8A, Ixia IxNetwork and Spirent TestCenter.

Three vendors participated in the symmetric interface-full EVPN test with the following DUTs: Cisco Nexus 9300 Series, Ixia IxNetwork, and Nokia 7750SR.

Spirent TestCenter and Ixia IxNetwork acted as the traffic generator.

Cisco Nexus 7702 participated as the router server.

MAC Mobility. MAC mobility is a mechanism that is used to detect host moving from one Ethernet Segment to another. It is achieved by adding a sequence number into the MAC/IP advertisement route (RT-2). When a host moves once, the sequence number is increased by one. The new PE sends an Ethernet MAC/IP advertisement route (EVPN RT2) to inform the other PEs to withdraw the route that has a smaller number.

The host was simulated either by Spirent TestCenter or Ixia IxNetwork. We moved the host by setting up the same MAC address of the previous host on a different device and removed the MAC address from the previous device. We first verified that the sequence number was 0 (or not set) before the MAC (host) was moved. When the MAC (host) moved once, we observed that the sequence number was exactly increased by one as expected.

We also measured the out of service time during the host movement. We used a constant rate of packets (1,000 packets/s) and performed the host movement. Then we measured the out of service time which we calculated based on the lost number of packets. The out of service times from different vendors were between 22 ms to 130 ms.

The following vendors participated the in test: Arista 7050SX, Huawei-NE40E-X8A, Ixia IxNetwork and Nokia 7750SR.

Nokia 7750SR joined as BGP route reflector in the EVPN overlay network.

Cisco Nexus 7702 joined as route server.

Spirent TestCenter and Ixia IxNetwork acted as the traffic generator.

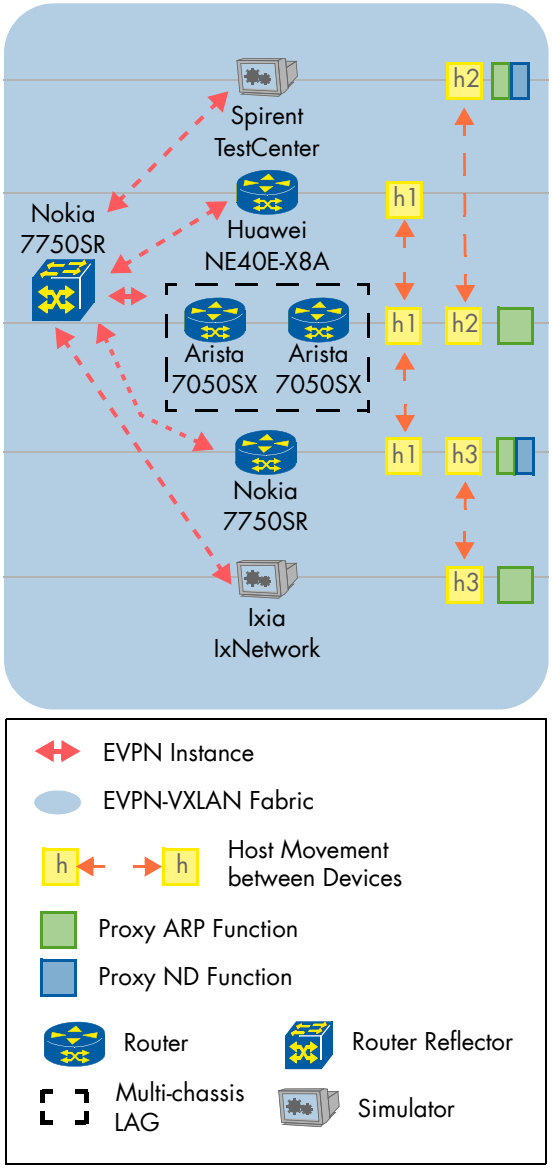

Figure 3: MAC Mobility and Proxy ARP/ND

Initially, one vendor failed to learn the first update when the host was moved from its own device to a remote peer. The DUT replied with a sequence number that was higher than the received one. The DUT got a RT2 route sent by the remote site to notify the changes, indicating that the learning failed. This error prompted the process to be repeated one more time. Afterwards, the DUT successfully learned the updates from the second RT2 route sent by the remote peer.

Single Homed EVPN with Proxy ARP and Proxy ND. The goal was to use ARP Proxy/Proxy ND (Neighbor Discovery) to learn MAC addresses across the EVPN instance. This is achieved by using RT-2 to exchange the MAC address and to store it locally. When one PE receives an ARP request/Neighbor solicitation, it will search locally first. Only if no proper result is found, the ARP request/Neighbor solicitation floods to the other remote PE.

The first step was to discover the prefix by having the PE connected at first with a host that was emulated by Ixia IxNetwork or Spirent TestCenter.

Then we sent a gratuitous ARP from the emulated host to the PE via CLI (see Figure 3). We started the capture via the traffic generator through the entire process. To complete this step and in order to verify that the ARP flooding was performing on the PE after a new host has entered the network, we checked the packet capture and observed that a RT-2 was generated from the PE to the remote PEs.

In the next step, we verified that the ARP proxy limited the number of ARP messages for a MAC address that was learned. We sent an ARP request from the host connected to the PE in order to observe that this ARP request was responded by the local PE only (running proxy ARP function). We observed the reply messages from the PE via packet capture. Finally, we ran the IPv4 traffic and expected no packet loss.

The following vendors joined: Arista 7050SX, Ixia IxNetwork, Nokia 7750SR and Spirent TestCenter.

Nokia 7750SR acted as BGP route reflector in the EVPN overlay network.

Cisco Nexus 7702 was the router server.

Ixia IxNetwork participated as the traffic generator.

For the Proxy ND scenario, we did not send any Neighbor Advertisement (NA) packets from the host. Based on the proxy-arp-nd draft, the PE learns the MAC/IPv6 from the NA messages only, but the host cannot generate these packets. Since this was not our focus, we manually configured a static route on the PE. Then we performed the second step identically as described above using Neighbor Solicitation. We did not observe any packet loss.

The following vendors participated in the test: Nokia 7750SR and Spirent TestCenter.

Spirent TestCenter also participated as the traffic generator.

EVPN-MPLS

EVPN-MPLS is another common overlay solution for data center interconnect. It uses labels instead of UDP encapsulation. As it inherits all the features of the MPLS technology, it can communicate with a variety of MPLS related areas, such as segment routing. In addition, the EVPN-MPLS can be used as part of the MEF-based E-Line and E-Tree implementation.

Ethernet Line (E-Line). The MEF defines a point-to-point Ethernet service as Ethernet Line (E-Line). The IETF now proposes a solution framework in order to support this service in MPLS networks using EVPN. The IETF discussed the features in the "VPWS support of the EVPN" draft and requires the use of VPWS to meet the E-Line requirements. In addition, EVPN also supports inherited functions to make the VPWS implementation more effective.

With the EVPN multi-homing function, EVPN offers resilience on multi-redundancy service for VPWS. There are two redundancy modes: single-active and all-active. When a device is multi-homed to two or more PEs, only one of the PEs can forward the traffic. This is called single-active redundancy mode. When an all-active redundancy mode is implemented, all PEs that connect to the same device can forward the traffic and perform per flow based load-balance by using the LAG protocol.

In the test, we setup a point-to-point connection between two PEs. We configured an EVPN instance between both devices and enabled VPWS inside EVPN instances.

In single-homing mode, we verified that the ESI field was set to zero and that the Ethernet Tag field mapped to the VPWS identifier, both of which were carried in the EVPN AD per EVI route.

In multi-homing mode, we sent Ethernet traffic for the E-Line service, and disconnect the cable between the CE and PE. During the failover, we measured out of service time up to 2 ms using single-active mode, and under 1 ms using all-active mode. We measured 2 ms - 10 ms out of service time during the recovery.

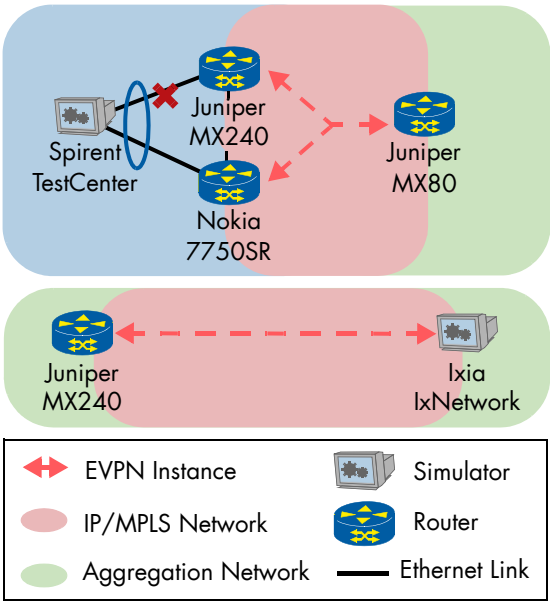

Figure 4: EVPN E-Line Service

The following vendors participated in the single-homing scenario: Ixia IxNetwork, Juniper MX240. The following vendors joining multi-homing (both single-active and all-active): Juniper MX240, Juniper MX80 and Nokia 7750SR.

Spirent TestCenter and Ixia IxNetwork acted as the traffic generator.

Initially, one DUT was not sending the ESI-label extended community with the AD per-EVI route in all-active multi-homing mode, even though the extended community is mandatory in multi-homing mode. The vendor soon fixed the issue by updating the software. After that all DUTs successfully forwarded EVPN traffic without any packet loss.

Ethernet Tree (E-Tree). The MEF defines a rooted-multipoint Ethernet service as Ethernet Line (E-Line). Again, the IETF proposes a solution for supporting this service in MPLS networks by using EVPN.

In the setup we verified that the EVPN technology can support E-Tree functional requirements. Based on the current IETF draft, we tested two scenarios: root/leaf per PE and root/leaf per AC (attached circuit). The difference between both is the type of endpoint that is connected to the PE: it is either root or leaf (in the first case), or a mixture of root and leaf (in the latter case).

We started the test with the verification of the imported route type-3 (inclusive multicast Ethernet tag route) via CLI on the devices under test. As expected, all leafs imported the routes only from the root whereas the roots imported the routes from the leafs.

We generated unicast Ethernet traffic between the traffic generator connected to the root and traffic generators connected to the leafs as well as between the leafs. As expected, we observed no packet loss between root and leafs, and 100% packet loss between the leafs.

The following vendors joined both test scenarios: Juniper MX240 and Nokia 7750SR.

Spirent TestCenter participated as the traffic generator.

At the beginning, we observed that the data plane traffic was dropped. The participating vendors found out that they had a different understanding of the encoding of the label field in the E-Tree extended community sent along with the AD per-ES route. The two vendors agreed on the correct behavior and one of them fixed it. In addition, due to this findings, a paragraph has been added to the EVPN-ETREE draft and will be included in the next published version.

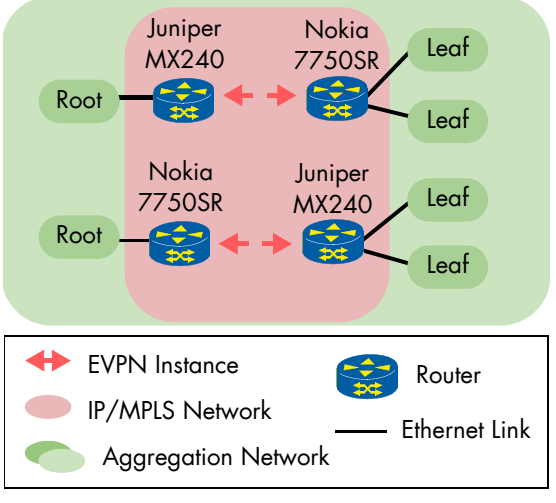

Figure 5: E-Tree Service - Leaf or Root Sites per PE

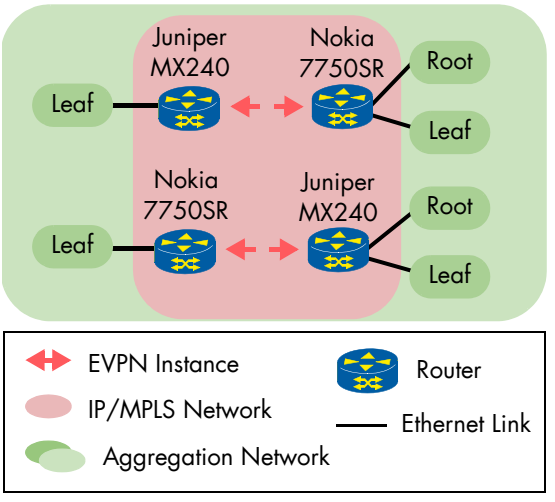

Figure 6: E-Tree Service - Leaf or Root Sites per AC

EVPN with MPLS/Segment Routing Transporting. EVPN provides a separation between the data plane and control plane which allows the use of different encapsulation mechanisms in the data plane such as MPLS, Virtual Extensible LAN (VXLAN), etc.

The test proofed EVPN functionality over MPLS data plane based Segment Routing (SR).

We started with the verification of the state of Interior Gateway Protocol (IGP) and exchanged Segment Routing information on each of the network nodes. We checked that the node and adjacency segment identifiers (SIDs) were exchanged and the appropriated labels were installed in the Forwarding Information Base (FIB) on each of the network nodes.

Afterwards, we tested the state on the EVPN control plane, namely the imported route type-3 (inclusive multicast Ethernet tag route) via CLI on the network nodes. As expected, each of the network nodes imported the route type-3 from the other network nodes.

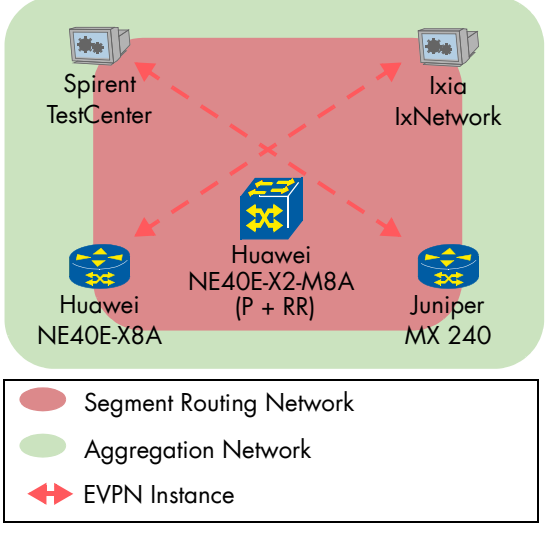

Figure 7: EVPN with MPLS/Segment Routing Transporting

We generated unicast and broadcast Ethernet traffic between the emulated devices and network nodes. Afterwards, we verified that the route type-2 carrying the remote MAC addresses/IPv4 routes were learned on the other devices via CLI. As expected, the traffic was forwarded without packet loss.

The following vendors participated in the test: Huawei NE40E-X8A, Ixia IxNetwork, Juniper MX240, Spirent TestCenter.

Huawei NE40E-X2-M8A acted as Provider router in the segment routing and as BGP route reflector in EVPN.

Spirent TestCenter and Ixia IxNetwork participated as the traffic generator.

EVPN Interworking

Data center interconnect is a crucial part in providing end-to-end services across data centers as well as internet access for a data centers. In this section, we tested EVPN-VXLAN interconnect with common IP/MPLS technologies including IP-VPN and EVPN-MPLS.

EVPN-VXLAN and EVPN-MPLS Interworking. The goal was to verify interworking between a WAN network that was based on EVPN-MPLS and data center networks that were based on EVPN-VXLAN. As shown in the figure, the devices in the data center network provided EVPN-VXLAN connections. The IP/MPLS edge routers provided the interconnection between the EVPN-VXLAN network and EVPN-MPLS for the control plane and data plane interoperability.

Once we tested the BGP sessions in the underlay and overlay, we checked the BGP routing table on each of the IP/MPLS edge devices (PE). Each of the PE devices received route type-3 from each remote EVPN PE.

We generated unicast Ethernet traffic between the sites and started to verify the MAC/IP advertisement (route type-2) routes in each of the network nodes. The MAC/IP advertisement routes in the IP/MPLS network segment were carrying the RD of common EVI, the MAC addresses, and the MPLS label associated with the MAC. The MAC/IP advertisement routes in the VXLAN network segments were carrying the RD of common EVI, the MAC addresses, the VNI associated with the MAC, and VXLAN encapsulation as extended community.

The following devices acted in this test as interworking devices: Huawei NE40E-X8A, Juniper MX240 and Nokia 7750SR.

Arista 7050SX, Ixia IxNetwork, and Spirent TestCenter participated as PE devices in the EVPN-VXLAN network.

Juniper MX80 joined as PE in the EVPN-MPLS network.

Nokia 7750SR participated as BGP route reflector in the EVPN overlay network.

The Cisco Nexus 7702 was the BGP router in the IP underlay network.

Finally, Spirent TestCenter and Ixia IxNetwork participated as the traffic generator.

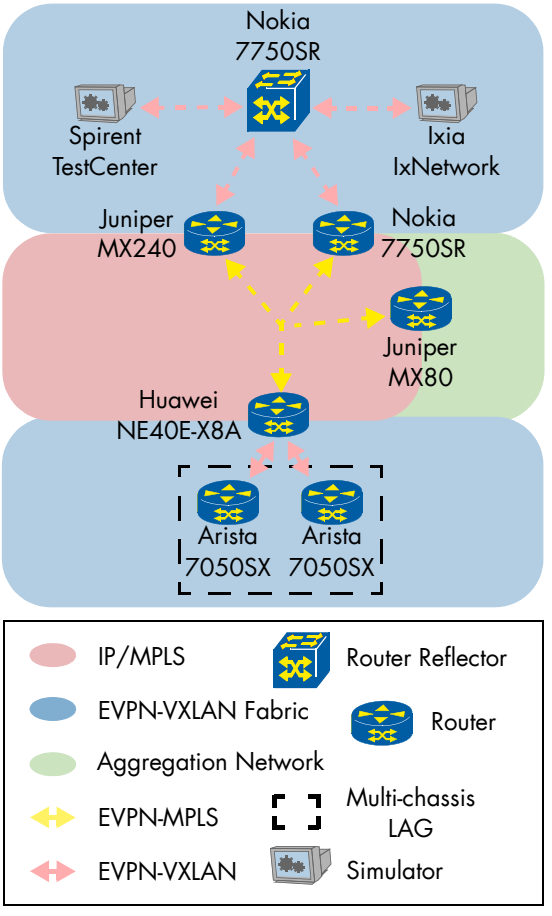

Figure 8: EVPN-VXLAN and EVPN-MPLS Interworking

EVPN and IP-VPN Interworking. In this test we verified interworking between IP VPN and EVPN networks. Two network edge devices provided control plane and data plane interworking.

First, we confirmed that the network edge devices translated the EVPN route type 5 from EVPN networks to IPv4-VPN routes in the IP/MPS core network. Then, we sent full mesh bidirectional traffic between all sites. The traffic was forwarded without packet loss.

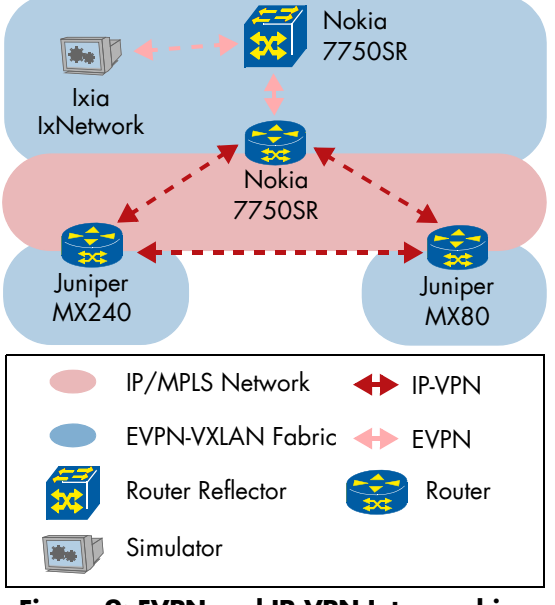

Figure 9: EVPN and IP-VPN Interworking

Ixia IxNetwork, Juniper MX240, Juniper MX80 and Nokia 7750SR participated.

Nokia 7750SR acted as BGP route reflector in the EVPN overlay network.

Cisco Nexus 7702 participated as BGP router in the IP underlay network.

Ixia IxNetwork joined as the traffic generator.

Remote Loop Free Alternates FRR. The goal of IP Fast ReRoute is to reduce the failure reaction time to tens of milliseconds by using pre-computed loop free alternate next hops. LFA does not provide coverage for all network topologies, in particular the commonly-found ring topology is not supported. Hence, an additional flavor of LFA is required, known as Remote LFA.

The basic idea behind remote LFA is to find a set of nodes that can be reached by the Point of Local Repair's (PLR) immediate neighbors without traversing the primary next-hop (extended P-space) as well as nodes that can reach the destination by normal forwarding without traversing the failed link (Q-space). A set of nodes in extended P-space and Q-space are termed as PQ nodes. Tunnel technologies, such as Multiprotocol Label Switching (MPLS), IP in IP, GRE or Segment Routing can then be used to encapsulate traffic from the PLR to the PQ-nodes.

We first tested that the IGP (IS-IS) and LDP sessions were correctly set up between all the devices. We then confirmed that the repair tunnel was established between the PLR and the PQ-node. While generating traffic at a fixed rate of 1,000 packets/s we introduced a link failure by disconnecting the cable and measured the time the PLR needed to switchover the service traffic to the backup path. The service traffic was sent in a VPWS service.

We tested four test combinations each consisting of four devices. The following devices participated in the test as Point of Local Repair (PLR): Ericsson Router 6672, Huawei NE40E-X8A, ECI NPT-1800 and Juniper MX104. The following devices acted as PQ node: Huawei NE40E-X8A and ECI NPT-1800.

Upon link failure, 3 out of 4 network devices switched over to the backup path within 25ms and one device exceeded the expectations. After restoring the link, one device demonstrated hitless recovery, two devices reacted within 60 ms on average and one device exceeded the expectations.

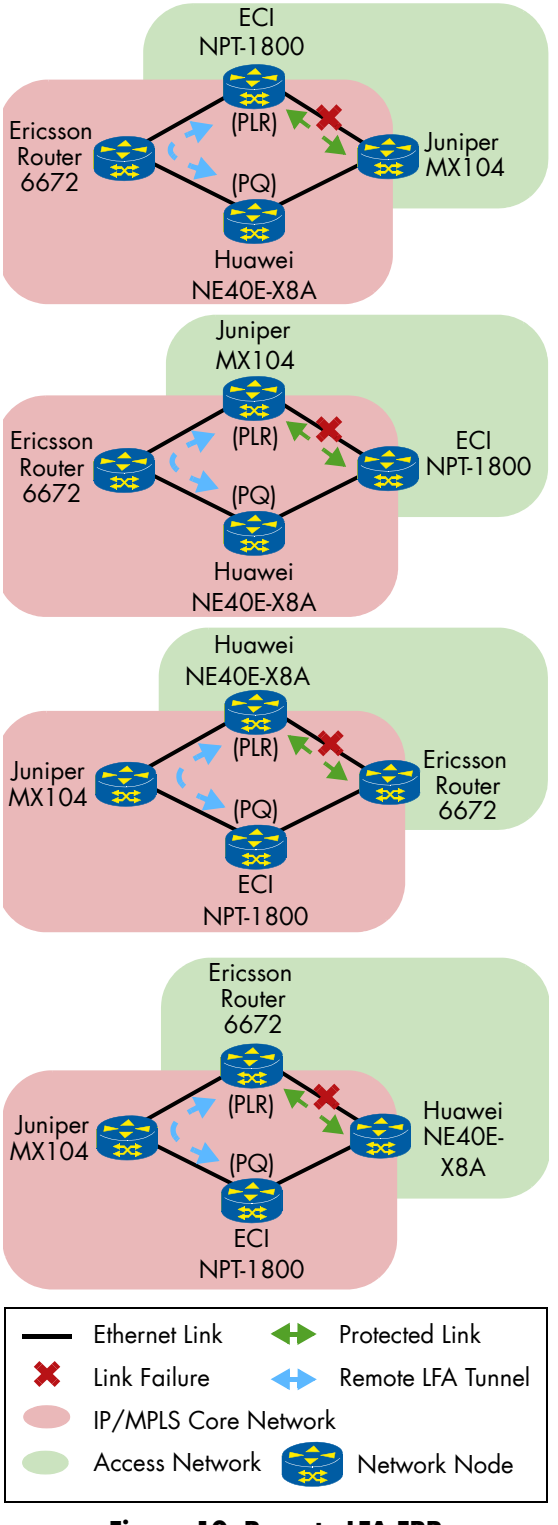

Figure 10: Remote LFA FRR

Segment Routing Fast Reroute (FRR).

Segment Routing utilizes the local repair properties of IP Fast Reroute (IP FRR) in conjunction with explicit routing to protect and maintain end-to-end connectivity without requiring additional signaling upon failure.

In this test, we focused on Loop Free Alternate (LFA) in a Segment Routing network. The LFA approach is applicable when the protected node or Point of Local Repair (PLR) has a direct neighbor that can reach the destination without looping back traffic to the PLR. When the protected path fails, the traffic will be sent to the neighboring node which in turn forwards the traffic to the destination.

We started this test by verifying the distribution of node and adjacency segment identifiers (SIDs) on each of the network nodes. We also confirmed the Forwarding Information Base (FIB) when the LFA was not configured. As expected, there was a single forwarding entry for each of the node segment IDs. We performed baseline measurement for packet loss using bidirectional traffic from a traffic generator.

Afterwards, we configured LFA on the network nodes and verified that the network nodes installed backup forwarding entry in FIB. While the traffic was running via the primary path we disconnected the link and measured the service interruption time based on the packet loss. We saw that the traffic was taking the backup path.

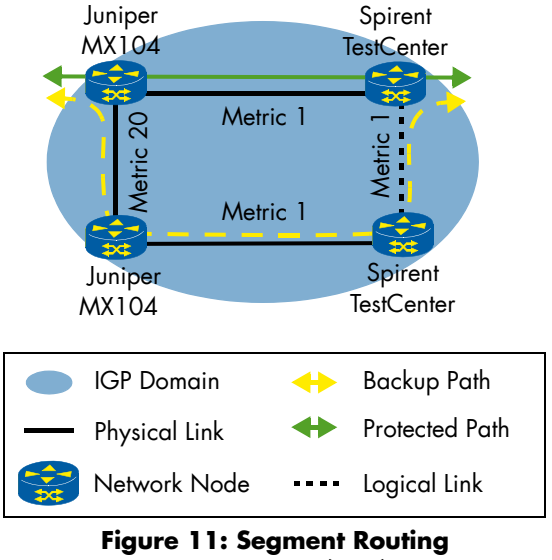

Figure 11: Segment Routing Fast Reroute (FRR)

Juniper and Spirent participated in this test. Two Juniper MX104 and two Spirent TestCenter acted as Network Nodes. Spirent was also used for traffic generation.

Segment Routing and LDP Interworking.

Segment Routing architecture needs to provide interworking mechanisms for seamless interconnection with existing IP/MPLS networks.

The interworking architecture defines two functional blocks, the SR mapping server and mapping client. The SR mapping server provides interworking between SR and LDP networks. The mapping server advertises remote-binding Segment ID for prefixes from non-SR capable LDP nodes. The SR mapping client uses this mapping to forward the traffic to the LDP nodes.

We started the test with the SR mapping server – the function that interconnects the SR capable device and the non-SR capable LDP device. We verified the advertised prefix is corresponding to the non-SR capable LDP device and its associated Segment ID.

We confirmed that the SR capable device (mapping client) processed the mapping and programed its MPLS forwarding table accordingly.

We also tested the control and data plane operations of a L3VPN service between the SR and non-SR capable devices. We generated bidirectional traffic between the attachment circuits to verify the proper encoding of the data path.

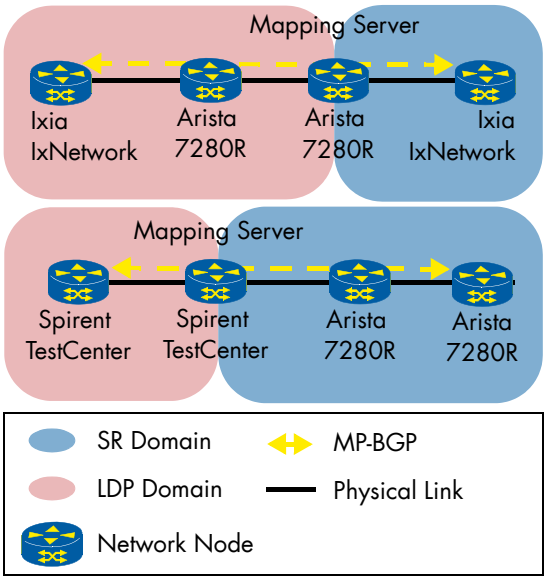

Figure 12: Segment Routing and LDP Interworking

Arista, Ixia and Spirent participated in this test. In the first scenario Arista 7280R acted as segment routing mapping server, Arista 7280R and Ixia IxNetwork were present in the LDP domain and Ixia IxNetwork was present in the SR domain. In the second scenario Spirent TestCenter acted as segment routing mapping server, two Arista 7280R were present in the SR domain and one Spirent TestCenter was present in the LDP domain.

Segment Routing Anycast Segment – Disjointness in Dual-Plane Networks.

Segment Routing provides a new solution for disjoint paths within dual-plane networks. Disjointness allows to transport different services of traffic across disjoint paths. Here we verified that this can be achieved by using SR Anycast segments in SR routers.

The Anycast segment Identifier (Anycast SID) is specified for a set of routers within a data plane. Each SR router in the Anycast set advertises the same Anycast-SID which represents an ECMP-aware, shortest-path IGP route to the closest node of that Anycast set.

One network node was located in the anycast SID domain 101 and two network nodes were located in the SID domain 102. We verified that the IGP (IS-IS) sessions were successfully established. We tested that the PE routers learned the Node-SID, the Adj-SID and the Anycast-SID from the other nodes. We configured each PE router to push the Node SIDs of the other PE for service traffic 1 and service traffic 2 and then verified that the traffic was load-balanced between the networks nodes of the anycast SID domains 101 and 102. We set the PE routers to push the anycast SID 101 for service traffic 1 and the anycast SID 102 for service traffic 2 (on top of the Node SIDs). We checked that the service traffic 1 was forwarded along the nodes of the anycast SID 101 and the service traffic 2 was forwarded along the nodes of the anycast SID 102.

In the first test combination, two Arista 7280R were acting as PE routers and one Juniper MX104 was acting as P router. One Juniper MX104 router was located in the Anycast domain with SID 101 and two Juniper MX104 routers were located in the Anycast domain with SID 102.

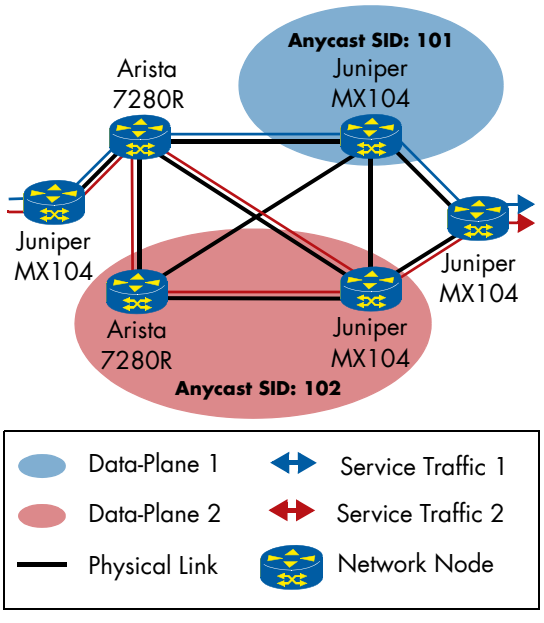

Figure 13: Disjointness in Dual-Plane Networks - Anycast Segment

In the second test combination, two Juniper MX104 acted as PE routers and one Arista 7280R was acting as P router. One Juniper MX104 was located in the Anycast domain with SID 101, one Juniper MX104 and one Arista 7280R were located in the Anycast domain with SID 102. Spirent TestCenter was used as traffic generator.

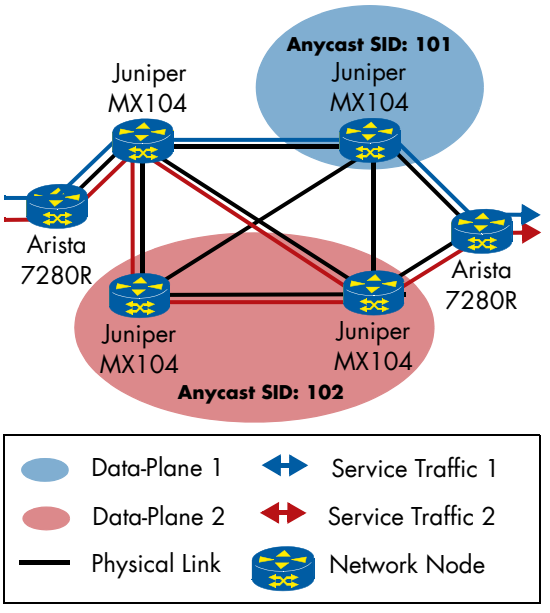

Figure 14: Disjointness in Dual-Plane Networks — Anycast Segment & CoS Based Traffic Engineering

Segment Routing Anycast Segment — CoS Based Traffic Engineering. Some applications and service traffic require special network characteristics such as bandwidth and latency. The Segment Routing Anycast segment provides a Class of Service (CoS) based traffic engineering technology that allows to steer traffic with heavy load away from the shortest path toward a higher latency or higher bandwidth path.

One network node was located in the anycast SID domain 101 and two network nodes were located in the anycast SID domain 102. We verified that the IGP sessions were successfully established. Then we confirmed that the PE routers learned the Node-SID, the Adj-SID and the Anycast-SID from the other nodes. We configured each PE router to push the Node SIDs of the other PE for service traffic 1 and service traffic 2 and tested that the traffic was load-balanced between networks nodes of the anycast SID domains 101 and 102. Following, we set the PE routers to push the anycast SID 101 for service traffic 1 marked with CoS ID = 1 and the anycast SID 102 for service traffic 2 marked with CoS ID = 2 (on top of the Node SIDs). We checked that the service traffic 1 was forwarded along the nodes of the anycast SID domain 101 and the service traffic 2 was forwarded along the nodes of the anycast SID domain 102.

Arista and Juniper participated in this scenario, with two Arista 7280R acting as PE routers performing the CoS based Traffic Engineering and one Juniper MX104 acting as P router. One Juniper MX104 router was located in the Anycast domain with SID 101 and two Juniper MX104 routers were located in the Anycast domain with SID 102.

Microwave Transport

As operators upgrade their networks with LTE-A in preparation for a 5G future, questions are being asked around the role Microwave transport will play. We see a trend of integrating the more specialised microwave devices into the standard IP/MPLS router domain, and thus looked into two particular aspects of this in the following tests.

Microwave: Bandwidth Notification Message. Today Mobile Backhaul networks are often built as an overlay with routers sitting 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.

The goal of this test was to check that an aggregation router can process the bandwidth notification messages (ETH-BN) and accordingly apply service policies to the traffic sent to the microwave system, based on the bandwidth information within the ETH-BN.

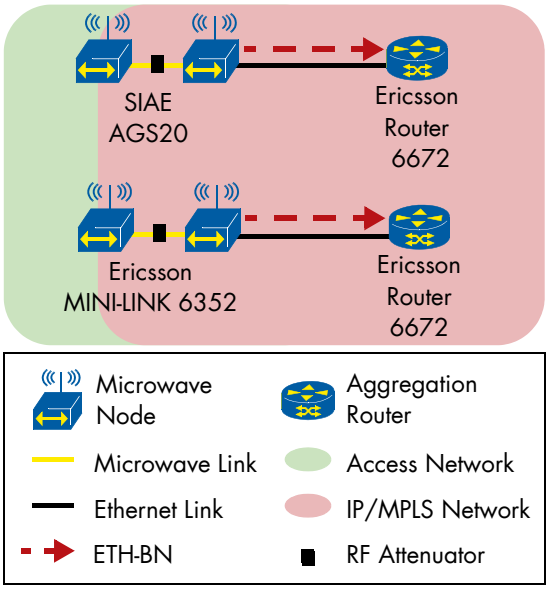

Figure 15: Bandwidth Notification Message

We used Spirent TestCenter in order to generate traffic for two Ethernet services from the aggregation to the access network and verified that no packets were lost while the microwave nodes were using the maximum modulation scheme available (shown in Table 3).

Table 3: Maximum Modulation Used

Device	Maximum modulation used	Channel Width
Ericsson MINI-LINK 6352	256 QAM	750 MHz
SIAE MICRO-ELETTRONICA AGS20	2048 QAM	56 MHz

To emulate severe weather conditions, we reduced the bandwidth between the two nodes of the microwave system using an RF attenuator. Then we verified that the microwave system generated a PDU with ETH-BN, indicating the reduction of the microwave link bandwidth. The aggregation router applied a traffic policer according to the ETH-BN packet received. While at this point it was still not possible to change the CIR or the path of a specific service, we observed that the Ericsson Router 6672, upon receiving the bandwidth notification message, was able to apply a policer on the traffic, adjusting the outgoing bandwidth according with the information received.

In both test combinations Ericsson Router 6672 acted as aggregation router. The Microwave link was established between two AGS20 nodes from SIAE MICROELETTRONICA and two MINI-LINK 6352 from Ericsson.

Layer 3 Microwave: MPLS-based Services. The aim was to confirm the capability to establish IP/MPLS service on a microwave platform crossing or terminating on existing infrastructure.

We tested over 20 different combinations relying on different transport profiles, and verified that a VPWS or L3VPN service can be set up between IP/MPLS capable microwave systems and IP/MPLS aggregation routers in multi-vendor scenario.

We created both end-to-end services between two different microwave vendors with standalone routers participating as aggregation routers, as well as between a microwave vendor and a standalone router acting as PE.

Table 4: Microwave MPLS-based Services: Profiles Used

	Profile ID		
	1	2	3
IGP Protocol	OSPFv2	IS-IS	OSPF-TE
Transport Signalling protocol	LDP	LDP	RSVP-TE
MPLS Service	L3VPN	L3VPN	VPWS
Service Signaling Protocol	mp-BGP	mp-BGP	t-LDP
PE-CE Protocol	Static routing	Static routing	N/A
number of services	2	2	2

Table 6 depicts in detail the protocols used for each of the three profiles.

Aviat CTR 8540, Ericsson MINI-LINK 6691, SIAE MICROELETTRONICA AGS20 participated in this test as microwave nodes. Ericsson Router 6672, Huawei ATN950C, Juniper MX104 participated in the PE router role, while Huawei ATN950C, Juniper MX104 participated as aggregation router.

Table 5: Maximum Modulation Used

Device	Maximum modulation used	Channel Width
Aviat CTR8540 (radio device ODU600)	1024 QAM	56 MHz
Aviat CTR8540 (radio device WTM 4000)	4096 QAM	112 MHz
Ericsson MINI-LINK 6691	4096 QAM	112 MHz
SIAE MICRO-ELETTRONICA AGS20	2048 QAM	56 MHz

Table 6: MPLS-based Services: Tested Combinations per Profile

	Aviat CTR 8540	SIAE AGS20	Ericsson MINI-LINK 6691
Aviat CTR 8540		1,2,3	1
SIAE AGS20	1,2,3		1
Ericsson MINI-LINK 6691	1	1	
Ericsson Router 6672	1,2	1,2,3	
Huawei ATN 950C	1	1	1
Juniper MX104	1,2	1,2	1
Nokia 7750 SR		2	

Layer 3 Microwave: Transport Resiliency.
Bringing IP/MPLS to the microwave network provides additional resiliency options in the access network and increases the end-to-end service availability.

The goal here was to show that a microwave node can react to degradation of the radio link by re-routing the traffic via a different path.

We used Spirent TestCenter to act as CE and send bidirectional traffic to the DUTs. We verified that the microwave nodes were using the main path with the maximum modulation scheme available and that no packets were lost. We then emulated severe weather conditions by reducing the available bandwidth of the channel with a RF attenuator. Ericsson MINI-LINK 6691 and SIAE MICRO-ELETTRONICA AGS20 participated as microwave stations and successfully performed a failover to the backup path when the available bandwidth of the link was equal to zero (Ericsson) or reduced below a specific threshold (SIAE). Juniper MX104, Nokia 7750SR, SIAE MICROELETTRONICA AGS20, and Spirent TestCenter participated as P and PE routers.

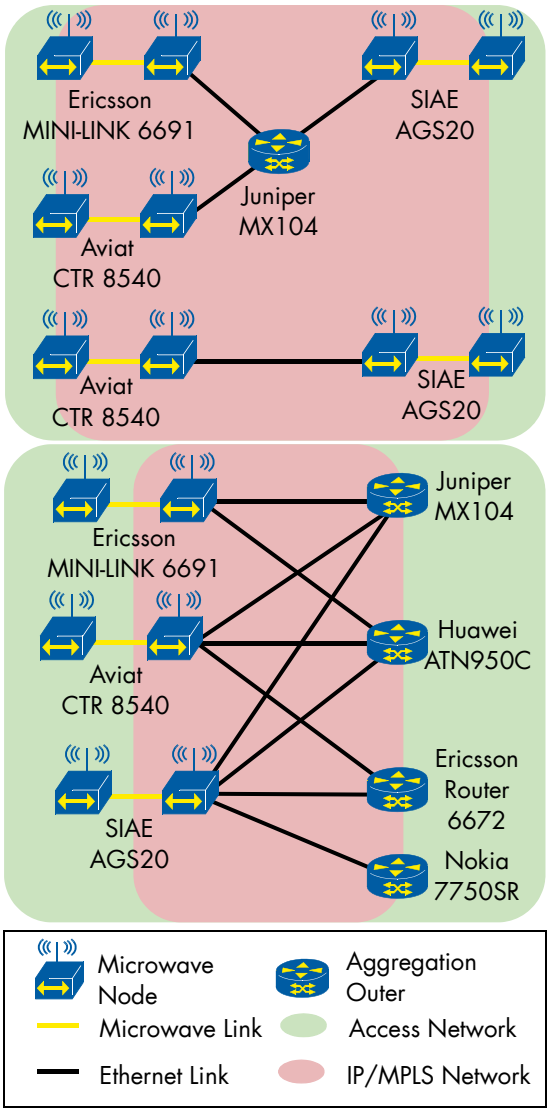

Figure 16: Layer 3 Microwave: MPLS-based Services

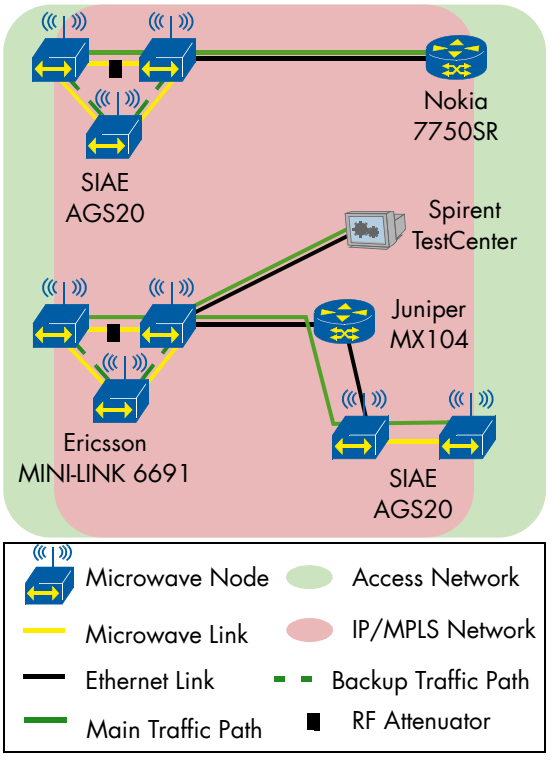

Figure 17: Layer 3 Microwave: Transport Resiliency

Software Defined Networking

This chapter focuses on two topics: NETCONF/YANG and Path Computation Element Protocol (PCEP). In this year's NETCONF/YANG section we introduced the testing of Virtual CPE Integration in WAN SDN Networks. We noticed an enhanced attention for NETCONF/YANG by a large number of vendors. However, the support of standardized models still remains outstanding. We also noticed an increased interest in the openconfig project from multiple participants. New vendors showed a great interest in the orchestration role and a test covering multi-vendor controllers for VPN creation was covered. The PCEP section witnessed a wide support of the different tests this time. PCC-initiated as well as PCE-initiated (using Segment Routing LSP) were successfully verified in multiple combinations.

NETCONF

The NETCONF/YANG protocol defines an easy to use mechanism through which a network device can be managed, configuration data information can be retrieved and new configuration data can be uploaded and manipulated. YANG is the data modeling language used by NETCONF. YANG data models are being standardized for various scenarios such as BGP/MPLS L3 VPNs. Therefore, we relied on proprietary YANG models that were provided by the participants. The YANG models supported by the NETCONF servers were provided to the vendor bringing the NETCONF client, and were compiled previous to the test.

Device Configuration Using NETCONF/YANG. In this test we used NETCONF/YANG functionality to manage configurations and retrieve the operational state on a network device.

We started the test with the verification of the NETCONF session between the NETCONF client and the NETCONF server. Afterwards, we retrieved the running configuration from the network devices. We also proved the usage of the sub-tree filtering to recall the running interface configuration. Subsequently we assessed a configuration change and then a configuration deletion on the device. Finally, we tested that the operational state could be retrieved and that the NETCONF session could be closed successfully.

In the first test setup, Huawei Agile Controller was managed by an emulator (postman) via the Northbound RESTCONF Application Program Interface (API). The emulator was used to simulate the northbound communication with the Agile Controller. We verified that Ericsson Router 6672 and RAD ETX were successfully operated by the Huawei Agile Controller.

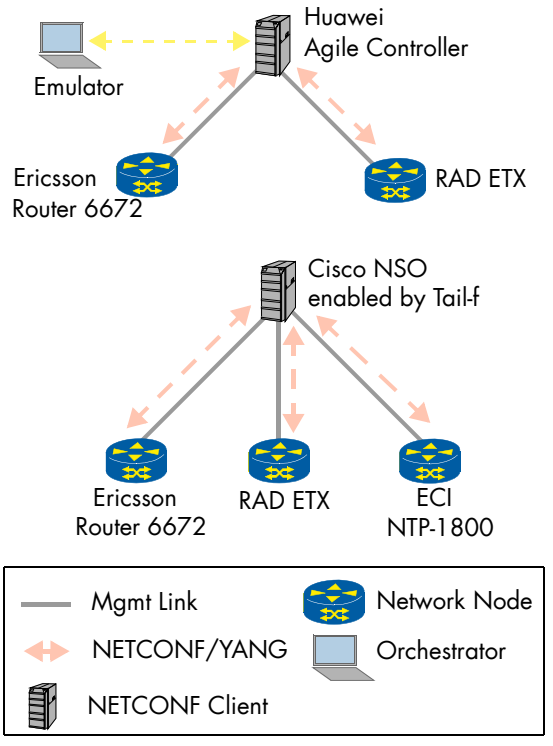

Figure 18: Device Configuration Using NETCONF/YANG

In the second setup, three NETCONF servers were successfully involved. Cisco enabled by Tail-f acted as the NETCONF client and we tested it against ECI NPT-1800, Ericsson Router 6672 and RAD ETX. The operational state could not be retrieved with one device, as its YANG model only contained the configuration information.

L3VPN Service Creation Using NETCONF/YANG. Each vendor provided their L3VPN service models and service delivery models. This test confirmed that those models can be used to configure and manage BGP L3VPNs.

The scenario was similar to the previous one but here we focused on the verification of the L3VPN service creation. We tested that the VRF specific parameters and BGP specific parameters that are applicable for L3VPNs were correctly configured. Traffic was generated inside the configured service and the service was finally removed using NETCONF.

Four participants joined three test setups.

In the first test setup, Cisco NSO enabled by Tail-f acted as the NETCONF client, ECI NTP-1800 and Ericsson Router 6672 as the NETCONF servers.

In the second test setup, Huawei CloudOpera IES was involved as the orchestrator interacting with the Northbound API of Cisco NSO using the NETCONF protocol. The service configuration was initiated by the Huawei CloudOpera IES through the Cisco NSO controller. Cisco CSR 1000V and Ericsson Router 6672 acted as NETCONF servers.

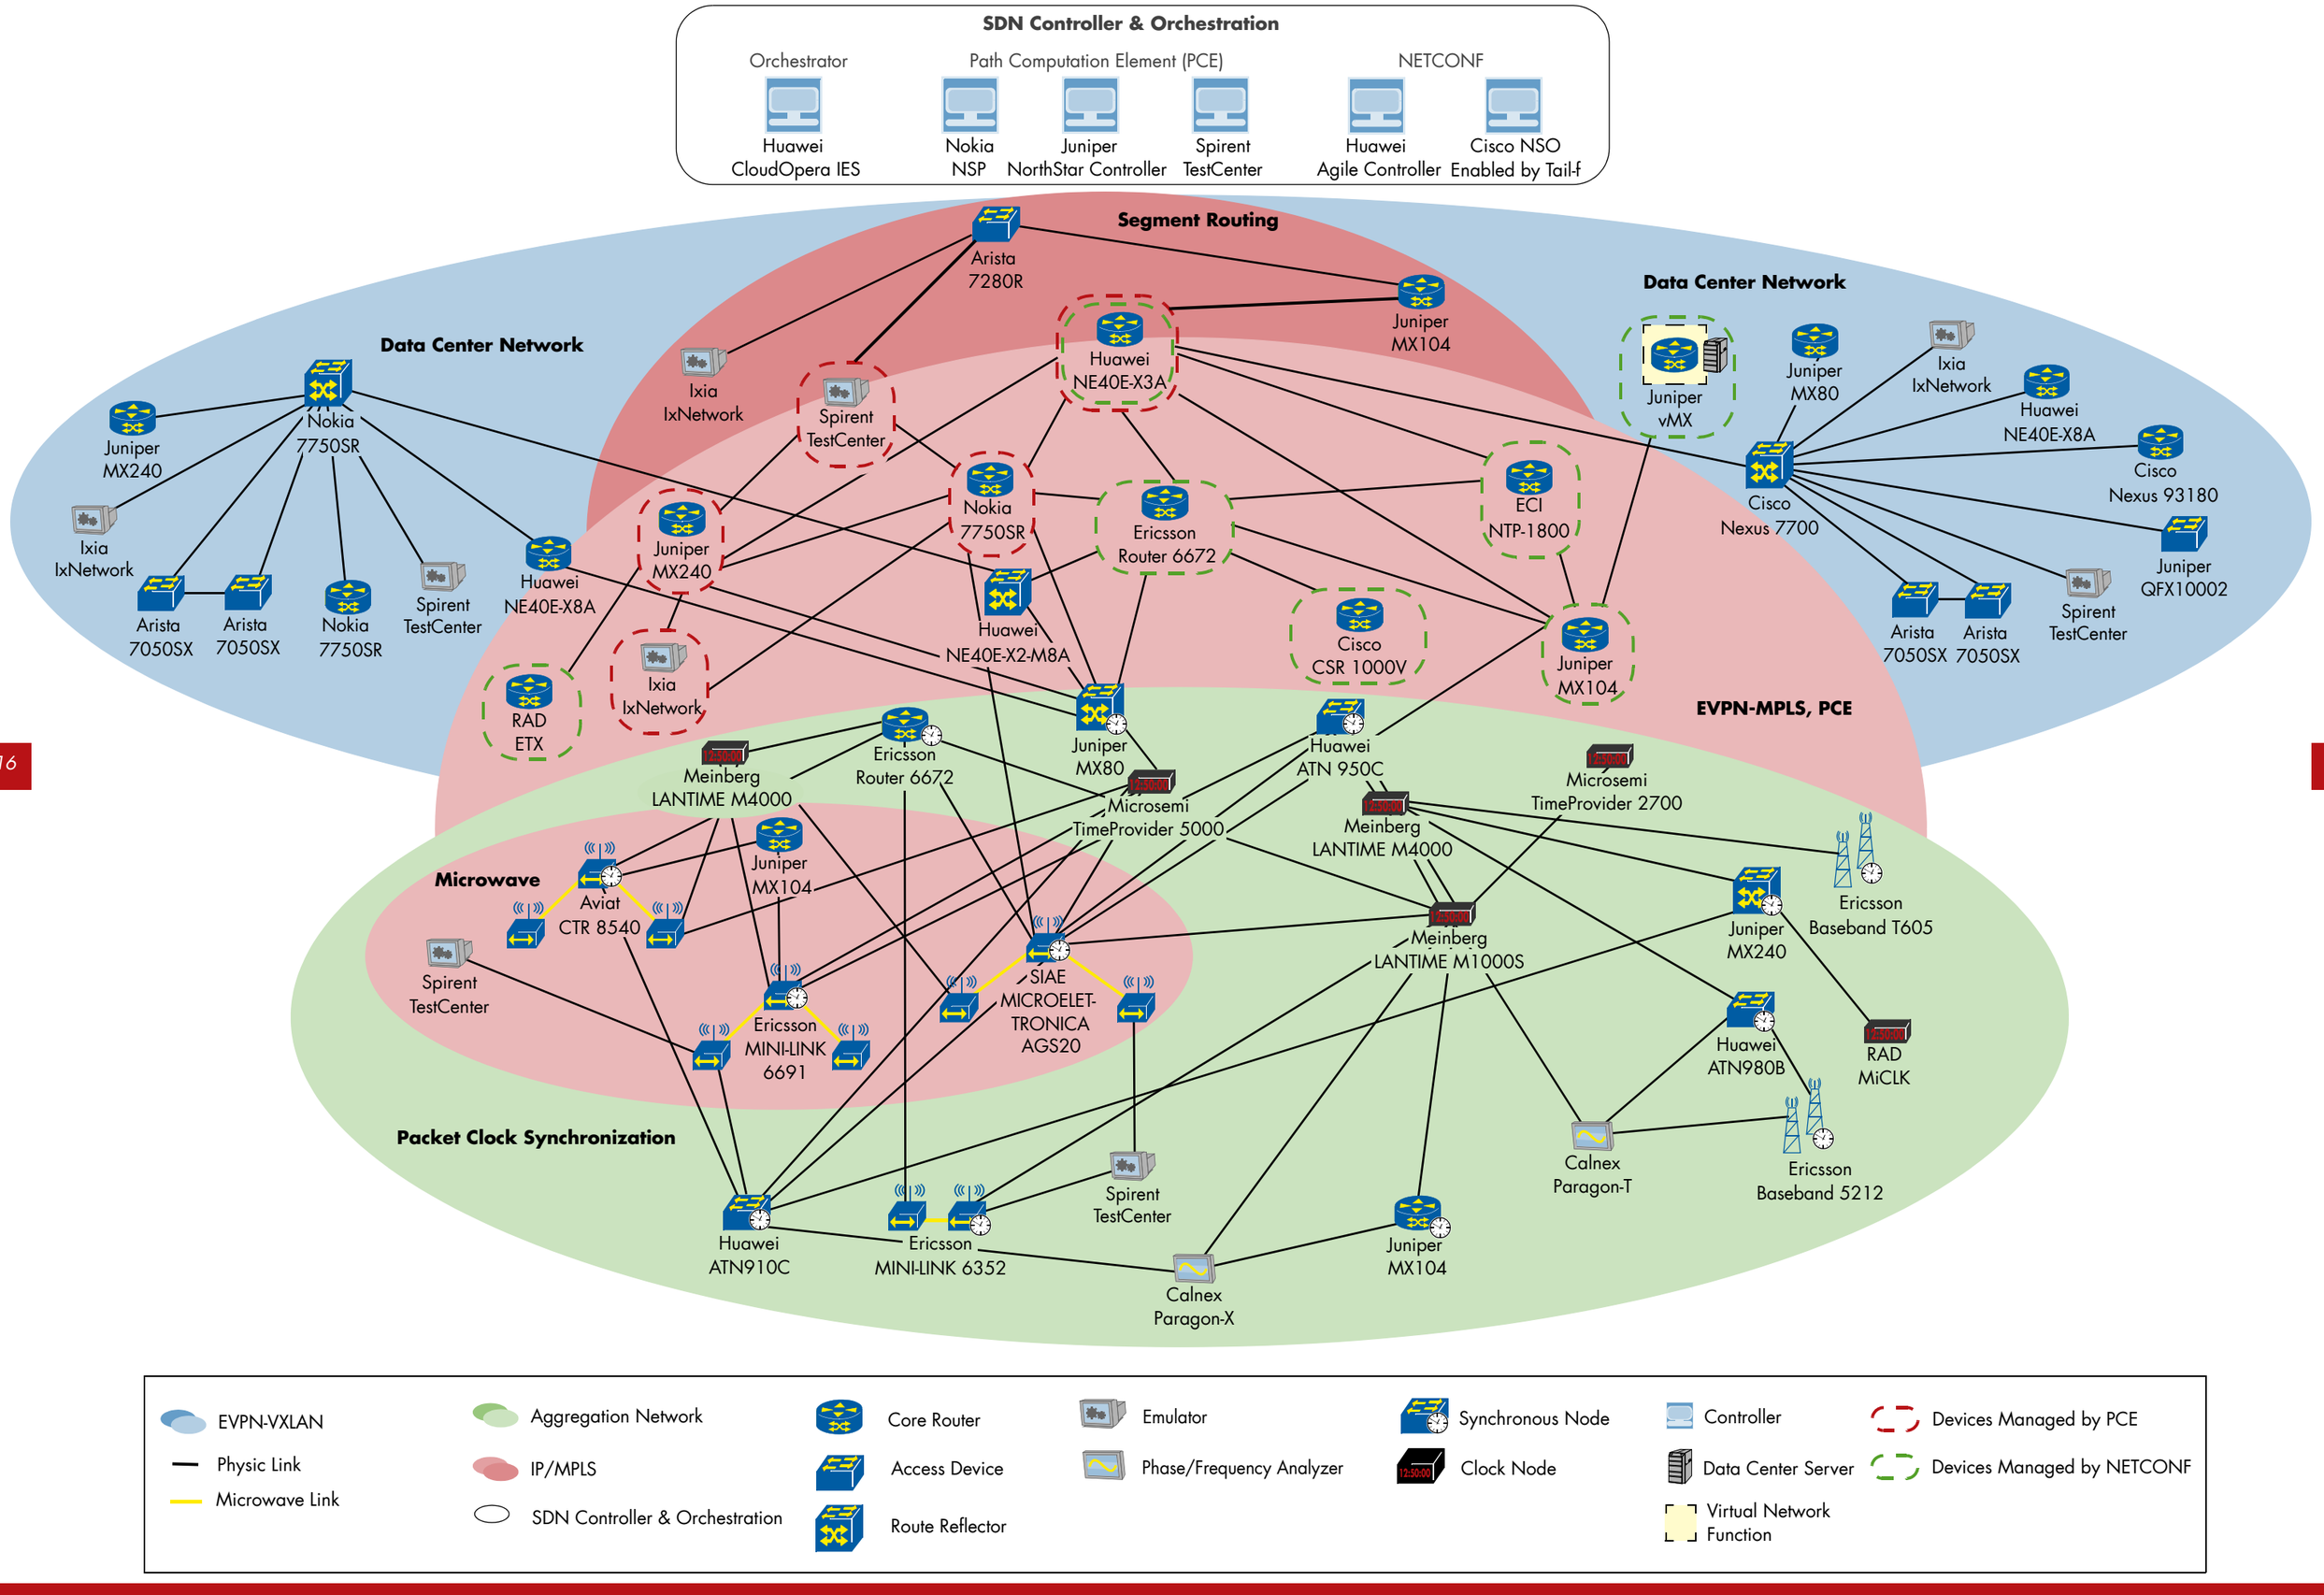
SDN Controller & Orchestration
Orchestrator
Huawei CloudOpera IES
Path Computation Element (PCE)
Nokia NSP
Juniper NorthStar Controller
Spirent TestCenter
NETCONF
Huawei Agile Controller
Cisco NSO Enabled by Tail-f
Segment Routing
Data Center Network
Arista 7280R
Juniper MX104
Ixia IxNetwork
Spirent TestCenter
Huawei NE40E-X3A
Juniper MX240
Nokia 7750SR
Ixia IxNetwork
Arista 7050SX
Arista 7050SX
Nokia 7750SR
Spirent TestCenter
Huawei NE40E-X8A
Juniper MX240
Nokia 7750SR
Ericsson Router 6672
ECI NTP-1800
Huawei NE40E-X2-M8A
Cisco CSR 1000V
Juniper MX104
RAD ETX
Ixia IxNetwork
Juniper MX80
Ericsson Router 6672
Meinberg LANTIME M4000
Microsemi TimeProvider 5000
Huawei ATN 950C
Microsemi TimeProvider 2700
Meinberg LANTIME M4000
Microwave
Aviat CTR 8540
Juniper MX104
Spirent TestCenter
Ericsson MINI-LINK 6691
SIAE MICROELET-TRONICA AGS20
Meinberg LANTIME M1000S
Juniper MX240
Ericsson Baseband T605
Huawei ATN980B
RAD MiCLK
Packet Clock Synchronization
Calnex Paragon-T
Ericsson Baseband 5212
Huawei ATN910C
Ericsson MINI-LINK 6352
Spirent TestCenter
Juniper MX104
Calnex Paragon-X
Data Center Network
Juniper vMX
Juniper MX80
Ixia IxNetwork
Huawei NE40E-X8A
Cisco Nexus 93180
Cisco Nexus 7700
Juniper QFX10002
Arista 7050SX
Arista 7050SX
Spirent TestCenter
EVPN-MPLS, PCE
EVPN-VXLAN
Physic Link
Microwave Link
Aggregation Network
IP/MPLS
SDN Controller & Orchestration
Core Router
Access Device
Route Reflector
Emulator
Phase/Frequency Analyzer
Synchronous Node
Clock Node
Controller
Data Center Server
Virtual Network Function
Devices Managed by PCE
Devices Managed by NETCONF

In the third test setup Huawei Agile Controller acted as NETCONF client and was involved to configure the services. The emulator (postman) was set to simulate the northbound communication to the Agile Controller using RESTCONF. Ericsson Router 6672 and Huawei NE40E-X3A acted as NETCONF server.

In one case the device configuration and the controller config database could not be synchronized using NETCONF/YANG.

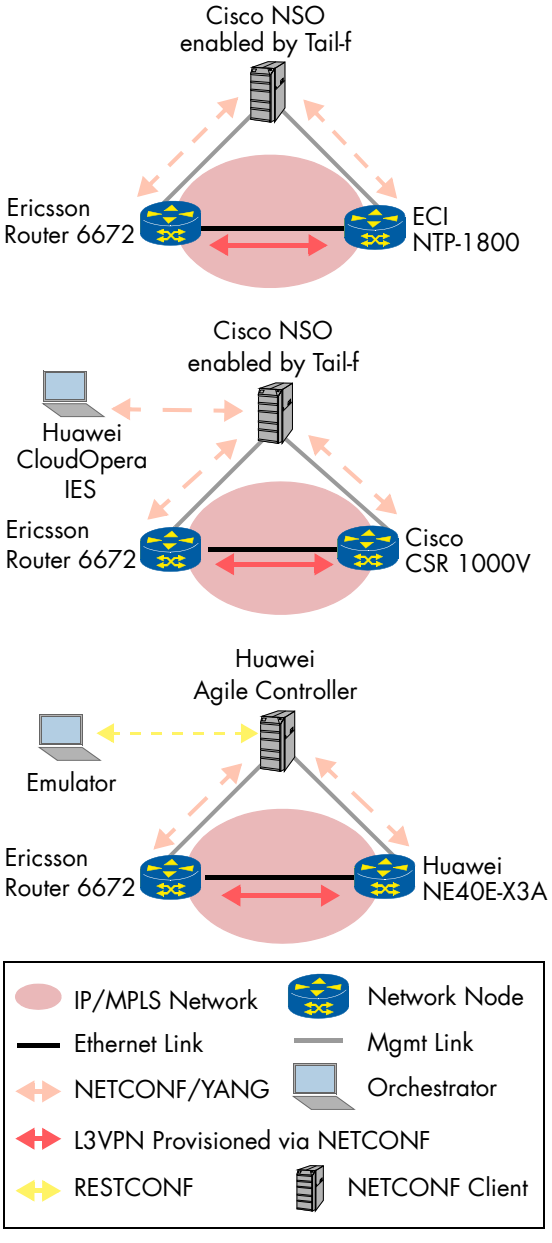

Figure 19: Test Setup 3— L3VPN Service Creation Using NETCONF/YANG

Virtual CPE Integration in WAN SDN Network. Virtual customer premises equipment (vCPE) aims at providing efficient service delivery capability through CE virtualization and is especially beneficial in virtual Private Cloud (vPC).

In this test we verified the integration of a vCPE in a WAN SDN network. The use case of the vCPE was an L3VPN deployment. As the standardization of the YANG model still remains outstanding (the latest draft "BGP IP VPN Virtual CE" remains unfinished) we used a proprietary YANG model.

We proved the IGP and RSVP-TE statuses on the network nodes in the WAN. Then we verified the NETCONF sessions between the controller and the CEs as well as PEs.

Later we tested that the SDN controller could retrieve the network topology. The orchestrator requested the creation of the L3VPN service. Upon reception of the request, the SDN controller created the service instance and implementation and sent it to the network nodes and vCPE. We tested that the PE-CE static routing was successfully established between the vCPE and the PE. Traffic was generated to check the correctness of the service and of the configurations. Finally we verified the termination of CE-PE static routing.

Four participants were successfully involved in two test setups.

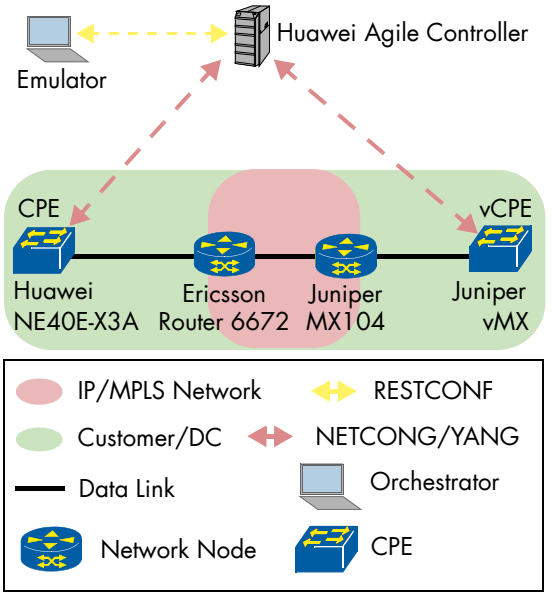

Figure 20: Test Setup 1 — Virtual CPE Integration in WAN SDN Network

In the first test setup Huawei Agile Controller acted as NETCONF client to configure the services. The emulator (postman) was used to simulate the northbound communication to the Agile Controller using the RESTCONF protocol. Huawei NE40E-X3A acted as CPE, Juniper MX104 and Ericsson Router 6672 joined as PE and Juniper Virtual MX acted as vCPE. NETCONF was used to configure the static PE-CE routing. Spirent participated as traffic generator. The vCPE (Juniper vMX) was hosted on a server.

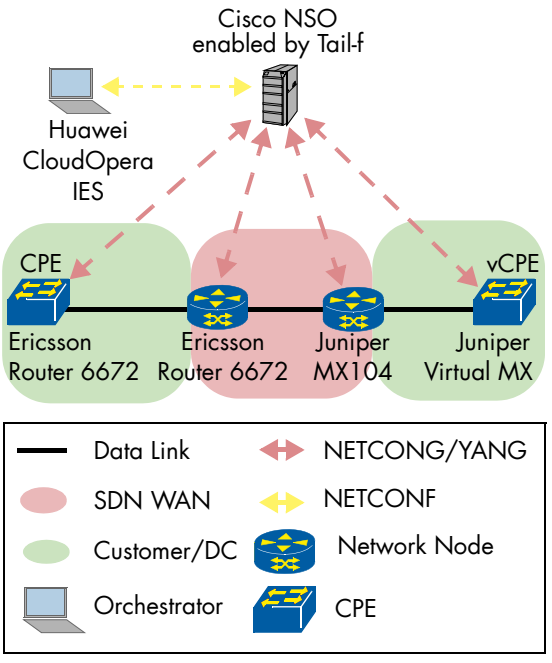

Figure 21: Test Setup 2— Virtual CPE Integration in WAN SDN Network

In the second test setup Cisco NSO enabled by Tail-f acted as NETCONF client and Huawei CloudOpera IES was involved as the orchestrator interacting with the Northbound API of Cisco NSO using the NETCONF protocol. Ericsson Router 6672 acted as CPE, Juniper MX104 and Ericsson Router 6672 acted as PE and Juniper Virtual MX was the vCPE. NETCONF was used to configure the static PE-CE routing and the L3VPN on the PEs. Spirent participated as traffic generator. The vCPE (Juniper vMX) was hosted on a server.

L2 Service Performance Monitoring using NETCONF/YANG. The ITU-T specification Y.1731 introduces three types which are defined to measure frame loss, frame delay and frame delay variation for a point-to-point EVC. Since 2011, the standard also added the synthetic frame loss measurement of a multipoint-to-multipoint EVC.

In the test, we used different CoS services per EVC to verify the functionalities by introducing different impairment profiles to the CoS services. These impairment profiles included: 10% and 15% frame loss - for two different CoS services; and 10 ms and 20 ms delay for two different CoS services; then 5 ms followed by 15 ms delay between every two frames within the CoS service which acted as delay variation. We confirmed that the network device launched the measurements and retrieved the values as defined in the above impairment profile. In addition, we expected to read and display the measured values via the NETCONF/YANG controller.

The following DUTs successfully participated in the test: Cisco NSO enabled by Tail-F and RAD ETX.

We successfully measured the frame loss (ETH-LM), the two-way delay measurement (ETH-DM) and the two-way delay variation measurement (ETH-DM) from the performance monitoring tool-kits (based on the ITU-T Y.1731) per point-to-point EVC and CoS ID. In addition, we also successfully measured the synthetic frame loss (ETH-SLM) using a multi-point-to-multipoint EVC between the RAD ETX. We used Calnex Paragon-X as impairment tool. The Cisco NSO enabled by Tail-f which acted as NETCONF controller, rightly read and displayed the measured values via NETCONF/YANG. We used Ixia IxNetwork and Spirent TestCenter to generate Ethernet traffic.

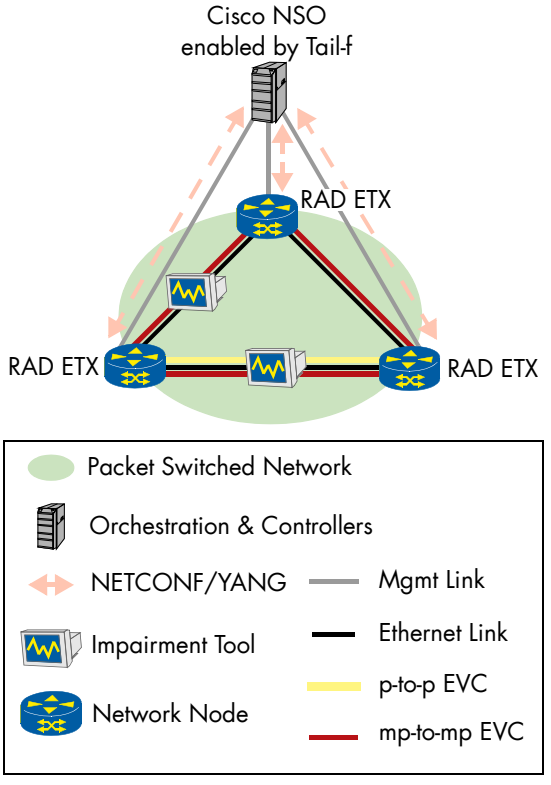

Figure 22: L2 Service Performance Monitoring Using NETCONF/YANG

L2 Service Activation using NETCONF. This test showed the service activation of Ethernet Virtual Connections (EVC) using NETCONF. During the test two EVPL services were configured within the service provider network. Service parameters for each service were also set, including CIR/CBS, EIR/EBS. NETCONF/YANG was used to configure CE, RE and maintenance End Points (MEP) to run Y.1564 SAT.

An access port was established on NID of the first DUT acting as generator that could non-disruptively drop and insert traffic into the DUT. The other port was plugged in to the customer edge and was configured to carry the live traffic. The second DUT acting as responder was set with the Latching Loopback to loop the traffic back for the specified VLAN services. We saw that Latching Loopback was activated on LLD and then we verified that CIR/EIR and CBS/EBS were correctly configured using NETCONF. We confirmed that the network configurations were stable with no issues. Finally, we confirmed that all service were within their SAC performance level.

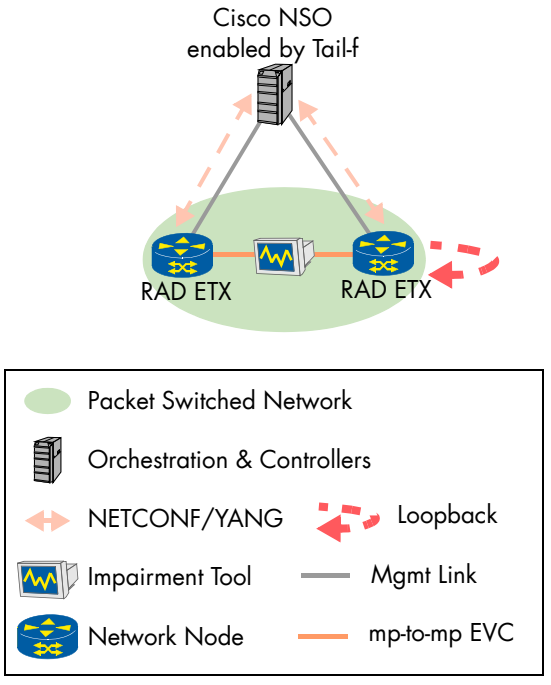

Figure 23: L2 Service Activation Using NETCONF

Cisco and RAD successfully participated in this test. Cisco NSO enabled by Tail-f acted as the NETFCONF Client and RAD ETX acted as NETCONF servers in the generator and responder roles.

Path Computation Element Protocol

The Path Computation Element Protocol (PCEP) was defined in RFC 5440 and specifies the communication between a Path Computation Client (PCC) and a Path Computation Element (PCE). Other drafts were developed to add stateful PCEP, PCE initiated LSP as well as extension to support Segment Routing LSP. A PCE is an entity (component, application or network node) that is capable of computing a network path or route based on a network graph and applying computational constraints. A PCC is a client application that will request a path computation to be performed by a Path Computation Element. The PCE controller receives knowledge of the network topology via the Traffic Engineering Database (TED) and of the previously established paths via the LSP database (LSPDB). The Path Computation Element Protocol facilitates the deployment of Software Defined Networks.

The LSP can be either initiated from a PCC or the PCE. When initiated from the PCC the LSP can later be delegated to the PCE which was covered in the first test of this section. The second test of this section focused on Segment Routing LSPs creation using the PCEP protocol with the LSPs being initiated from the PCE.

PCC-initiated Paths in a Stateful PCE Model. In the following setup we verified that the PCC could initiate an LSP and delegate it to the PCE. We also tested modification and revocation of the PCC-initiated LSP. In this test, the PCCs sent the Path Computation (PC) request to the PCE and the PCE computed the LSP path and sent the PC reply to the PCCs. Then a L3VPN service was established through the RSVP tunnel along the LSP path. LSP delegation to the PCE was granted from the PCC. Upon reception of the request, the PCE accepted the delegation of the LSP. We also tested modification and revocation of the PCC-initiated LSP. In both directions a new transport path was installed and the L3VPN service was established through it. Bidirectional traffic was sent through the L3VPN services established between all PCCs. Then we confirmed that the PCCs initiated the revocation of the LSP and cleared the LSP state provided by the PCE.

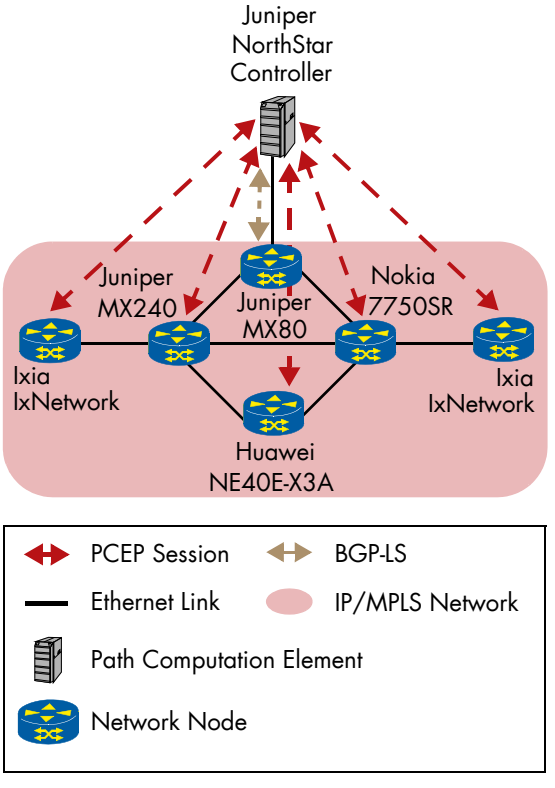

Figure 24: Setup 1- PCC-initiated Paths in a Stateful PCE Model

Five participants were successfully involved in three test setups.

In the first test setup Juniper NorthStar Controller acted as a PCE, Juniper MX240, Nokia 7750SR, Huawei NE40E-X3A and two Ixia IxNetwork acted as PCCs and Juniper MX80 acted as transit Node. Ixia was used as traffic generator. The PCE did not support Path Computation (PC) requests and reply, therefore the PCCs initiated the LSP path first and then the PCCs delegated the LSP to the PCE. L3VPN services were configured between all routers acting as PCCs. One PCC did not manage to accept the path update triggered by the PCE.

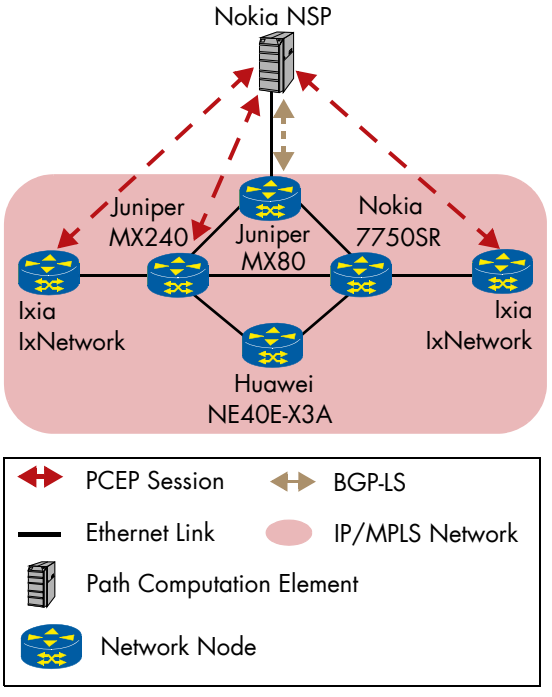

Figure 25: Test Setup 2 — PCC-initiated Paths in a Stateful PCE Model

In the second test setup Nokia NSP acted as PCE, two Ixia IxNetwork and Juniper MX240 acted as PCCs and Juniper MX80 and Huawei NE40E-X3A acted as transit nodes. L3VPN services were configured between all routers joining as PCCs. One PCC did not manage to accept the path update triggered by the PCE.

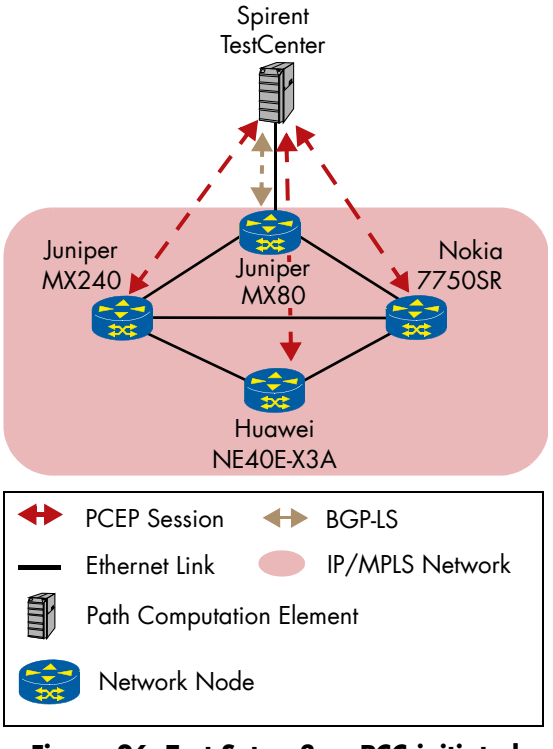

Figure 26: Test Setup 3 — PCC-initiated Paths in a Stateful PCE Model

In the third test setup Spirent TestCenter acted as PCE, Juniper MX240, Nokia 7750SR and Huawei NE40E-X3A acted as PCCs and Juniper MX80 acted as transit Node. L3VPNs services were configured between all routers functioning as PCCs. Spirent was used as traffic generator. One PCC did not manage to accept the revocation triggered by the PCE.

In one scenario the PCC sent the PCReq and the PCE responded with a PCRep message including a No-Path object (cannot find a path that meets constraints). The PCE showed that RSVP was not enabled on a transit link between the LSP headend and tailend,, however RSVP was actually enabled on this link. Inside the BGP-LS update in the PCE for this link was an invalid attribute at the start of the Link Attribute TLVs. The initial assessment was that this was potentially causing this issue.

In a second scenario, one PCC did not manage to accept the path update triggered by the PCE. When the PCE sent a PCUpd to this PCC to reroute the LSP, the PCC recorded an error of "bad strict route" and sent a PCRpt with RSVP Error Code 0/0, even though the ERO appeared correct.

PCE-initiated Segment Routing Paths. The draft "PCEP extension for Segment Routing" by the IETF defines how the Segment Routing Explicit Route Object (SR-ERO) can be used to carry a segment routing path. In this test we verified the setup of end-to-end service using a standard service control plane, while the transport is derived using segment routing without utilizing hop-by-hop signaling techniques (LDP or RSVP-TE). The segment routing path is derived from a PCE controller. The PCE controller learns the network topology via the Traffic Engineering Database (TED) and previous established paths via LSP database.

We verified that the IGP information was correctly exchanged between the nodes and the PCEP session was established successfully between PCE and PCCs. We then checked that the PCE correctly computed the paths, using a shortest path scheme first and that the paths were correctly installed on the PCC nodes. Later we saw that the PCE could change the paths, using an explicit path different from the shortest one. Bidirectional traffic was sent through the L3VPN services established between all PCCs.

Five participants were successfully involved in three test setups.

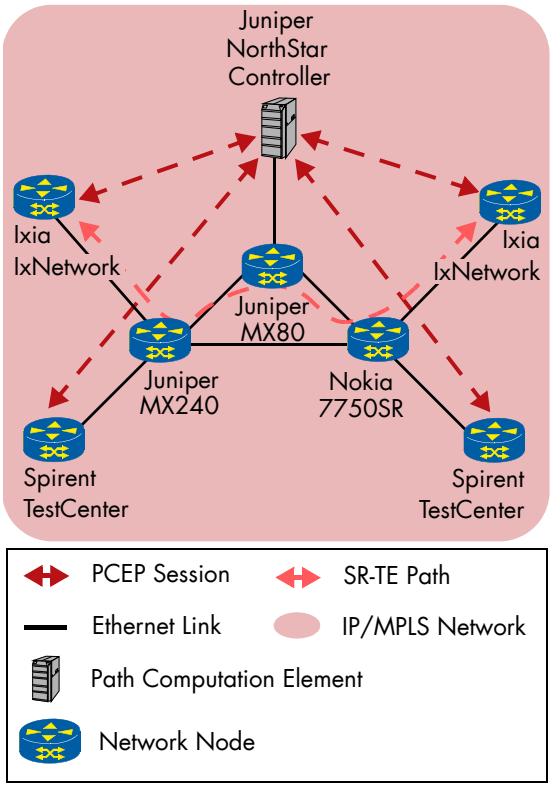

Figure 27: Test Setup 1 - PCE-initiated Segment Routing Paths

In the first setup Juniper NorthStar Controller acted as PCE, two Ixia IxNetwork and two Spirent TestCenter acted as PCCs and Juniper MX80, Juniper MX240 and Nokia 7750SR acted as transit Node. Spirent and Ixia were used as traffic generator.

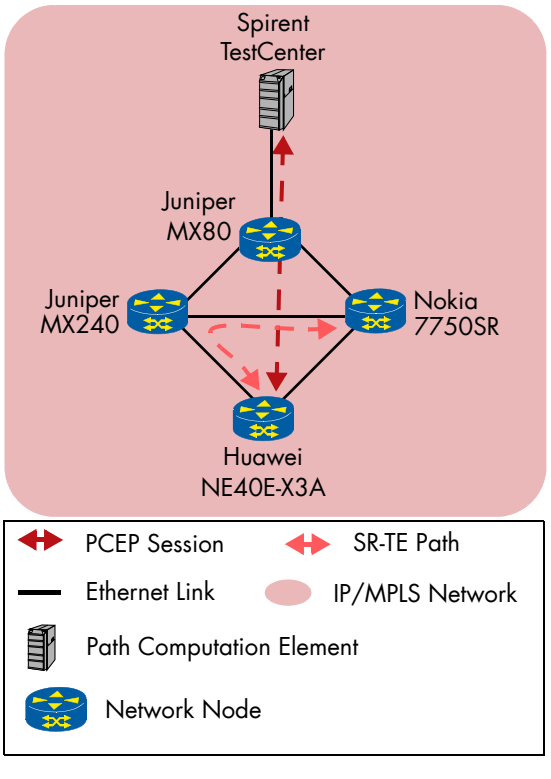

Figure 28: Test Setup 2 - PCE-initiated Segment Routing Paths

In the second test setup Spirent TestCenter acted as PCE, Huawei NE40E-X3A acted as PCC and Juniper MX80 Juniper MX240 and Nokia 7750SR acted as transit Node. In this setup the LSP were PCC-initiated. Spirent was used as traffic generator.

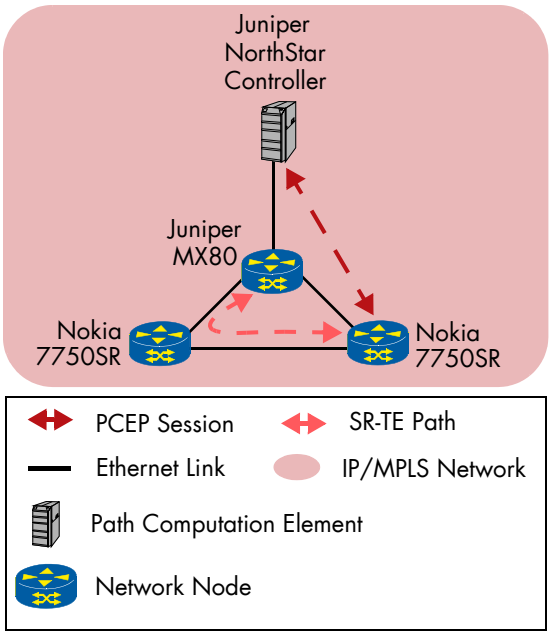

Figure 29: Test Setup 3 - PCE-initiated Segment Routing Paths

In the third scenario Juniper NorthStar Controller acted as PCE, Nokia 7750SR acted as PCC and Juniper MX80 and Nokia 7750SR acted as transit Node. L3VPNs services were configured between all routers acting as PCCs. Ixia was used as traffic generator.

In one setup we noticed multiple interoperability issues. One PCC expected the P flag enabled in the objects of PCInitiate message, even though the P flag is meant for the PCReq message and the draft does not specify anything for PCEInitiate messages. One PCC did not support PCEInitiated LSP. The delegation of a PCC-initiated LSP did not work either as it was not having any path information. One PCC did not support the latest SR-ERO sub-object type 36.

Clock Synchronization

The vendors participating in the synchronization part this year were really eager for multi-vendor scenarios; we tested almost 70 different combinations. We found several interoperability issues and small errors in the implementations, but all vendors' R&D departments reacted very fast to any issue found and provided quick fixes which allowed us to keep on testing at high speed.

Last year the ITU-T published the new G.8275.2 standard that defines a telecom profile for phase/time synchronization with partial timing support from the network, i.e. when not all nodes are PTP aware.

Besides the functionality tests for the new profile, our clock synchronization event this time was focused on the performance measurement, using full and assisted partial timing support scenarios, or a combination of the two.

We tested the behavior of the time signal delivery in optimal and suboptimal conditions: delay asymmetry, hold-over performances, source failover between two grandmaster clocks.

While there are no agreed standards yet on the accuracy requirements, 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. For frequency synchronization, we continued using the G.823 SEC mask as a requirement.

The primary reference time clock (PRTC) was GPS using an L1 antenna located on the roof of our lab.

Phase/Time Partial Timing Support

In this test we verified if a slave clock could lock to the boundary or grandmaster clock when the PTP telecom profile for time/phase synchronization with partial timing support from the network (ITU-T G.8275.2) is used without any physical frequency reference – such as SyncE.

In test setup 1 only the grandmaster and slave clocks were present, while in test setup 2 a boundary clock participated as well.

The slave and boundary clocks started from a free running condition and no PTP configured. Following, we enabled PTP while the impairment was emulating a path lacking PTP support i.e. PDV according to the profile defined in G.8261 test case 12. After conforming that the slave clock was able to lock within 30 minutes, we used the Calnex Paragon-X to verify that phase accuracy requirement of ±1.1 µs and frequency requirements were satisfied.

We successfully tested the following combinations: in test setup 1 Meinberg LANTIME M4000 and Microsemi TimeProvider 5000 participated as grandmaster clock, Ericsson Router 6672, Huawei ATN910C and Huawei ATN950C as slave clock.

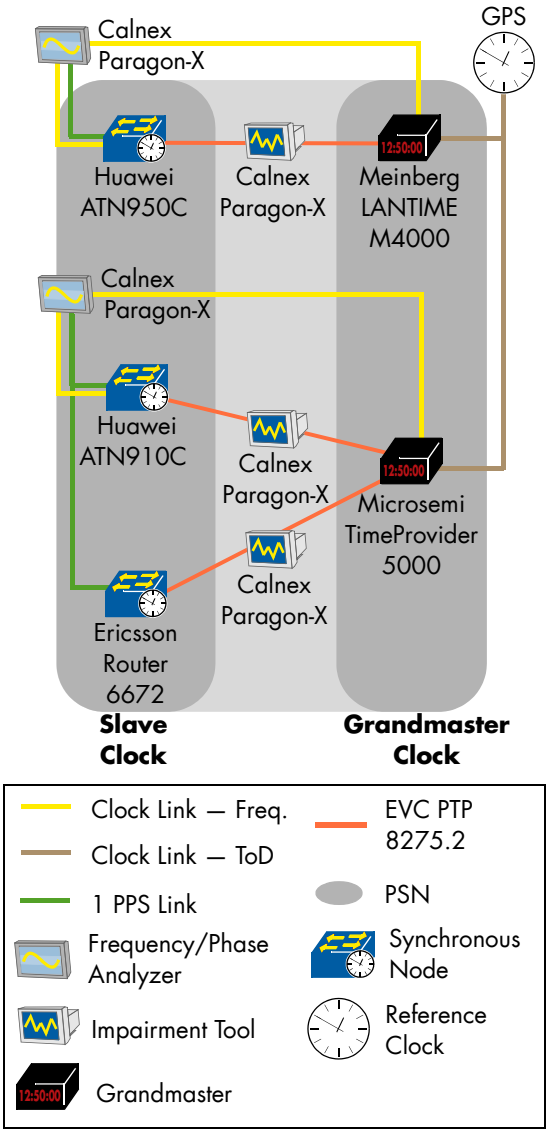

Figure 30: Test Setup 1 — Phase/Time Partial Timing Support

In test setup 2 Microsemi TimeProvider 5000 participated as grandmaster clock, Microsemi TimeProvider 2700 as boundary clock, and Meinberg LANTIME M1000S as slave clock.

In one of the tested combinations in test setup 2 the boundary clock did conversion between G.8275.2 used in the uplink with the grandmaster and G.8275.1 in the downstream to the slave. During this test we discovered that the downstream PTP packets were malformed and not accepted by the slave clock.

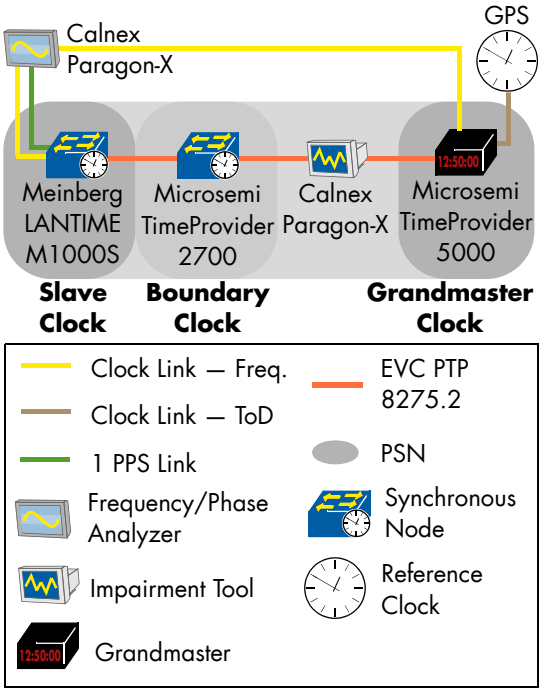

Figure 31: Test Setup 2— Phase/Time Partial Timing Support

Phase/Time Assisted Partial Timing Support

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 boundary and slave clocks.

We began with both the grandmaster and boundary clocks connected to GPS and the slave clock locked via PTP to the boundary clock. We then started the impairment which emulated a PDV according to the profile defined in G.8261 test case 12. Upon disconnecting the GPS from the T-BC-P we verified that it started using PTP. We performed the measurement using the Calnex Paragon-X and confirmed that the output met the phase accuracy requirement of ±1.1 µs and frequency requirements.

In the last step we reconnected the GPS antenna to the T-BC-P and repeated the measurement.

Meinberg LANTIME M4000 and Microsemi TimeProvider 5000 successfully participated as grandmaster clock, Ericsson Router 6672 and Meinberg LANTIME M1000S as boundary clock, and Huawei ATN950C, Juniper MX104 and Meinberg LANTIME M1000S as slave clock.

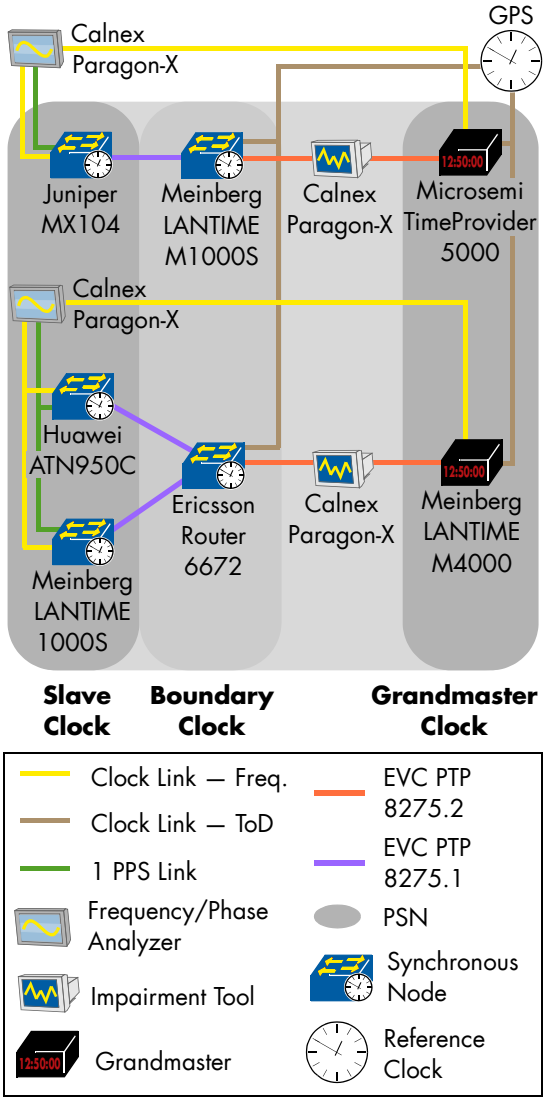

Figure 32: Phase/Time Assisted Partial Timing Support

In one of the test combinations, we observed that the T-BC-P was overwriting the grandmaster clock's PTP data set information downstream to the slave. According to the standard, the boundary clock should let the slave know the identity of the T-GM and its parameters.

After disconnecting the GPS from the boundary clock, we observed that the slave went in holdover status as it did not accept the combination of clock-Class 7 and clockAccuracy "unknown", when the timeTraceable flag was set to false.

In another test combination, we observed that the T-BC-P did not set the PTP_UNICAST flag as required by G.8275.2 and thus the PTP was not established between T-BC-P and slave clock.

Both problems were solved by the R&D departments of the vendors during the hot staging.

Phase/Time Assisted Partial Timing Support: Delay Asymmetry

This test was performed using the ITU-T G.8275.2 profile in the upstream between boundary and grandmaster clock and ITU-T G.8275.1 in the downstream link between boundary and slave clock.

Both the grandmaster and boundary clocks were connected to GPS. We started the impairment emulating a PDV according to the profile defined in G.8261 test case 12.

After disconnecting the GPS from the boundary clock, we used the Calnex Paragon-X to introduce a delay asymmetry of 125 µs and verified that the boundary could calculate and compensate the asymmetry introduced.

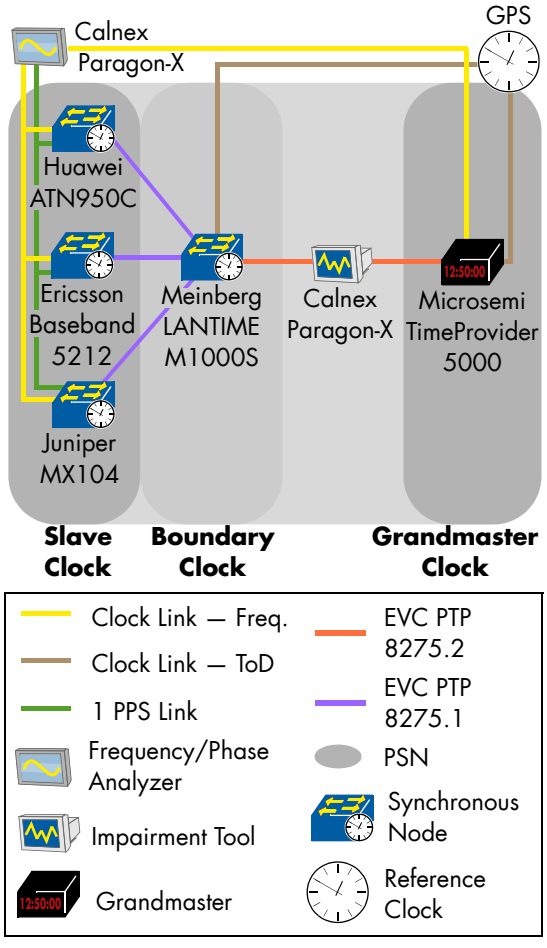

Figure 33: Phase/Time APTS: Delay Asymmetry - Scenario 1

Microsemi TimeProvider 5000 participated as grandmaster clock, Meinberg LANTIME M1000S as boundary clock, Ericsson Baseband 5212, Huawei ATN950C and Juniper MX104 as slave clock.

All the combinations depicted in Figure 33 complied with the G.823 SEC mask and the ±1.1 µs absolute phase error requirements.

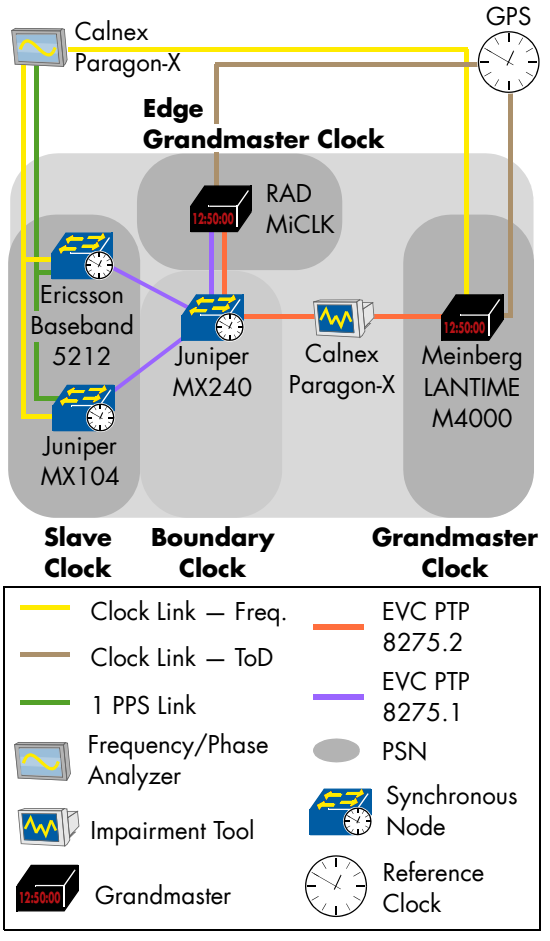

Figure 34: Phase/Time APTS: Delay Asymmetry - Scenario 2

In the second scenario, where the edge grandmaster was present, we first introduced the asymmetry using the Calnex Paragon-X and then disconnected the GPS antenna from the edge grandmaster clock. We observed that the edge grandmaster sent out to the boundary clock clock-Class 7.

The edge grandmaster did not support the display of the detected asymmetry, but the slaves kept their "lock" to the edge grandmaster and the measurement performed with the Calnex Paragon-X met the criteria set in the ITU-T G.823 SEC mask and the ±1.1µs absolute phase error requirement.

Meinberg LANTIME M4000 participated as grandmaster clock, RAD MiCLK as edge grandmaster clock, Juniper MX240 as boundary clock and Ericsson Baseband 5212 and Juniper MX104 as slave clock.

Hold Over Performance

In the following section we will highlight the two hold over performance tests. The hold over time is considered to be the longest period that a slave clock maintains the required accuracy. The measurements were performed over night with the evaluated duration of 12 hours.

Phase/Time Assisted Partial Timing Support: Hold Over Performance. This test was performed using the ITU-T G.8275.2 profile between the Grandmaster clock and boundary or Edge Grandmaster clock. The participants had the choice to use APTS or convert to the FTS profile in the downstream towards the slave clock.

We started the test case with allowing the boundary clock to lock to GPS. We then enabled PTP on the boundary and applied the PDV impairment according to the profile defined in G.8261 test case 12.

After disconnecting the GPS antenna from the boundary clock, we verified the holdover performance of a slave clock in relation to phase/time stability using the Calnex Paragon-T to perform the measurements.

Microsemi TimeProvider 5000 and Meinberg LANTIME M4000 participated as grandmaster clock, Meinberg LANTIME M1000S and Microsemi TimeProvider 2700 as boundary clock, Ericsson Baseband 5212, Juniper MX 104 and Meinberg LANTIME M1000S as slave clock.

All test runs depicted in Figure 35 complied with the ±1.1µs absolute phase error requirement as well as the G.823 SEC frequency mask.

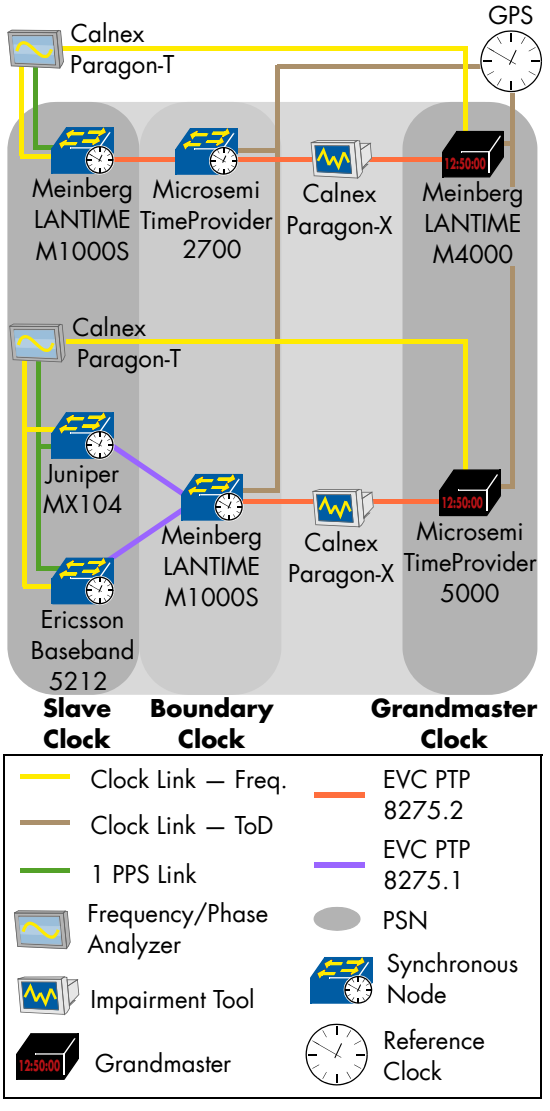

Figure 35: Phase/Time APTS: Hold Over Performance - Scenario 1

In the scenario featuring the edge grandmaster (scenario 2), we observed that the edge grandmaster advertised clockClass 7 upon losing the GPS signal. The measurement performed with the Calnex Paragon-T met the criteria set in the ITU-T G.823 SEC mask and the ±1.1µs absolute phase error requirement.

Meinberg LANTIME M4000 participated as grandmaster clock, RAD MiCLK as edge grandmaster clock, Juniper MX240 as boundary clock and Huawei ATN910C as slave clock.

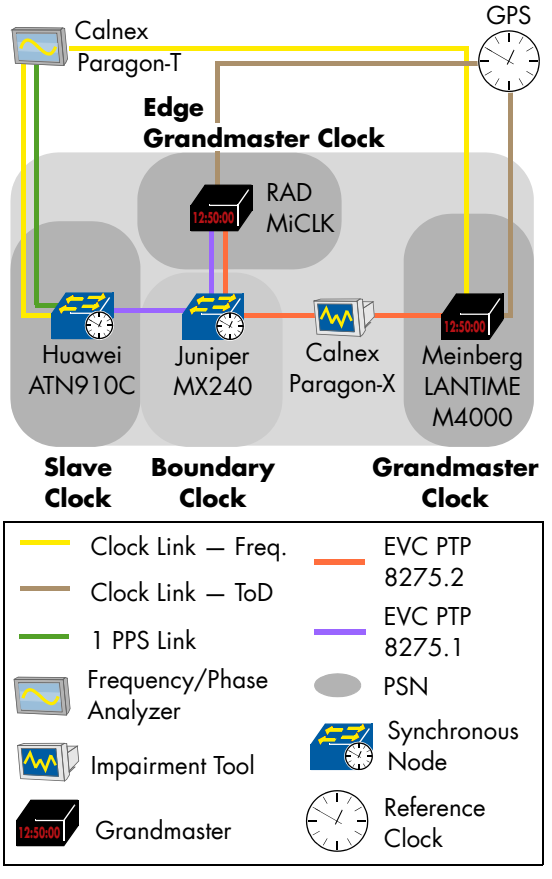

Figure 36: Phase/Time APTS: Hold Over Performance - Scenario 2

Frequency Synchronization: Hold Over Performance. In this test case the slave clock started in a free running setup, except for one test combination. Here the slave clock started in frequency holdover status, as it was not able to enter "free-run" on request nor after a complete node reload. We then enabled PTP and let the slave lock to the grandmaster clock. We emulated a PTP impairment and verified the holdover performance of a slave clock in relation to frequency stability.

We introduced the impairment by disconnecting the cable to the grandmaster and thus breaking the PTP signal. Calnex Paragon-T was used to perform the measurements.

After we verified the overnight measurement, we removed the impairment on the PTP stream and verified that the transient response of the slave clock matched the requirements after it re-locks to the grandmaster.

The following devices participated in this test: Meinberg LANTIME M4000 and Microsemi TimeProvide 5000 as grandmaster clock, Ericsson Baseband T605, Huawei ATN950C, Juniper MX104 and Meinberg LANTIME M100S as slave clock.

The measurement performed with the Calnex Paragon-T showed that Ericsson Baseband T605, Juniper MX104 and Meinberg LANTIME M1000S met the criteria set in the ITU-T G.823 SEC mask, while Huawei ATN950C met the ITU.T G.8263 mask criteria.

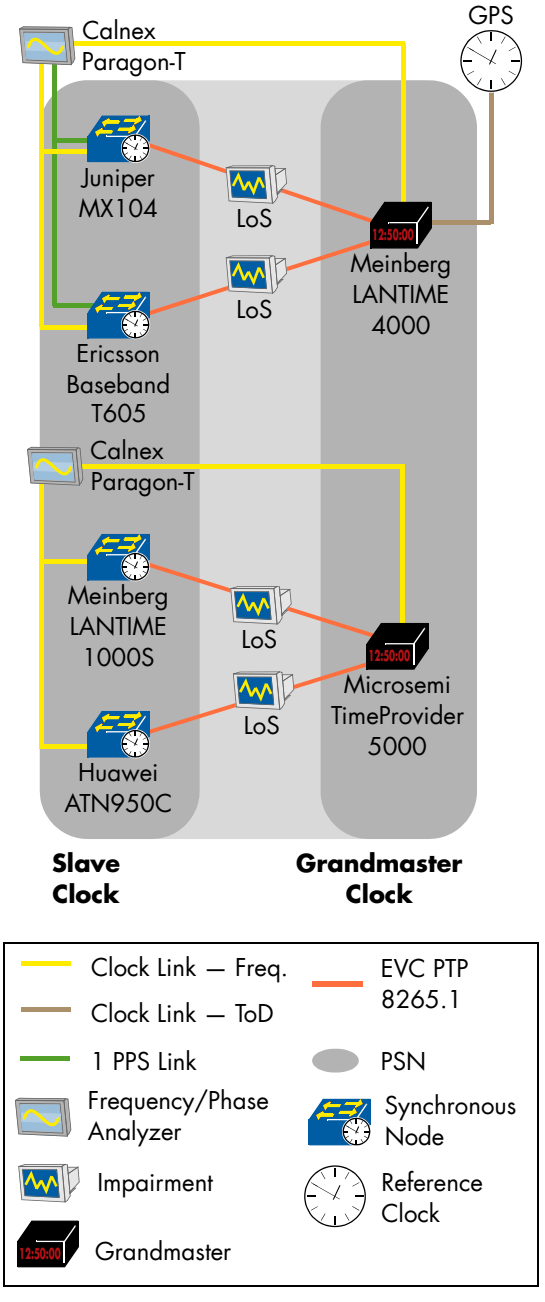

Figure 37: Frequency Synchronization: Hold Over Performance

Source Failover

In this setup, both grandmasters were provided with a GPS signal from a common GPS antenna. We allowed the slave clock to lock to the primary grandmaster and then degraded the primary grandmaster's quality by disconnecting its GPS input. We verified that the slave (boundary) clock switched over to the secondary grandmaster and measured the slave clock's transient response. We also tested if the correct clockClass values are being signalled by the grandmasters according to the telecom profiles, which allows the alternate best master clock algorithm running on the slave clock to correctly select the best grandmaster during each step of the tests.

We used the priority-2 field as tie-break parameter. One vendor's slave device choose the T-GM based on local priority.

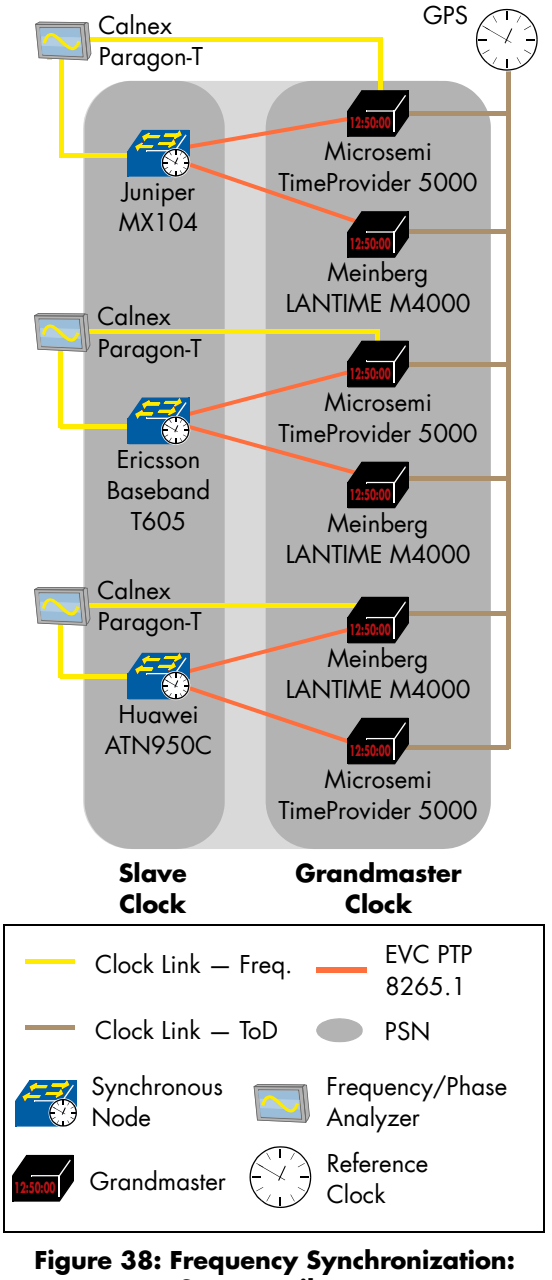

Figure 38: Frequency Synchronization: Source Failover

Frequency Synchronization: Source Failover. This test was performed with the G.8265.1 frequency profile.

The following devices successfully participated: Meinberg LANTIME M4000 and Microsemi TimeProvider 5000 joined as grandmaster clock, Ericsson Baseband T605, Huawei ATN950C and Juniper MX104 as slave clock.

Phase/Time Synchronization: Source Failover. For this test we used the ITU-T G.8275.1 profile on all devices.

The following solutions successfully participated in this test: Meinberg LANTIME M4000 and Microsemi TimeProvider 5000 as grandmaster clock, Ericsson Baseband T605, Ericsson Router 6672, Ericsson MINI-LINK 6691, SIAE MICROELETTRONICA AGS20 as boundary clock and Ericsson Baseband 5212, Ericsson MINI-LINK 6691, Huawei ATN950C, Huawei ATN980B, Juniper MX104 and Meinberg LANTIME M1000S as slave clock.

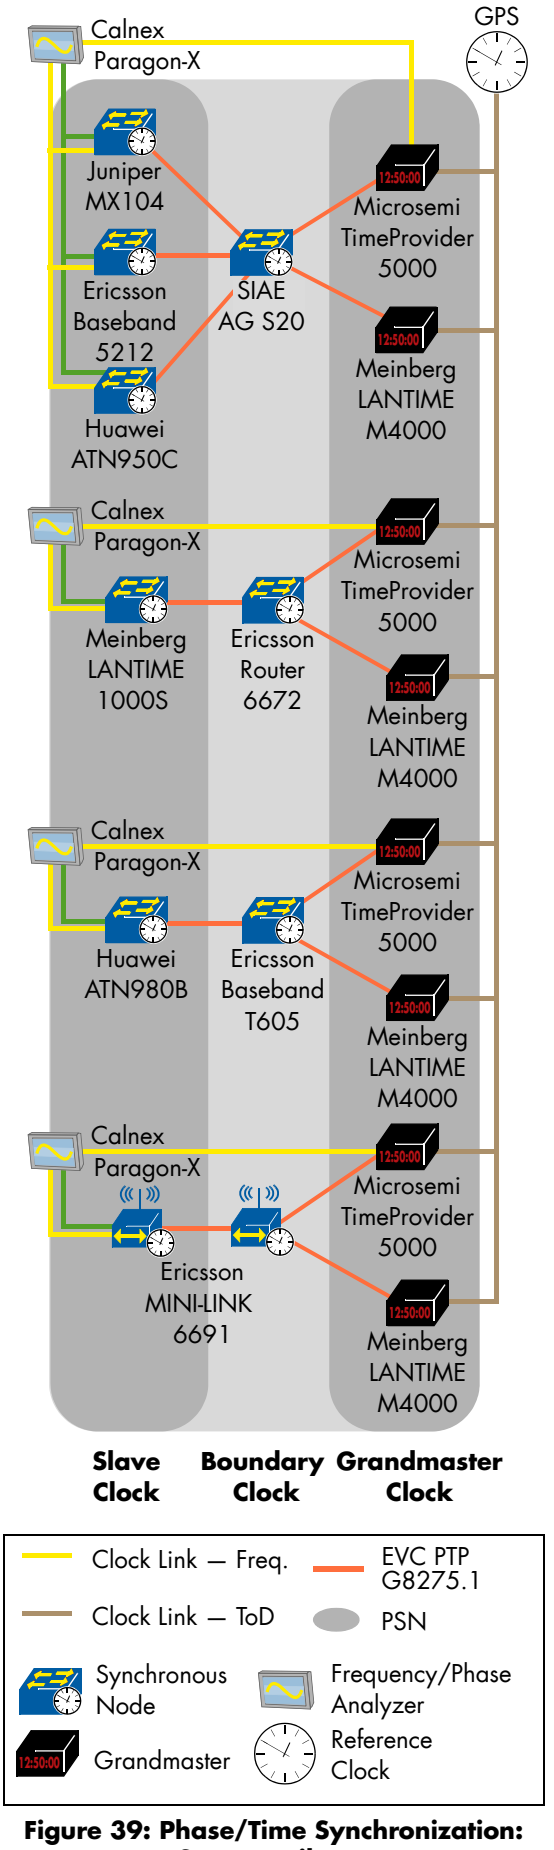

Figure 39: Phase/Time Synchronization: Source Failover

We observed one interoperability issue, when one grandmaster was sending clockClass 160 instead of 7 upon losing GPS. This problem was solved with a quick software upgrade during the hotstaging and we successfully progressed with the testing.

Phase/Time Synchronization with Full Timing Support

Microwave Transport. 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 for phase and frequency from the slave clock outputs. 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 retains the required quality level. Since the bandwidth decreased accordingly, we saw that data packets were dropped according to the available bandwidth.

In the first setup, the microwave stations acted as boundary clocks.

The following devices successfully participated in this test: Meinberg LANTIME M4000 and Microsemi TimeProvider 5000 as grandmaster clock, Aviat CTR 8540, Ericsson MINI-LINK 6691, SIAE MICROELETTRONICA AGS20 as boundary clocks, Aviat CTR 8540, Ericsson MINI-LINK 6691, Huawei ATN910C, Huawei ATN950C, Meinberg LANTIME M1000S as slave clocks.

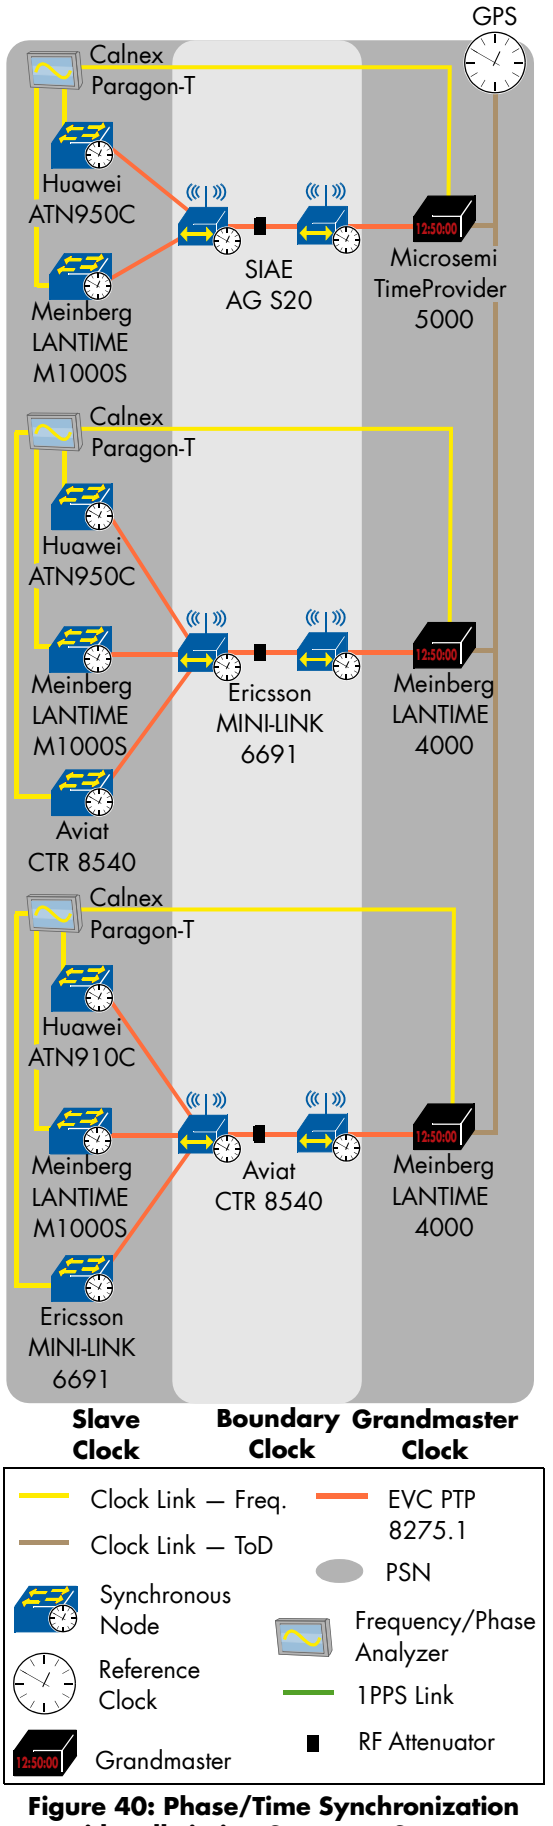

Figure 40: Phase/Time Synchronization with Full Timing Support - Setup 1

In the second setup, the microwave stations acted as transparent clocks.

Meinberg LANTIME M4000 and Microsemi TimeProvider 5000 participated as grandmaster clock, Aviat CTR 8540, Ericsson MINI-LINK 6352 and Ericsson MINI-LINK 6691 as transparent clocks, Huawei ATN910C, Huawei ATN980B, Meinberg LANTIME M1000S as slave clocks.

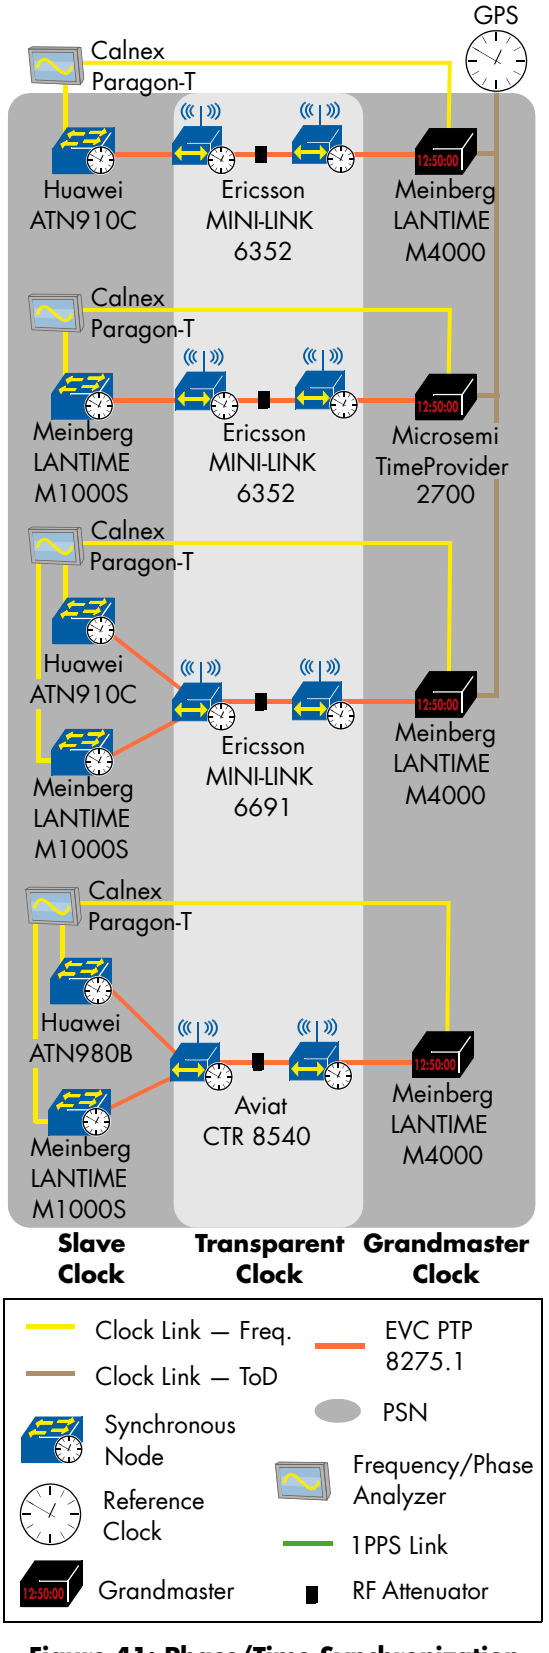

Figure 41: Phase/Time Synchronization with Full Timing Support - Setup 2

In one of the tested combination, during the sudden change of the modulation, we saw that the microwave system acting as boundary clock stopped sending the PTP packets to the downstream boundary clock. This caused a holdover state of the slave clock. Later the test was successfully retested by introducing the attenuation in the transmission channel more slowly.

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.

This test was performed using the ITU-T G.8275.1 profile.

Both, the grandmaster and one of the boundary clocks, were provided with a GPS signal. We allowed the grandmaster and the boundary clock to lock to GPS input. The boundary clock acted as primary grandmaster for the upstream boundary clock.

We then disconnected the antenna of the boundary clock to emulate a GPS failure and verified that both boundary clocks locked via PTP to the central grandmaster. In the last step, we recovered the GPS of the boundary clock and verified that the upstream boundary clock locked again to the downstream boundary clock.

The following devices successfully participated in this test: Microsemi TimeProvider 5000 as grandmaster clock, Ericsson Baseband T605 as boundary clock A and Huawei ATN950C as boundary clock B.

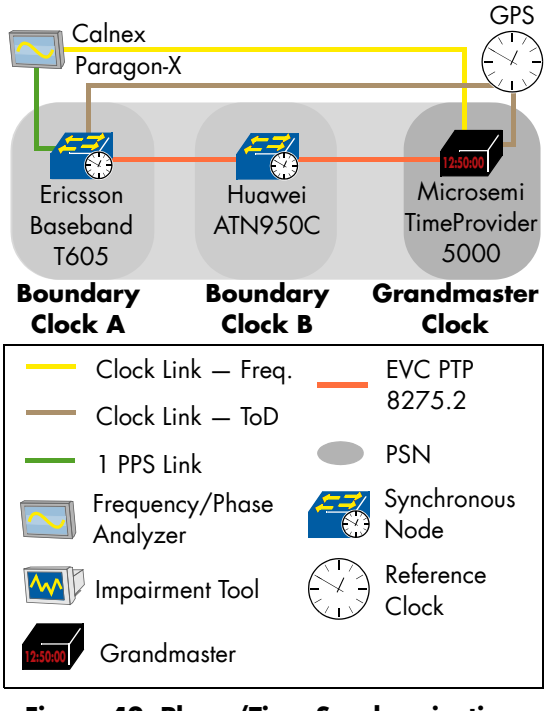

Figure 42: Phase/Time Synchronization: Degradation of Primary Source

In one of the test combinations, we observed a second shift when the clock reference changed from the downstream boundary clock to the grandmaster. The vendor assumed that this was a leap second problem.

In another test combination, we saw that one of the boundary clocks was sending unexpected value of 127 for *logMessagePeriod*.

On of the combinations was conducted to the end, but regarded as unsuccessful, as boundary clock B correctly switched reference from boundary clock A to the grandmaster clock and vice versa, but boundary clock A remained in hold over when missing its own GPS reference. Another test failed because boundary clock B did not manage to lock to boundary clock A.

Edge Grandmaster Clock. RAD proposed an extension of the test setups for the partial timing support test cases by adding an "Edge Grandmaster Clock" which is implemented as an SFP that can be hosted by another router.

Juniper MX240 hosted the RAD SFP in all tests that were performed with the edge grandmaster. In the link towards the grandmaster clock, the Juniper MX240 acted as a PTP not capable router, while in the downstream towards the slave, the Juniper MX240 acted as a boundary clock. In order to create different PTP EVCs, we used VLAN encapsulation on the interface used for the edge grandmaster clock.

An edge grandmaster clock can have one or more ports in slave mode that are working in a different PTP domain and possibly with a different PTP profile. Compared to a boundary clock the edge grandmaster does not forward the attributes of the upstream central grandmaster towards the slave clock, but rather sends its own attributes based on its current state.

We observed that under normal GNSS reception, it publishes clockClass 6 (T-GM connected to GNSS) towards the slave node. Upon losing its GNSS reception and falling back to APTS, it starts advertising clockClass 7 (T-GM in holdover, within holdover specification) toward the slave clocks, even if the signal can be tracked back to a grandmaster clock provided with GPS.

One combination failed because one slave clock did not accept the combination of clockClass 7 and clockAccuracy below 100 ns. ITU-T G.82751.1-201606 mandates that slaves must support all the values of clock_Accuracy upon reception (shall not discard) defined in the full IEEE 1588 range.

The problem was solved by the R&D department of the vendor during the hot staging by statically setting the accuracy value sent from the edge grandmaster to the slaves to 254 when the GPS signal is not available.

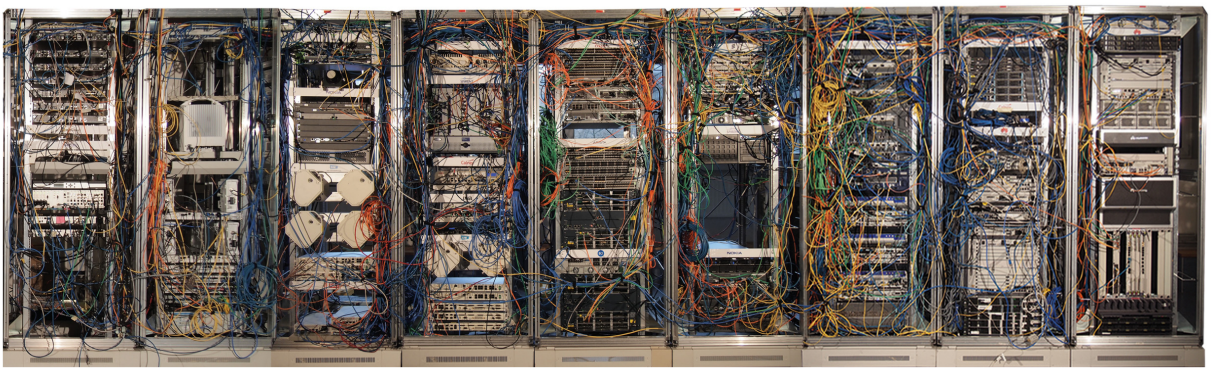

EANTC	upperside conferences
EANTC AG European Advanced Networking Test Center	Upperside Conferences
Salzufer 14 10587 Berlin, Germany Tel: +49 30 3180595-0 Fax: +49 30 3180595-10 info@eantc.de http://www.eantc.com	54 rue du Faubourg Saint Antoine 75012 Paris - France Tel: +33 1 53 46 63 80 Fax: + 33 1 53 46 63 85 info@upperside.fr http://www.upperside.fr

20170314 v2