EANTC

EUROPEAN ADVANCED NETWORKING TEST CENTER

Transport & Cloud Networks Interop Test Report

2026

Use Case Scenarios

EVPN

Orchestration and Automation

Segment Routing

Time Synchronization

Access the Full Test Report Online
https://eantc.de/interop26

UPPERSIDE WORLD CONGRESS
24/26 MARCH 2026 PARIS

Editor's Note

Carsten Rossenhoevel
Co-Founder & CTO

The 2026 edition of our multi-vendor interoperability test has been a massive undertaking of the **twelve leading vendors participating this year:** Arista Networks, Calnex Solutions, Ciena, Cisco Systems, Ericsson, HPE, Keysight, Microchip, Nokia, Raisecom, Ribbon, and ZTE – with the EANTC engineering team. Together, we have spent more than 1,600 person days (equivalent to seven person years) designing and implementing test cases with 56 device types, creating more than 1,300 result datasets. More than three tons of equipment were moved across the globe to facilitate this undertaking and are shown live in Paris.

Why invest all this effort? Because the transport network innovations that we evaluate – Segment Routing, SRv6, EVPN, Orchestration and Network Automation, and Time Synchronization – are the foundation of today's and tomorrow's Internet backbones, including mobile networks and data centers. Fortunately, thanks to the outstanding continued work by the IETF and other SDOs, these technologies are fully standardized and open. To ensure that multi-vendor networks remain a viable option for network operators, we validate the latest innovations with manufacturers in advance before they hit service provider proof of concept labs and enterprise acceptance procedures. Our joint testing **increases trust, accelerates innovation, and enables diverse supplier sourcing**. Digital sovereignty is built on these pillars: Enabling global alternatives, minimizing dependencies on single vendors or regions. In our tests, vendors from the Americas, Europe, and Asia are collaborating to implement open standards.

What are the specific innovations this year? Outside our lab, most operators are still in the process of deploying Segment Routing in the near future. In the test lab, we have declared interoperability success for many basic test areas in Segment Routing and EVPN years ago. All previous reports are available from our website. Nowadays, our protocol tests focus on closing the remaining functional, reliability, scalability, and network automation gaps: Advanced multicast EVPNs; multi-homing in SRv6; telemetry collections and digital twins; large-scale partial/full timing architectures, etc. These topics are crucial proof points to enable Segment Routing migration for many custom network configurations out there.

By far the most complex integration effort this year was spent on **use case testing**: We crafted two large-scale, realistic use case scenarios:

- 5G xHaul with Segment Routing Interworking
- EVPN service automation and assurance

Table of Contents

These use case scenarios involved almost all participating vendors, creating realistic architectures that can serve as **blueprints**, guiding operators towards vendor-independent network design. Being very demanding and time-consuming, the use case tests rewarded us with strong results. They are the foundation for the live demos shown at the Upperside Congress in Paris this year, are well documented, and will be expanded on next year.

Of course, we must not forget about AI these days. Our tests included both relevant sub-topics this year:

AI-enabled networking is governed by standardized provisioning (via PCE, Yang models, and BGP-SR — check), extensive telemetry data (via BGP-LS and TWAMP—check), and automated optimization checks (via digital twins—check). The vendors participating in this test area are on a steady path; that said, it is still a long way towards multi-vendor Autonomous Networks. Today, partial automation of specific service aspects in SR and EPVN is possible; it is important to require standardized methods in RfPs in detail.

Networking for AI workloads is a topic we wanted to cover more intensively, but it was too early. The next generation of data center transport has been defined by the Ultra Ethernet Consortium (UEC). The implementations naturally take time to get ready because they require major hardware innovations. We only covered a small aspect this time and plan to expand the UEC integration next year.

This 16-page report is only the *very* short version of all results. Please check out the QR codes with many more test results. We hope our joint effort is beneficial for any WAN, mobile x-haul, and data center network architects!

If you have any detailed questions, suggestions for next year's test coverage, or would like to tap our brains for an individual network design, please contact us.

Access the Full Test Report Online
https://eantc.de/interop26

EANTC's Mission

EANTC is a leading independent test lab dedicated to validating the interoperability, performance, robustness, and security of network solutions across platforms and applications. With 35 years of expertise, EANTC supports innovation by strengthening the reliability and operational readiness of vendor solutions. Through transparent, reproducible assessments, we help the industry ensure compliance with standards, reduce operational risk, and enable stable, trustworthy network deployments.

Interop Test Working Process

Preparations for the EANTC Transport & Cloud Networks Interop Test 2026 began in September 2025 with a series of technical calls involving all interested vendors.

During these phase, potential test cases were identified and refined, with vendors contributing new ideas, while the focus remained on exploring innovative testing approaches and ensuring alignment with the latest industry standards and IETF drafts.

The hot-staging took place in Berlin over three weeks in January/February 2026. The second and third week focused on intensive testing on-site. After in-depth discussions and rapid problem-solving during this period, vendors created all the test results, peer-reviewed by the EANTC team.

Interoperability Test Results

EANTC engineers closely supported and validated every test combination, following strict procedures and predefined steps. The resulting report presents only results that were consistently logged, submitted, and verified by EANTC specialists, ensuring accuracy and preventing misinterpretations or false positives.

Our focus is on multi-vendor testing and single-vendor demos are generally excluded. An exception is made if a previously agreed multi-vendor test case fails during hot staging, leaving only one vendor with a working, standards-compliant implementation. In such situations, EANTC acknowledges the effort and includes single-vendor results in the report.

This test report highlights successful test combinations, clearly identifying the participating vendors and devices. "Tested" in this context refers specifically to multi-vendor interoperability. Combinations that did not pass are not shown in the diagrams but are mentioned anonymously to provide insight into the industry's current state. Maintaining confidentiality is essential to encourage vendors to present their latest, often still in beta, solutions, creating a safe environment for testing, learning, and advancing network interoperability.

The test results will be presented live at the Upperside World Congress (previously the "MPLS World Congress") in Paris, March 24–26. For 24 years, EANTC has showcased its interoperability testing at Upperside conferences, highlighting the latest advances in network technologies.

Participating Vendors and Devices

The participating vendors and the devices evaluated in this edition, with several product families tested in multiple configurations to explore different interface types and hardware options.

Participants	Devices
Arista Networks	7050X3 7060X6 7280R3 7280R4
Calnex Solutions	Paragon neo, Paragon neo-A Paragon neo-S SNE Ignite, Sentry
Ciena	5164 8192
Cisco Systems	Crosswork, ASR-9901, ASR-9902 8011-4G24Y4H-I 8201-24H8FH 8711-32FH-M, 8711-48Z-M 8712-MOD-M, 9300-Fx2 9300-FX3, 93400LD-H1 93600CD-GX N9348Y2C6D-SE1U
Ericsson	RAN Connect 6682 Router 6671 Router 6676 Router 6678
Hewlett Packard Enterprise	ACX7020, ACX7024 ACX7100-32C, ACX7100-48L ACX7332, ACX7509 MX204, MX301, MX304 PTX10002-36QDD PTX10002-60MR
Keysight Technologies	IxNetwork Time Sync Analyzer
Microchip Technology	TimeProvider 4500 TimeProvider 4100
Nokia Corporation	Network Services Platform (NSP) 7210 SAS-X, 7220 IXR-D2L 7250 IXR-e2, 7250 IXR-X3B 7730 SXR-1x-44s, 7750 SR-1
Raisecom Technology	iTN8800-A RAX721-T-4C24
Ribbon Communications	NPT-2507
ZTE Corporation	ZENIC ONE R22 ZXCTN 6120H-E ZXCTN 6120H-S ZXR10 M6000-3S Plus ZXR10 M6000-2S16 ZXR10 M6000-4SE

Use Case Test Results

Modern networks are becoming increasingly complex, integrating diverse technologies, vendors, and architectures. To ensure that such multi-vendor environments can operate seamlessly, it is crucial to validate interoperability and standards compliance under realistic conditions. Building on our tradition of aligning testing efforts with the latest industry standards and drafts, this year's initiative aimed to bridge the gap between theory and practice by combining a broad set of technologies, features, and domains into a unified end-to-end test setup.

The physical testbed brought together 11 leading vendors—Arista, Calnex, Ciena, Cisco, Ericsson, HPE, Keysight, Microchip, Nokia, Raisecom, and ZTE—providing network nodes, emulation and measurement tools, and centralized control. Using this environment, we validated two representative use cases that reflect real operational scenarios for next-generation networks: EVPN Service Automation and Assurance and 5G xHaul with Segment Routing Interworking, as illustrated in the following figure.

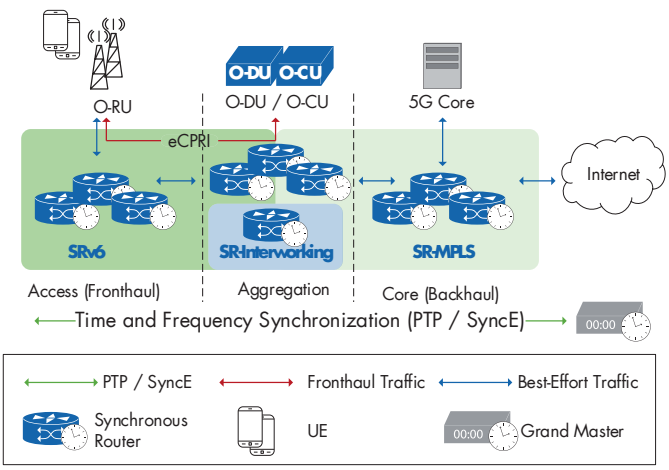

Figure 1: 5G xHaul Overview Topology

5G xHaul with Segment Routing Interworking

Mobile networks require precise time synchronization and deterministic transport to coordinate radio equipment and maintain service quality. In 5G deployments, fronthaul traffic must meet strict phase accuracy and low latency, while the transport network simultaneously carries mixed best-effort traffic. Quality-of-service mechanisms must ensure that congestion does not degrade time-sensitive streams. This use case was implemented and evaluated on the physical test bed illustrated in the following figure.

This test validates an end-to-end 5G transport scenario using a dual-domain Segment Routing architecture with SRv6-to-SR-MPLS interworking, representing a migration scenario. We configured two Segment Routing domains, one SR -MPLS and one SRv6, with IS-IS as the underlay IGP and BGP peering to a route-reflector for route distribution. At the domain boundary, we enabled SRv6-to-SR-MPLS interworking on the domain gateway so traffic could cross the encapsulation boundary and forward end-to-end between the access and core segments. On top of this transport, we enabled SyncE and Full Timing Support (ITU-T G.8275.1) PTP on the routers, established a baseline, and measured time error at the time synchronization test points; the results stayed within the limit of max $|TEL| = 1100$ ns from the Grandmaster to the output of the Access nodes. Additionally, the relative time error between the outputs of the Access nodes was measured, and all were within the +/-130 ns limit.

To validate behavior under congestion, we applied quality of service toolset to realize Time Sensitive Networking (TSN) Profile A, so that all DSCP CS6 traffic received priority treatment. We generated fronthaul eCPRI traffic between the O-RU and the O-DU marked with priority code point (PCP) 6. This was translated to DSCP CS6 in the outer IPv6 header upon encapsulation in the SRv6 domain. PTP synchronization packets were classified in the same traffic class. We then generated background best-effort traffic between the traffic generator endpoints to emulate internet traffic and saturate the links. Under saturation, we measured the change in PTP two-way time error compared to the baseline condition, during which no difference was observed. We still observed small latency (microseconds) and no packet loss for eCPRI or other DSCP CS6 traffic, such as PTP. High latency (milliseconds) and drops occurred only in the best-effort background traffic, which confirms that the prioritization and congestion handling worked as intended.

In addition to time synchronization and QoS validation, we used proactive OAM to verify 5G transport service health beyond control-plane checks. On the same testbed, participating nodes from Arista, Ciena, Cisco, HPE, Nokia, Raisecom, and ZTE were involved in TWAMP-based performance monitoring and gNMI telemetry collection via Cisco Crosswork Assurance.

Due to time constraints, we were unable to run additional scenarios. Topology variations, device role swaps, and alternative traffic paths remain valuable areas for future validation.

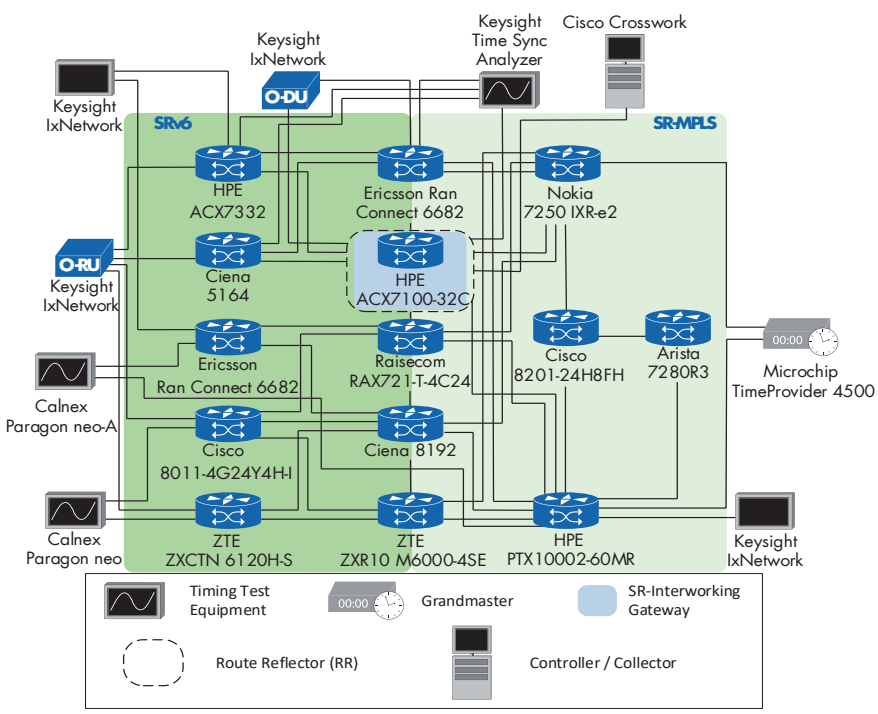

Figure 2: Use Case Physical Topology

EVPN Service Automation and Assurance

Operators need visibility into service health before customers experience degradation. Even if control-plane protocols indicate that an L2VPN is up, issues such as packet loss, latency spikes, or unstable connections can go unnoticed without active performance monitoring. Proactive Operations, Administration, and Maintenance (OAM) helps network administrators track performance indicators and identify problems early.

To verify that EVPN services remained healthy beyond control-plane checks, in this test, we set up proactive OAM monitoring on the same testbed, with participation from Ciena, Cisco, HPE, and ZTE nodes, and a Cisco Crosswork controller. The Cisco Crosswork Network Controller provisioned an SRv6-based EVPN ELAN service using NETCONF. We confirmed the service was active by checking EVPN MAC address learning on all nodes.

The logical topology for this test is illustrated in the following figure. The underlying physical test bed is shared with the 5G xHaul with Segment Routing Interworking test.

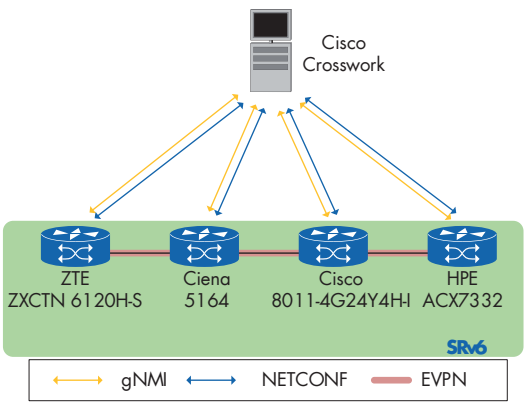

Figure 3: EVPN Service Automation and Assurance Overview Topology

Once the service was running, we enabled Ethernet Connectivity Fault Management (CFM) with Continuity Check Messages (CCM), Delay Measurement Messages (DMM), and Loss Measurement Messages (LMM). These OAM sessions targeted the remote Maintenance End Point (MEP) and used priority class 6 to match expected production traffic. Crosswork Assurance collected OAM performance metrics from the routers and provided a comprehensive view of service performance.

These three scenarios are showcased at the Upperside World Congress 2026 in Paris as live demonstrations.

For additional information and details regarding the individual test cases and the vendors involved, please visit the link or scan the QR code.

More Use Case Test Results
https://eantc.de/usecase26

EVPN Test Results

Ethernet VPN (EVPN) is a mature, widely adopted technology in data center (DC) networks, and has undergone extensive multi-year testing by EANTC. This year, we continued to verify interoperability across essential services, including E-Line, E-Tree, E-LAN, and redundancy. We also focused on gateway interoperability, including EVPN-SR-MPLS to EVPN-VXLAN, EVPN to IPVPN, VXLAN stitching, and the latest EVPN-to-EVPN gateway. The participating vendors include: Arista, Ciena, Cisco, Ericsson, HPE, Nokia, Raisecom, Ribbon, and ZTE.

EVPN Multihoming Support for L3 Services

Layer 2 and Layer 3 services can be commonly implemented within the same EVPN domain, as can multihoming. In this condition, for Layer 3 interfaces multi-homed to the same CE, we may encounter a situation where Layer 2 hashes ARP/ND requests to a single multihomed PE, leaving the other PE's ARP/ND table empty. A device with an empty ARP/ND table will drop traffic, disrupting load balancing. draft-ietf-bess-evpn-l3mh-proto introduces the solution for it. We tested the ARP/ND sync function in section 3.6. We simulated 4 subnets, 2 on each end, including both IPv4 and IPv6. And we sent ARP and ND from all hosts, and verified that ARP and ND entries were synced with EVPN RT-2. Both multihomed PEs had identical ARP and ND entries, including both local dynamic entries and EVPN-synched entries. We then sent bidirectional unicast traffic with Keysight IxNetwork and observed no packet loss and proper load balancing.

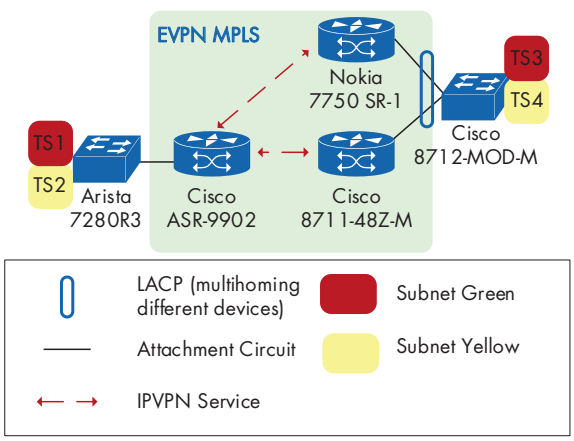

Figure 4: EVPN multihoming support for L3 services

EVPN-VPLS Layer 2 Attribute Extended Community

We have tested the EVPN-VPWS Layer 2 attribute last year. We continued testing this feature on the VPLS service this year. RFC 7432bis defined 3 Layer 2 attributes: Control Word, Flow Label, and Maximum Transmission Unit (MTU). This time, the Layer 2 attribute extended community is advertised along with EVPN RT-3 in VPLS service, whereas RT-1 is used in VPWS service. We verified the control-word-mismatch and control-word-match scenarios with unicast traffic under the VPLS service this year. When the control word mismatched, all traffic was dropped; when matched, all traffic flew without loss.

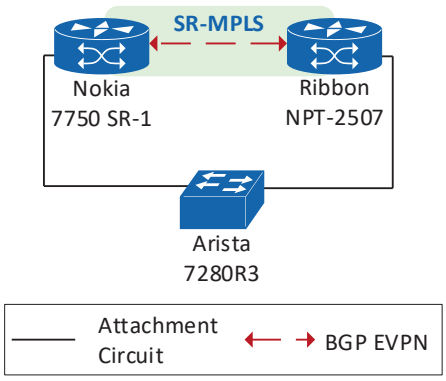

Figure 5: EVPN-VPLS Layer 2 attribute extended community

Dual-Stack Integrated Routing and Bridging (IRB) Interface

IRB has been tested for a long time under an IPv4 environment. This year, we have tested dual-stack IRB on the SR-MPLS test bed with VLAN-based services in symmetric working mode. The current networks are moving towards IPv6 today. Therefore, we have tested the essential step to move to IPv6: dual-stack. The underlay and overlay remain IPv4. However, the IRB interfaces enabled both IPv4 and IPv6. We sent both IPv4 and IPv6 traffic simultaneously, along with unicast routing and bridging traffic, with Keysight IxNetwork. We observed no packet loss, and bridging and routing traffic worked as expected.

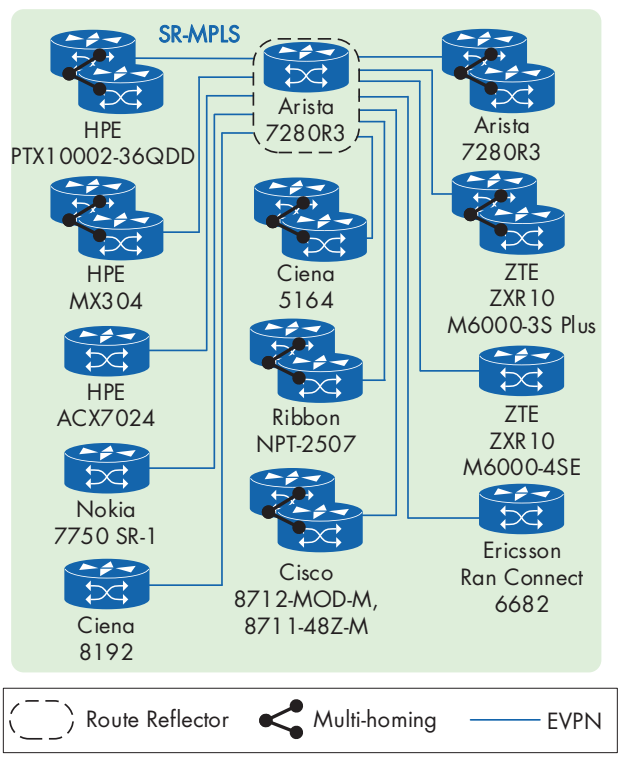

Figure 6: Dual-stack IRB interfaces

Interoperability Between Symmetric and Asymmetric IRB

Symmetric and asymmetric are the 2 working modes of IRB. In a real-world scenario, we're most likely to encounter a situation where some devices are configured with one IRB mode while others are configured with another within the same EVPN domain. This situation requires the two different IRB modes to interoperate with each other as per draft-ietf-bess-evpn-modes-interop.

In our test, the Arista leaf was configured with asymmetric IRB, while the Cisco leaf was configured with inter-op (hybrid) IRB mode. Then, we sent bidirectional unicast traffic, both inter- and intra-subnet, using Keysight IxNetwork between 2 VLAN-based IRB interfaces of these 3 symmetric and asymmetric PEs. There was no packet loss, confirming that bridging and routing functioned as expected. And we verified on the control plane that asymmetric PE had only 1 route target, whereas symmetric PE had 2. It means that both the control plane and the data plane worked properly in this interoperability scenario.

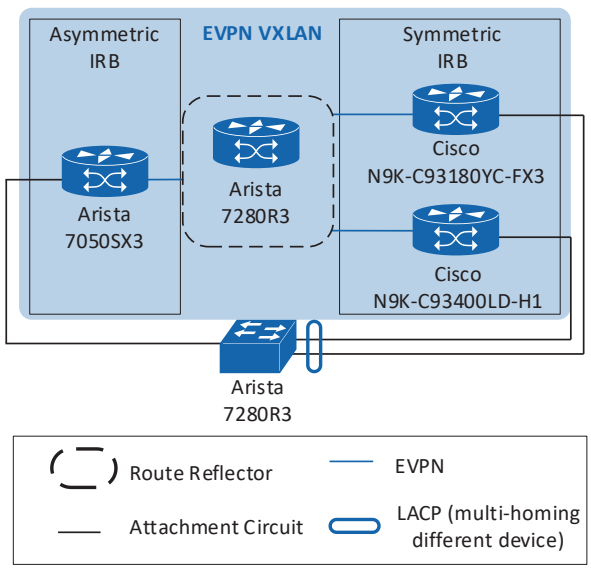

Figure 7: Interoperability between symmetric and asymmetric IRB

EVPN Group Policy

Micro- or macro-segmentation is a technique for dividing the network into multiple smaller segments (zones) by grouping endpoints based on their traffic patterns or access privileges, and controlling traffic between any two segments.

VXLAN uses group-based policy, based on draft-lrss-bess-evpn-group-policy, which specifies a mechanism for carrying Group Policy IDs (also known as Group Policy Tags) and VXLAN header extensions to enable micro- or macro-segmentation in the VXLAN fabric.

A test with four Tenant Systems (TS1, TS2 on PE1; TS3, TS4 on PE2) used group policies to permit only specific bidirectional traffic flows. Traffic was allowed between TS1–TS3, TS2–TS3, and TS2–TS4, while all other flows were denied. Keysight IxNetwork statistics showed that only permitted pairs communicated successfully, confirming group policy enforcement.

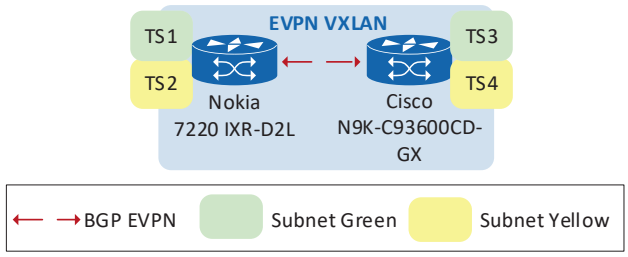

Figure 8: EVPN Group policy

OISM-Based Multicast Forwarding with EVPN-to-EVPN Gateway (EEG)

Multicast is essential for multimedia, video surveillance, and other scenarios, such as airports and stock exchanges. In past test setups, PIM/EVPN Gateway (PEG) connected only a single PIM domain to a single EVPN domain. To address the need for interconnecting multiple EVPN domains, we tested EVPN-to-EVPN Gateways (EEG), as described in draft-rabnic-bess-evpn-mcast-eeg.

In this test case, we verified that EEGs synchronize and redistribute multicast interests between EVPN domains using BGP EVPN RT-6 routes. The elected Supplementary Broadcast Domain – Designated Router (SBD-DR) ensures that traffic is transmitted across domains only when needed, minimizing unnecessary flooding and avoiding suboptimal routing. The external PIM domain had PIM-SSM enabled. The Interconnect Ethernet Segment (I-ES) links two VXLAN domains, while Domain Path (D-PATH) mechanisms prevent loops. To begin, we used Keysight IxNetwork to send IGMPv3 joins and observed that SMETs were generated, and the source was updated to the EEG's route ID. We then sent bidirectional multicast flows and confirmed there was no packet loss or duplication.

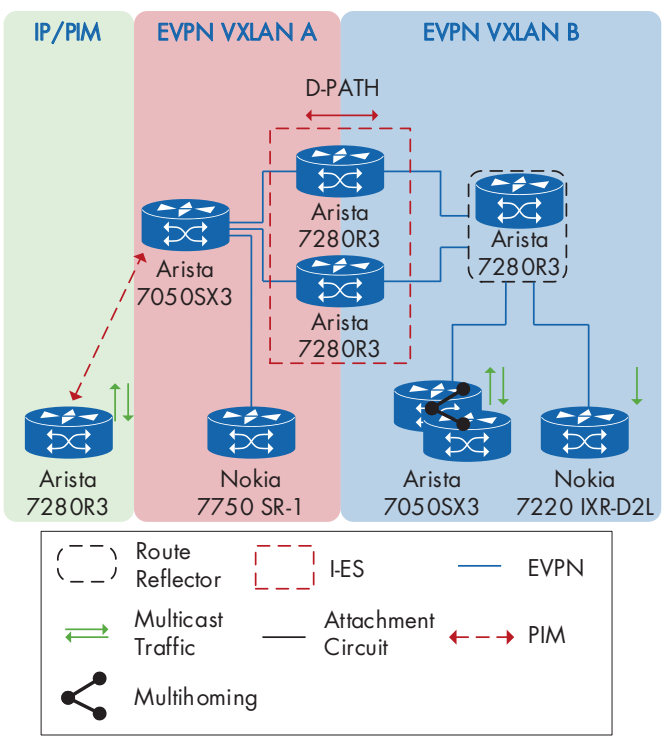

Figure 9: OISM-based multicast forwarding with EEG

Multi-Domain Interworking

In real-world networks, it is common to have multiple transport and EVPN domains distributed across different geographic locations. This year, we successfully recreated such a scenario in our test environment. Our topology consists of five interconnected domains: EVPN-SR-MPLS, EVPN-VXLAN, and IPVPNv4. The two VXLAN domains demonstrate VXLAN stitching. The two SR-MPLS domains interconnected via a VXLAN domain showcase the EVPN-SR-MPLS and EVPN-VXLAN Gateway features. The IPVPNv4 domain demonstrates EVPN IPVPN interworking. This multidomain topology covers several common scenarios in real-world use cases. We sent full-mesh unicast traffic with Keysight IxNetwork through all domains, testing both routing and bridging scenarios. All traffic was delivered end-to-end without loss.

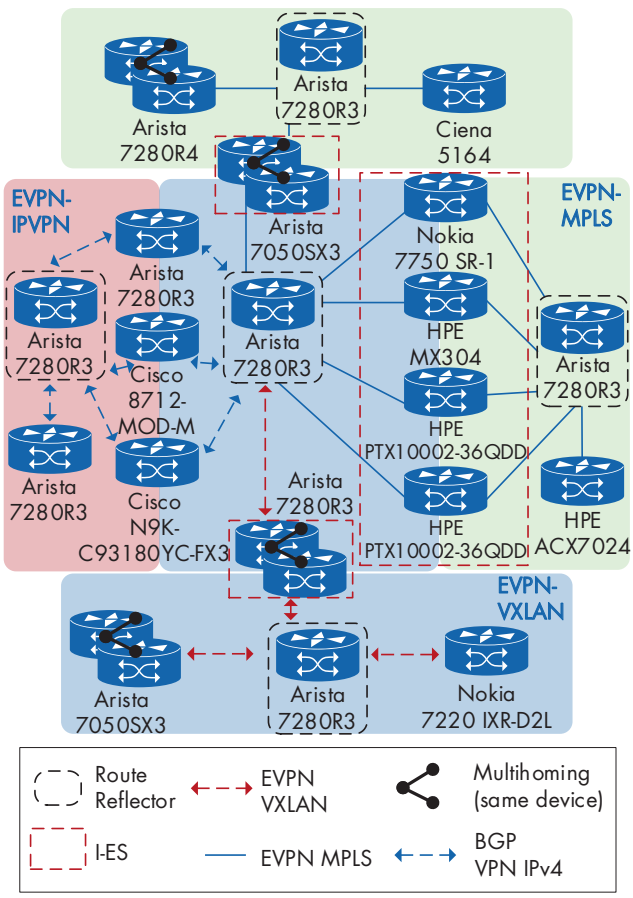

Figure 10: Multi-domain interworking

These last three scenarios are showcased at the Upperside World Congress 2026 in Paris as live demonstrations.

For additional information and details regarding the individual test cases and the vendors involved, please visit the link or scan the QR code.

More EVPN Results
https://eantc.de/evpn26

Physical Test

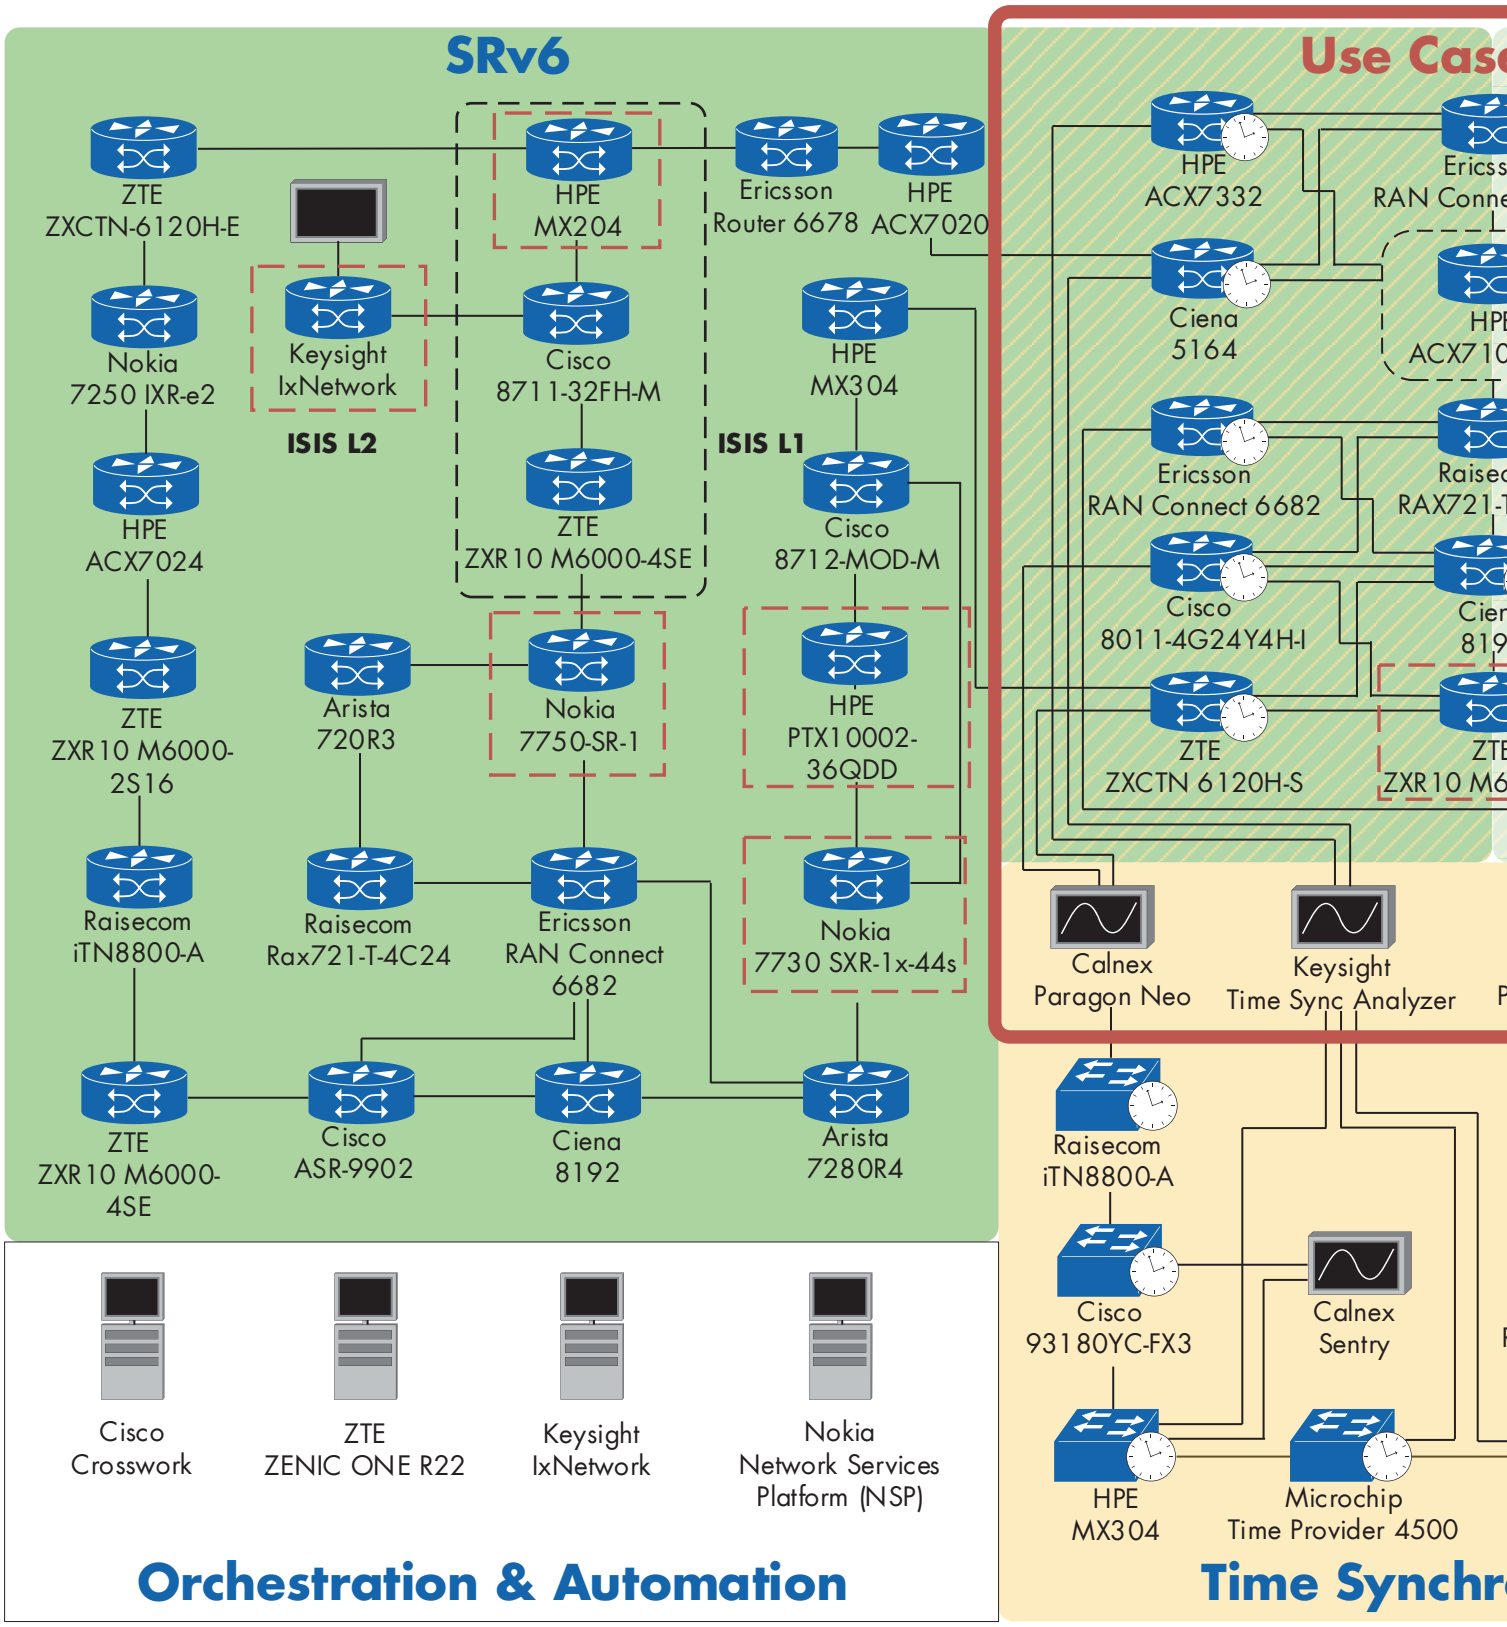

t Topology

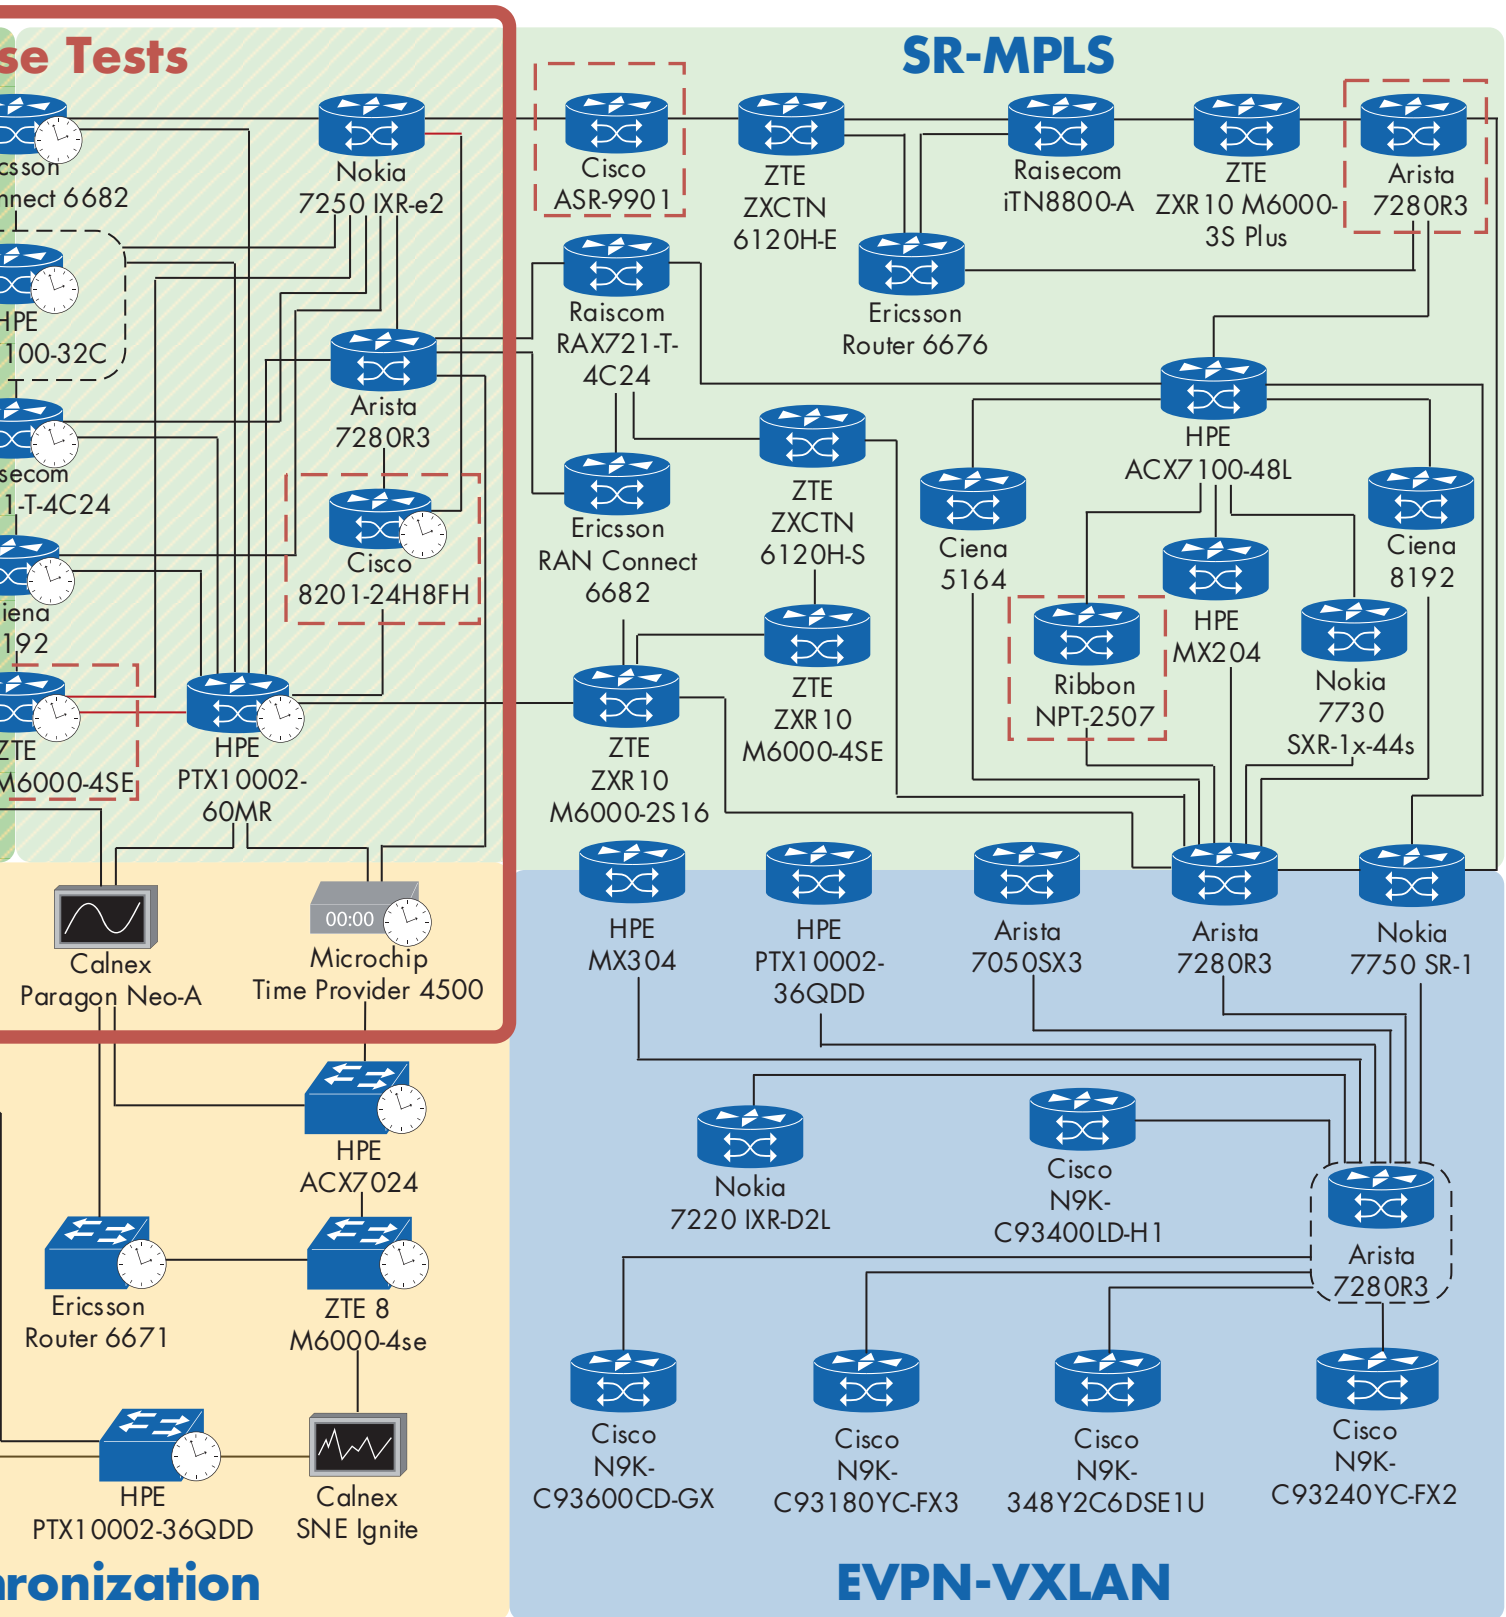

dmaster Clock	Managed by PCE / NETCONF		Controller / PCE
Generator	Route Reflector (RR)		Physical Link

Orchestration and Automation Test Results

The evolution of artificial intelligence is transforming how networks are provisioned, monitored, and troubleshot. Terms like autonomous networks are gaining more attention today. This development is mainly seen in individual components, such as a controller capable of generating network configurations for multiple devices simultaneously or analyzing thousands of log lines to detect failures and remediate them. A reliable controller that can handle these complex tasks requires robust protocols and a feature set that supports standardized machine-to-machine communication, enabling the collection of device status and telemetry, the implementation of device configurations, and real-time path computations and optimizations.

This year, in the orchestration and automation test area, we focused specifically on protocols such as NETCONF, which enables model-driven network device provisioning and configuration; gNMI for telemetry; PCEP for centralized path computation, reporting, and optimization; and BGP-based extensions such as BGP-LS and BGP-SR for topology discovery and SR Policy distribution. Seven leading vendors were involved in the testing: Arista Networks, Ciena, Cisco, HPE, Keysight, Nokia, and ZTE. We introduced several new test cases this year, including MSD reduction with Binding SID automation, digital twin network simulation, and L2/L3 service assurance using TWAMP.

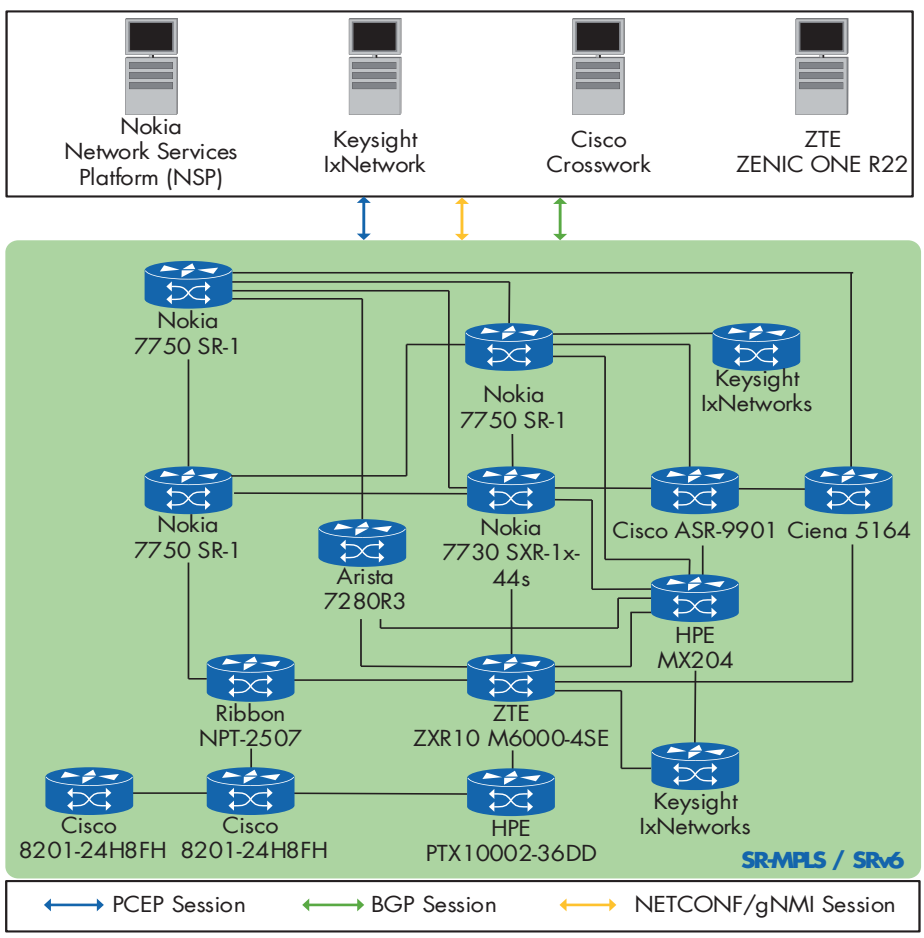

Figure 11: Physical Orchestration Test Setup

Path Computation Element Tests

We validated the core functionalities of a stateful PCE. The PCE computed SR Policies and SR-TEs and signaled them on headend routers via PCEP, while receiving path reports via PCEP to maintain a synchronized state. This year, SR policy signaling with SRv6 uSID was validated using both instantiation modes — PCE-initiated and PCC-initiated/PCE -delegated. Both PCEPv6 and PCEPv4 were used to test these SRv6 combinations. We also tested bidirectional path association, in which the PCE associated two unidirectional paths into a bidirectional path. Latency-based optimization was also tested: we simulated changes in network latency, and the PCE automatically recomputed and updated the affected path. The signaling of SR policies with Multiple Candidate Paths was also tested. We also verified that PCEP sessions can run over IPv6, which is critical for networks transitioning to IPv6-only operation.

A highlight of our path computation testing was the first-time validation of MSD (Maximum SID Depth) reduction. When a computed path exceeded the headend router's SID limit, the PCE automatically introduced a Binding SID at an intermediate anchor node to represent and compress the remaining segment list. During a simulated link failure, the PCE recomputed the path and dynamically updated the Binding SID, maintaining connectivity without requiring any intervention at the headend. For the first time, SR policy instantiation via API was also tested using gRPC, validated for a PCE-initiated SRv6 policy. In this model, the PCE translates API requests into PCEP, allowing controllers to instantiate SR/SRv6 policies across multi-vendor environments regardless of the headend vendor.

BGP-LS and BGP-SR

We used BGP-LS to collect and visualize network topology information. The controllers successfully gathered topology data, including information on SRv6 μSID and Flexible Algorithm. We also verified SR Policy distribution via BGP-SR, which provides a stateless alternative to PCEP for pushing computed paths to headend routers.

NETCONF

We executed multiple test cases for L3VPN provisioning over SRv6 locators using NETCONF and also used NETCONF to deploy routing policies that control how routes are advertised and selected throughout the network. gNMI (gRPC Network Management Interface)

For telemetry collection, we used gNMI to stream operational data

from network devices in real time. This gave the controller continuous visibility into device status, interface statistics, and protocol states.

Digital Twin Testing

This year, we tested digital twins for the first time. A digital twin is a virtual representation of the real network, automatically built from data collected from the actual network devices. We used such replicas to simulate configurations, metrics, paths, and failure changes without affecting the live network. This approach enables safe pre-deployment testing and what-if analysis without affecting the production network.

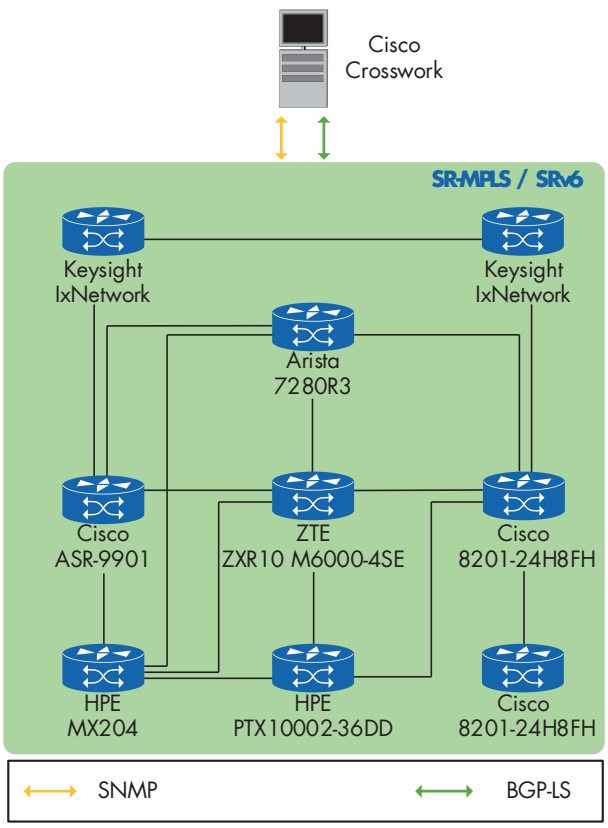

Figure 12: Digital Twin Test Topology

This scenario is showcased at the Upperside World Congress 2026 in Paris as live demonstrations.

Two-Way Active Measurement Protocol (TWAMP)

TWAMP testing was a new addition to this year's testing. We validated both L2 and L3 service assurance by measuring bidirectional delay, jitter, and packet loss between TWAMP sender and reflector endpoints. These measurements were collected and aggregated by the controller.

Most of the tests went as expected with interoperable implementation, but we identified some areas for improvement. Adoption of some recent IETF drafts is still maturing. For instance, we could not test bidirectional SR paths as described in "draft-ietf-pce-sr-bidir-path"; we could only combine the paths on the PCE side.

Some devices also showed limited SRv6 support, reducing the scope of SRv6-related scenarios we could validate. Although we executed more than 70 path-computation test runs this year, it remains clear that not all devices support all scenarios, such as SR-TE, SR Policies, and both PCE- and PCC-initiated paths.

For additional information and details regarding the individual test cases and the vendors involved, please visit the link or scan the QR code.

More Orch & Aut Results
https://eantc.de/or&au26

Segment Routing – SR-MPLS Results

Segment Routing over MPLS is a scalable and flexible transport technology for IP networks. By removing the need for traditional label distribution protocols and relying on source-based forwarding with segment identifiers, SR-MPLS allows operators to steer traffic deterministically while maintaining a straightforward operational model. For this year's event, we went through scenarios related to the failures and also enhancements of the control planes and operational stability on both IPv4 and IPv6. We verified in a multivendor environment that the services can be deployed for traffic steering, fast reroute, and scalability. Nine leading vendors were involved in the testing: Arista, Ciena, Ericsson, HPE, Keysight, Nokia, Raisecom, Ribbon, and ZTE.

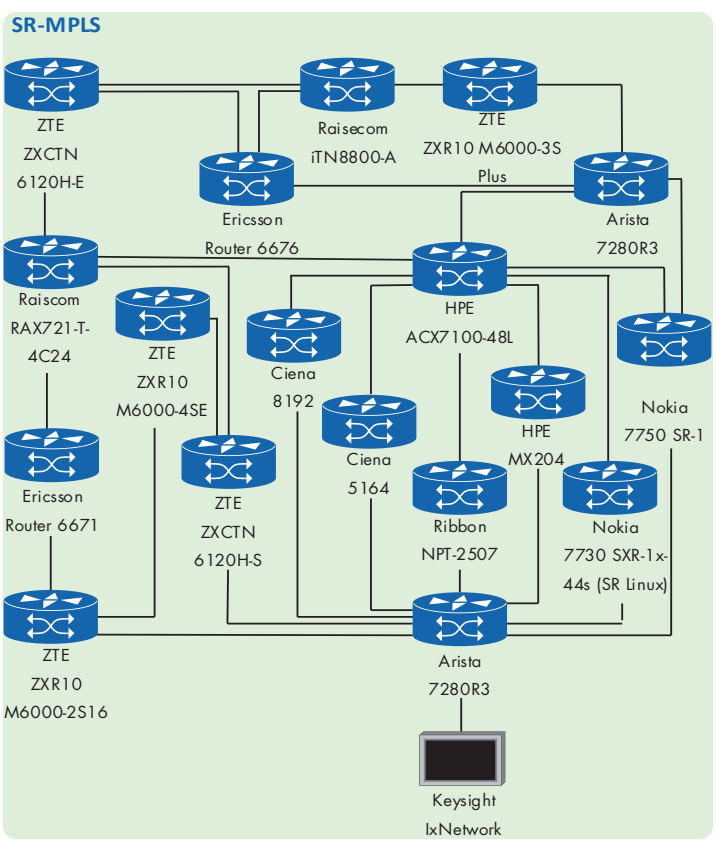

Figure 13: General SR-MPLS Topology

Nil-FEC (Forwarding Equivalence Class) Egress Validation for SR-MPLS Ping

The MPLS Operations, Administration, and Maintenance (OAM) methods are essential to confirm the reachability of the node in SR-MPLS networks. We validated Nil-FEC egress verification for SR-MPLS LSP ping and MPLS traceroute. We validated Nil-FEC egress verification for SR-MPLS LSP ping and MPLS traceroute. We created an SR-MPLS domain with one PE as the source and another PE as the egress, and made sure that PE could steer probes to other PE using a node-SID while keeping one of the transit routers on legacy code. From source PE, we sent MPLS Echo Request and MPLS traceroute probes that carried the Egress TLV with Nil-FEC and repeated the test over two paths. In the first path, the probes traversed HPE ACX7100 -48L, which processed the Nil-FEC information and returned the expected egress indication. In the second path, the probes traversed an Arista 7280R3, which responded with standard TTL-expired behavior for the hop. We observed return code 3 on the replies from HPE and ZTE PE devices, which is the backward compatible NIL-FEC validation behavior, and is defined in RFC 8029. Support for the updated RFC (RFC 9655) can be tested in future interoperability events. The figure below shows the Nil-FEC egress validation test topology and the two probe paths used during the run.

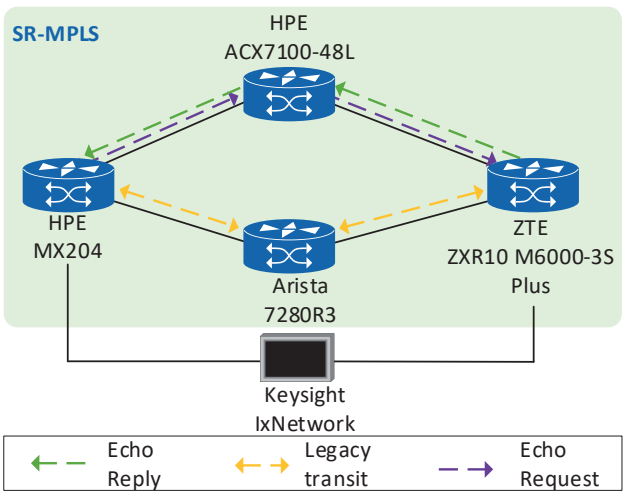

Figure 14: Nil-FEC Egress Validation for SR-MPLS Ping

TI-LFA with IPv6 SRLG using TLV 139

In transport networks, the same physical links can be used for different logical links, and a single failure in the physical link can have an impact on several logical links at the same time. In this test, we validated Topology Independent Loop-Free Alternate (TI-LFA) with Shared Risk Link Group (SRLG) in an SR-MPLS network using an IPv6 control plane. SRLG information was advertised using TLV 139, which supports SRLG signaling for IPv6-only topologies, including SRv6 and SR-MPLS with an IPv6 control plane. This enabled routers to calculate repair paths that avoid links sharing common physical risks. To check whether traffic on a failed link can be locally repaired within the same flexible algorithm topology, we introduced a link failure and verified the behavior as expected. It was without micro-loops and without reverting to a path that is not related. For networks with multiple links that share infrastructure and fiber, implementing SRLG -aware network protection would be important, as it would reduce failure impacts and help carrier-grade operators deploy networks with greater resiliency.

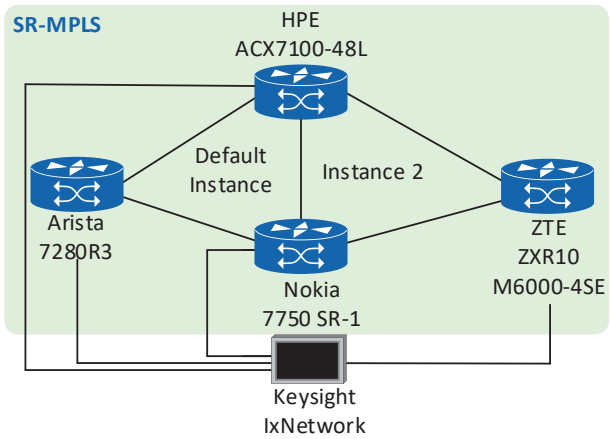

Figure 15: TI-LFA with IPv6 SRLG using TLV 139

Multi-Instance IS-IS Type-Length-Value (TLV) 7

In some scenarios, service providers require multiple logical topologies on the same physical infrastructure to isolate services or use it for migrations. We verified Multi-Instance ISIS using Instance Identifier TLV (type 7) in an SR-MPLS domain, and confirmed that TI-LFA protection operated independently per instance. We ran the test with two ISIS instances on the same physical link (and same logical link too) on HPE and Nokia devices, and confirmed that when a link that is being used by one instance fails, traffic for that instance would be rerouted using the other path and continued to traverse the DUTs as expected. In the first round, we disabled the link between Arista and HPE and observed traffic switching from Arista to Nokia and back to HPE. In the second round, we disabled the ZTE to Nokia link and observed reroute via HPE. We then swapped the instance placement and repeated the failures.

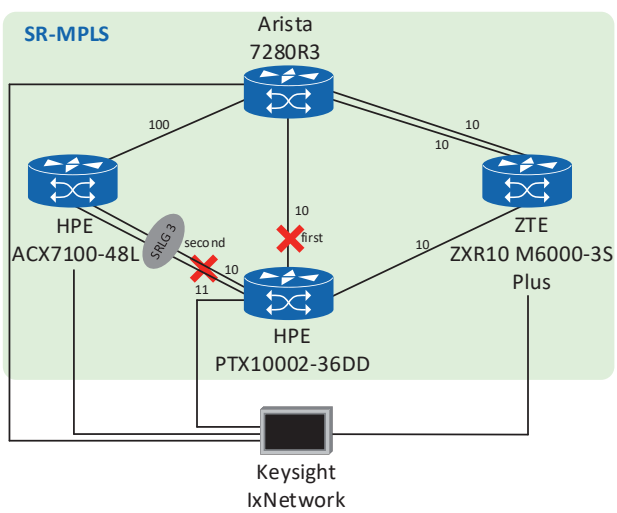

Figure 16: Multi-Instance IS-IS TLV 7

In addition to the test cases mentioned, we also executed test cases related to SR-TE traffic steering, Flexible Algorithm, Seamless BFD for SR policy liveness, Inter-AS SR-MPLS using Option C, SR-MPLS with IPv6 control plane for EVPN services, multi-vendor 400G ZR optical interoperability, and multicast services using BIER and IPMSI over SR-MPLS.

These last scenario is showcased at the Upperside World Congress 2026 in Paris as live demonstrations.

For additional information on the individual test cases and the vendors involved, please visit the link below or scan the QR code.

More SR-MPLS Results
https://eantc.de/sr-mpls26

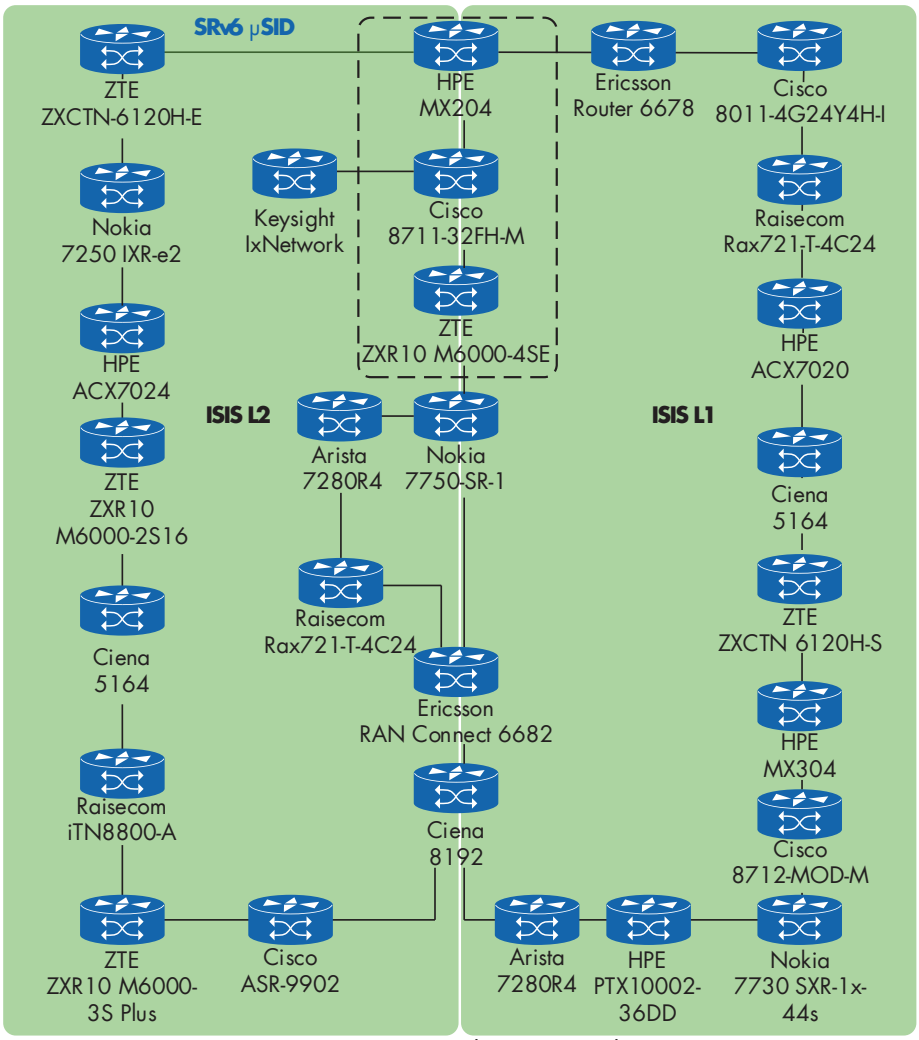

Figure 17: General SRv6 Topology

Segment Routing — SRv6 Results

The rapid expansion of AI workloads is transforming the AI ecosystem, including backend networks, data center interconnects, and service provider transport infrastructures. Currently, this traffic spans multiple network layers, each with distinct performance and reliability requirements. The increasing volume and diversity of such traffic show that no single networking approach addresses all operational requirements. In this context, SRv6 provides explicit mechanisms for traffic steering and forwarding control that can be applied where traditional approaches are insufficient.

This year, we continued verifying the interoperability status of SRv6 across multiple aspects, including VPN services (L3VPN, EVPN-ELAN, and EVPN IRB-Integrated Routing and Bridging), resiliency and convergence mechanisms (UPA — Unreachable Prefix Announcement, PIC — Prefix Independent Convergence, S-BFD — Seamless Bidirectional Forwarding Detection), routing and scalability features (Global Routing Table, route summarization, and Flex Algorithm), as well as traffic engineering capabilities. The interoperability tests were conducted using a dual-ring topology. The following vendors participated in building the test topology: Arista, Ciena, Cisco, Ericsson, HPE, Keysight, Nokia, Raisecom, and ZTE.

UPA and BGP PIC Edge in Summarized SRv6 Domains

Unreachable Prefix Announcement (UPA), as defined in draft-ietf-lsr-igp-ureach-prefix-announce, enables explicit signaling of unreachable prefixes in IS-IS to accelerate convergence.

The test designed a network that used summarization at domain boundaries to keep routing information manageable across the core. Locator and prefix summarization ensured control-plane scalability and reduced IGP state, as observed in real-world large-scale deployments.

This design would hide detailed reachability and failure information for individual nodes or prefixes outside the local domain. When an egress PE becomes unreachable, UPA explicitly signals the affected prefixes as unreachable, regardless of the summarized routing information. The ingress PE nodes are preconfigured with both primary and backup BGP paths toward multiple egress PEs. Upon receiving a UPA indication from the primary egress PE, the ingress PE immediately triggers BGP PIC (Prefix Independent Convergence) switchover and redirects traffic to the backup paths without waiting for BGP reconvergence, ensuring uninterrupted traffic forwarding.

The validation confirmed the correct use of the newly defined Prefix-SID flag bits 5 (U) and 6 (UP), verifying that UPA capability and unreachable prefix state are signaled and interpreted consistently across implementations.

This ensures explicit UPA signaling and guarantees that prefixes advertised as unreachable are promptly removed from the forwarding path, in line with the draft.

The following nodes showed interoperability in generating UPA and acting according to this signal:

ABR: Ericsson RAN Connect 6682 and HPE MX204.

Ingress PE: Cisco 8712-MOD-M, Ericsson Router 6678, HPE MX304.

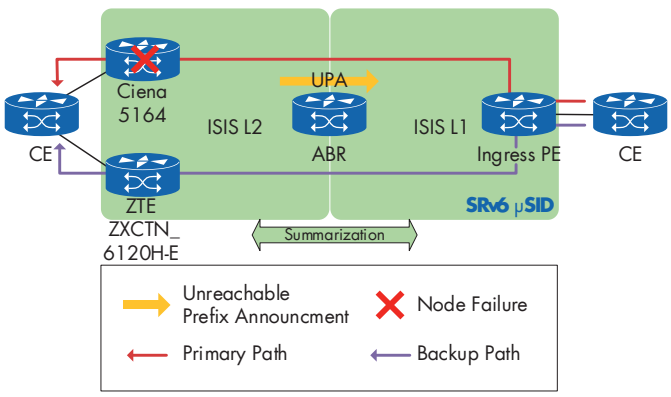

Figure 18: SRv6 Unreachable Prefix Announcement

encoding. RFC 9819 does not define the use of transposition for carrying the SRv6 Argument in the Prefix-SID of the AD/ES route. In our testing, we verified both scenarios: the transposition-based encoding as described in RFC 9252 and the non-transposition encoding as specified in RFC 9819.

The following pairs showed interoperability, encoding the argument ARG.FE2 and interpreting it correctly:

PE1 & PE2 : Cisco 8712-MOD-M and HPE MX304.

Ciena 5164 and Nokia 7750 SR-1 (SR OS).

Cisco 8712-MOD-M and Nokia 7750 SR-1 (SR OS).

Cisco 8712-MOD-M and Raisecom RAX721-T-4C24.

Ciena 5164 and Cisco 8712-MOD-M.

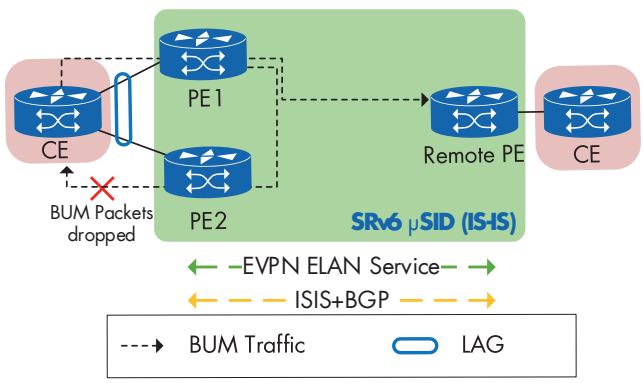

Figure 19: SRv6 ELAN MH using argument

SRv6 EVPN ELAN Multi-Homing Active-Active Using Argument

Two years ago, this mechanism was first validated with a limited set of vendor implementations. In this year's testing, the same capability was verified across a broader multi-vendor environment. The results showed consistent support for SRv6 SID argument signaling, indicating increased maturity and adoption.

The objective of the test is to validate interoperable split-horizon loop prevention for EVPN ELAN over an SRv6 data plane. The topology used an all-active multi-homing model. BUM traffic was forwarded traffic to the same Ethernet Segment, and split-horizon filtering was required to prevent this traffic from looping back toward the originating CE.

The test confirmed that loop prevention relied on the SRv6 SID Argument, specifically the Arg.FE2 value. The ingress PE encoded this information to identify the source Ethernet Segment of the BUM traffic. The egress PE compared the received value with the local Ethernet Segment identifier. If a match was detected, the packet was discarded and not forwarded.

According to RFC 9252, when the ESI filtering approach is used with the transposition encoding scheme, the 24-bit ESI Label field of the ESI Label Extended Community carries all or part of the Argument portion of the SRv6 SID, enabling the receiving PE to reconstruct the required SRv6 behavior without advertising the full SID. In contrast, RFC 9819 defines how the SRv6 Argument is encoded in the Prefix-SID of the AD/ES route when the ESI filtering approach is used, but it specifies only the non-transposition

SRv6 EVPN multi-homing support for L3 services with route sync for ARP/ND

draft-ietf-bess-evpn-l3mh-proto identifies a common challenge in L3VPN multi-homing designs: how a multi-homed CE performs Layer-2 hashing across its Link Aggregation Group (LAG) members. This hashing determines the outgoing interface for all Ethernet frames, including ARP and IPv6 Neighbor Discovery (ND) messages.

As a result, ARP or ND responses from the CE may consistently be directed to a single service PE. For instance, if the hashing algorithm always selects the link to PE1, only PE1 will populate its ARP/ND table, while PE2 will not obtain the corresponding neighbor information. Consequently, when unicast traffic from the core is load-balanced toward PE2, it may lack the required adjacency information and be unable to forward traffic correctly.

To resolve this issue, the EVPN L3 multi-homing draft specifies ARP/ND route synchronization. By exchanging ARP and ND information between peering PEs using EVPN MAC/IP Advertisement routes (EVPN Route Type 2), both PEs maintain a comprehensive and consistent view of CE adjacencies. This approach ensures correct unicast forwarding behavior regardless of the hashing of ARP/ND responses by the CE.

In this test, a single CE LAG was distributed across two independent PEs, both operating in all-active mode for the same L3VPN instance. The PEs established an EVPN Ethernet Segment toward the CE, which remained fully transparent to the CE. To ensure consistent Layer-3 forwarding in this topology, the PEs synchronized ARP and

IPv6 ND entries learned on their L3 interfaces using EVPN Route Type 2 advertisements.

During validation, ARP/ND entries learned on one PE were successfully advertised via EVPN and installed on the peer PE without requiring local relearning. Continuous CE-to-CE ICMP ping traffic verified data-plane behavior. When one of the CE-to-PE access links was intentionally disabled, the remaining PE immediately continued forwarding traffic for the L3VPN service. Traffic switched seamlessly to the redundant path, with no observable packet loss or convergence delay.

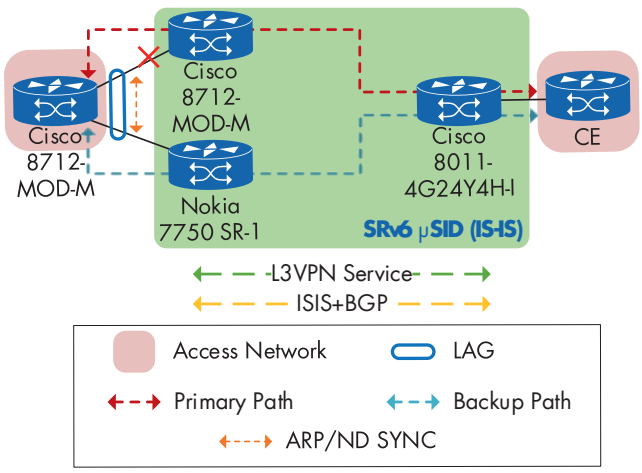

Figure 20: SRv6 EVPN multi-homing support for L3 services

These last two scenarios are showcased at the Upperside World Congress 2026 in Paris as live demonstrations.

For additional information on the individual test cases and the vendors involved, please visit the link or scan the QR code.

More SRv6 Results
https://eantc.de/srv6-26

Time Synchronization Test Results

Time synchronization in modern networks, such as 5G and the upcoming 6G, is facing increasing demands. Expectations for precision, system complexity, network size, and security are rising because 5G and 6G introduce new low-latency and coordinated radio use cases, along with more distributed architectures, all of which depend on highly precise and secure time synchronization; therefore, time synchronization solutions need to evolve to keep pace with these stricter requirements.

To verify the evolution of time synchronization solutions to meet increasing requirements, such as timing accuracy and stability, security, support for higher bandwidths, and interoperability, we introduced new time synchronization test cases into the EANTC interoperability tests and reworked older ones. All test cases were developed in accordance with ITU-T and IEEE standards, as well as O-RAN Alliance specifications.

This year, nine vendors participated in the time synchronization test area: Calnex Solutions, Ciena, Cisco, Ericsson, HPE, Keysight Technologies, Microchip Technology, Raisecom, and ZTE.

ITU-T G.8275.2 Packet Forwarding

In real-world synchronization setups, it's common to find multiple PTP profiles running in parallel, especially in mobile transport and xHaul networks. The choice of which profile to use often depends on the standard recommendation, network design, hardware in use, and the specific accuracy requirements of each segment.

For example, access and fronthaul segments usually depend on ITU-T G.8275.1 due to their need for high-accuracy phase and time delivery with Full Timing Support (ITU-T G.821, FTS). On the other hand, aggregation or IP-based segments might use Partial Timing Support (ITU-T G.8275.2, PTS), as routers in these segments are often not FTS-compliant, leaving PTS as the only option for time synchronization using PTP.

It's not unusual for networks to implement each of these profiles in different network segments, traditionally using Interworking Functions (IWFs) to convert between FTS and PTS segments and vice versa.

However, it should also be possible to make a PTS T-BC communicate directly with a Grandmaster that is providing both FTS and PTS PTP messaging to receive its PTP packets, rather than having FTS T-BCs process and convert G.8275.1 to G.8275.2, which can introduce additional time error.

This test verifies exactly that: the FTS T-BC-1 should operate as boundary clocks per G.8275.1 while simultaneously allowing G.8275.2 PTP packets to pass through transparently, like ordinary data traffic, as the G.8275.2 PTP packets are ordinary IP packets sent throughout the network and should therefore be forwarded unchanged, without processing or modifying timestamps.

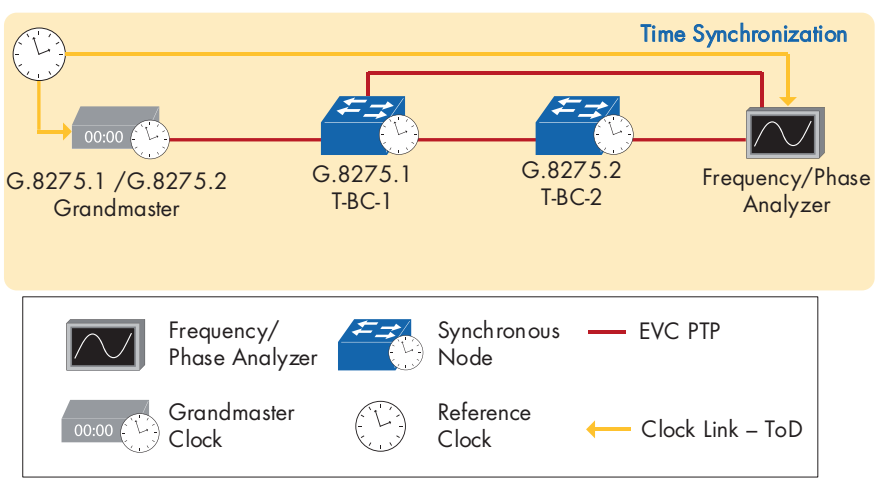

Figure 21: ITU-T G.8275.2 packet forwarding

Negative Leap Second

The International Earth Rotation and Reference Systems Service (IERS) is planning to remove the concept of leap seconds, including negative leap seconds, by 2035.

A negative leap second, as the name suggests, is the opposite of a positive leap second; during a positive leap second, one extra second is added, e.g., instead of the time going from 23:59:59 to 00:00:00, it goes from 23:59:59 to 23:59:60 and then to 0:0:00. During a negative leap second, however, one second is skipped, meaning time jumps from 23:59:58 to 00:00:00. This is done to align the Coordinated Universal Time (UTC) with the Universal Time (UT1), which is based on earths rotation.

To date, there has been no recorded instance of a negative leap second; however, while it has not yet been confirmed by the IERS, there is still a chance of one occurring by 2035. Since only positive leap seconds have occurred to date, the impact of a negative leap second on various computer systems and networks is unknown.

The goal of this test case was to verify that a time-synchronized network consisting of Telecom Boundary Clocks (T-BCs) and Precision Time Protocol (PTP) would correctly process a negative leap second and propagate this information downstream.

Since PTP itself is not affected by leap seconds, courtesy of using Temps Atomique International (TAI; eng. International Atomic Time), as leap seconds only affect UTC time, not TAI, the goal was to check whether all the flags in the announce messages were properly set and removed, as the T-BCs process the negative leap second information.

Additionally, test cases demonstrating PTP over MACsec, PTP over DWDM optics, PTP over 800GbE interfaces, APTS & delay asymmetry compensation up to 12 µs, and Interworking Functions were executed, with almost all T-BCs used in testing being T-BC Clock Class D-compliant. In addition, many of these features and test cases were combined into a single test case, showcasing the culmination of this year's Time Synchronization testing in a single setup.

This last scenario is showcased at the Upperside World Congress 2026 in Paris as live demonstration.

For additional information on the individual test cases and the vendors involved, please visit the link or scan the QR code.

More Time Sync Results
https://eantc.de/timesync26

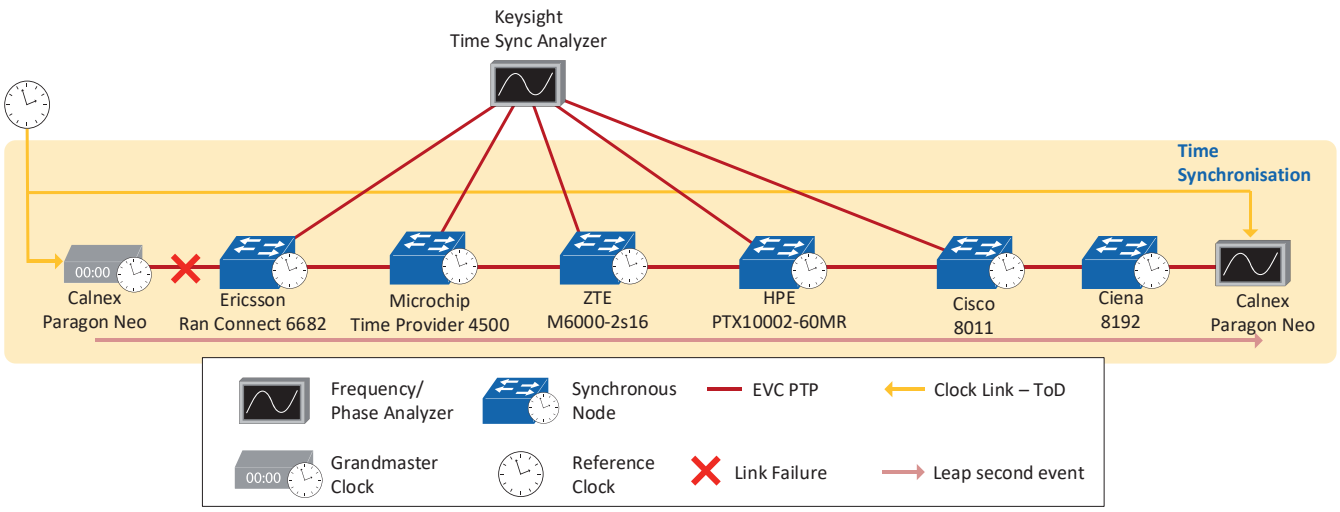

Figure 22: Negative Leap second in a chain of class D T-BCs

Access the Full Test Report Online
https://eantc.de/interop26

EANTC AG, Salzufer 14, 10587 Berlin, Germany

info@eantc.de / www.eantc.de

V1.0 20260312