EANTC

EUROPEAN ADVANCED NETWORKING TEST CENTER

Multi-Vendor MPLS & SDN Interoperability Test Report 2025

EVPN, Orchestration and Automation, Segment Routing, and Time Synchronization

MPLS&SRV6AINET
WORLD25

Editor's Note


Carsten Rossenhoevel
Co-Founder & CTO

Welcome to the latest edition of the EANTC multi-vendor interoperability test report! It has been an honor for us to gather the leading network equipment manufacturers again—Arista Networks, Arrcus, Calnex Solutions, Ciena, Ericsson, H3C Technologies, Huawei Technologies, Juniper Networks, Keysight, Microchip, Nokia, and Ribbon Communications participated. In February, a team of more than 70 engineers from vendors and EANTC conducted an intense three-week test event at our lab in Berlin.

This year, **we split the report into a paper-based overview and a much more extensive online version.** Our website has full details of all test results. **Please scan the QR codes,** enter the URLs displayed on this and the following pages, **or browse EANTC's website to access the individual report areas**. The online report is more detailed than the printed report in previous years!

This year, we have continued our coverage of multi-vendor interoperability of segment routing, EVPN, time synchronization, network management, and orchestration. Meanwhile, these technologies have matured considerably, albeit to different degrees (see more details in each report section). Once we had built a baseline of interoperable services between participating vendors, we focused on the latest innovations in each area—cloud data center and data center interconnect services, and advanced 5G/6G transport use case tests.

Table of Contents

Segment Routing over IPv6 (**SRv6**) has evolved from a challenger to a major pillar of our transport protocol tests. Its advantages—a single, scalable transport solution spanning from data centers through access and aggregation to the core network—are convincing. The number of vendors participating in this test area is stable. In parallel, **SR-MPLS** (Segment Routing over MPLS) remains a stronghold for service provider networks because it works on many routers without needing hardware upgrades. For the foreseeable future, SRv6 and SR-MPLS will coexist. Likewise, VXLAN maintains its role in data center (DC) networks, supported by standard switch hardware and simple to manage for DC operators. Consequently, vendors tested gateways between SR-MPLS, VXLAN, and SRv6 again.

EVPN (Ethernet VPN Service) is the workhorse carrying all the end-to-end application data across all segment routing and VXLAN transport networks. We validated operational improvements and focused on interworking EVPN gateways between VXLAN and SR-MPLS. Additionally, we helped vendors expand the interoperability of many of the core EVPN service attributes.

Access the full test report
https://eantc.de/testreport25

5G networks continue to evolve, with packetized transport becoming the de facto standard across all segments. Transformation particularly occurs in fronthaul networks, where Ethernet-based transport solutions substitute traditional CPRI-based connectivity, for example, in the O-RAN Alliance architecture. The migration enables greater flexibility and cost savings; at the same time, it introduces stringent transport network performance and reliability requirements. Critical parameters such as ultra-low latency, minimal packet delay variation, and precise timing synchronization must be maintained to ensure optimal radio access network operation. Our interoperability tests reflected multi-vendor, real-world deployment conditions: We evaluated x-haul SRv6 environments with parallel PTP traffic, secure PTP transport using MACsec, constraint-based routing based on policies, and more.

Many 2025 MPLS & SRV6AI Net World conference presentations focus on artificial intelligence (AI). In preparation for our interop test, we encouraged vendors to contribute AI-enabled network solutions. However, today's early AI-based implementations focus on single-vendor-focused domain-specific optimizations, and some use proprietary telemetry collection methods. From EANTC's point of view, the strength of today's Internet routing solutions (and the theme of our event) is inter-operability, which drives innovation, protects investments, and enables a global market.

We feel that our continued coverage of multi-vendor solutions helps progress the industry more than jumping on the AI bandwagon too early: Interoperable automated provisioning with PCEP and NETCONF/Yang, and interoperable automated network optimization with Segment Routing policies and flexible algorithms are a reality already today. We look forward to including Yang push and other improved standardized telemetry mechanisms next year—and potentially AI-based solutions if they will fit in standardized, multi-vendor environments.

I sincerely thank all the participating vendor teams for their outstanding commitment to this year's test over many months, through an intense three-week hot staging at EANTC in Berlin and culminating in the public live display in Paris. It is an honor for EANTC to see that our interoperability efforts contribute a little to the success of many service provider, enterprise, and government networks worldwide—enabling a wide choice of innovative solutions and a lively global market. The EANTC team hopes that this in-depth, factual and unbiased test report provides new insights and proof points for future network deployments.

EANTC's Mission

EANTC provides vendor-neutral testing to validate the interoperability, performance, robustness, and security of network solutions, platforms, and applications. Since 1991, we have been committed to transparent and reproducible assessments, helping the industry ensure standards compliance and minimize operational risks. With over 30 years of experience, we accelerate technology development and enhance the stability of vendor solutions, enabling reliable and high-quality network deployments.

Interoperability Test Working Process

Preparations for the MPLS & SDN interoperability event began in September 2024, with initial discussions on test areas and potential test cases held across multiple technical calls with all participating vendors. These sessions focused on reviewing test case details, exploring new testing ideas, and aligning with the latest standards to ensure the test plans are up-to-date with industry advancements.

The hot-staging event with all vendors was held in Berlin in February, where the latest hardware arrived from all over the world. Three weeks of intensive testing, detailed discussions, and quick problem-solving led to excellent outcomes—more than 1,300 result pairs achieved.

EANTC engineers carefully observed and validated each test combination, following strict procedures and predefined test steps. This report includes only those results that were consistently submitted, logged, and verified by EANTC specialists to ensure accuracy and avoid misinterpretations or false positives.

Test Area Selection

EANTC defines the test areas in collaboration with participating vendors to cover all key aspects of service provider networks. Vendors contribute new test cases, often with input from IETF RFC and draft editors within their teams. Given the extensive scope, we prioritize test cases that receive broad implementation support.

Our primary focus is on multi-vendor testing, so test cases implemented by a single vendor are generally excluded. However, an exception is made if a previously validated multi-vendor test fails during hot staging, leaving only one vendor with a working, standards-compliant implementation. In such cases, we recognize their commitment and report the result.

Interoperability Test Results

As always, this test report highlights the successful test combinations, along with the corresponding vendor and device names. In this context, the term "tested" refers to multi-vendor interoperability tests.

Test combinations that did not pass are not shown in the diagrams but are referenced anonymously in the report to illustrate the current state of the industry. Maintaining confidentiality is key to encouraging vendors to bring their latest—often still in beta—solutions to the table, ensuring a protected space for testing and learning.

The Test Results will be presented live at the 26th edition of the MPLS & SRv6 AI Net World Congress. For over 20 years, we have showcased interoperability tests at Upperside's conferences.

Participating Vendors and Devices

The following vendors and devices were tested this time. In some cases, multiple fixed configurations of the same product families were tested, to explore different interface types or hardware options.

Participants	Devices
Arista Networks	7050SX3 7280R2 7280R3
Arrcus	S9600-72XC S9610-36D
Calnex Solutions	Paragon-neo PAM4 Paragon-neo S Sentinel Sentry SNE Ignite
Ciena	5134 8140 Coherent Metro Router Navigator Network Control Suite
Ericsson	Router 6671, 6676, 6678
H3C Technologies	AD-WAN CR16000-M1A CR16003E-F CR16005E-F S12500R-48Y8C S12500R-48C6D
Huawei Technologies	ATN 910C-G ATN 910D-A NetEngine A816 NetEngine A821 NetEngine 8000 M14 NetEngine 8000 M8 NetEngine 8000 X4
Juniper Networks	ACX7024, ACX7024X ACX7100-32C, ACX7100-48L ACX7348, ACX7509 MX10004, MX204, MX304 PTX10002-36QDD QFX5120-48Y QFX5130-32CD Cloud-Native Router (JCNR)
Keysight	IxNetwork Time Sync Analyzer
Microchip	TimeProvider® 4500
Nokia	7220 IXR-D2L 7250 IXR-e2 7250 IXR-X3b 7750 SR-1 7730 SXR-1x-44s Network Services Platform (NSP)
Ribbon Communications	NPT-2100

EVPN Test Results

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

Ethernet Virtual Private Network (EVPN) is a well-established technology used in data centers (DC) and across service provider networks. EVPN was invented in 2015, and we have accompanied its evolution with interoperability tests since the same year. EVPN has become the dominant solution in the data center and data center interconnect (DCI) market due to its performance, redundancy, multi-service integration, and ability to interwork across various transport domains. The Integrated Routing and Bridging (IRB) feature enables Layer 3 routing functionality, accommodating complex Layer 2 and Layer 3 routing scenarios.

EANTC divided the EVPN tests into the transport domains:

- EVPN-SR-MPLS, where Arista, Arrcus, Ciena, Ericsson, H3C, Huawei, Juniper, Nokia, and Ribbon participated
- EVPN-VXLAN with participating vendors Arista, H3C, Juniper, and Nokia
- EVPN over SRv6 results are documented in the SRv6 section (further below).

Spirent TestCenter was used as the traffic generator for all EVPN SR-MPLS and EVPN VXLAN test cases.

This year, new tests included LSP Ping mechanisms for EVPN, EVPN auto-Ethernet Segment Identifier (ESI), and D-path attributes with gateway interworking. We also performed several essential legacy test cases, such as E-Line service, E-LAN service, Integrated Routing and Bridging (IRB) interface, MAC mobility, EVPN IP-VPN gateway, and interoperability testing between EVPN-VXLAN and EVPN-SR-MPLS gateways.

LSP Ping for EVPN

RFC 9489 defines a Label-Switched Path (LSP) Ping Mechanism for EVPN, a simple yet essential tool for analyzing connectivity issues on the EVPN services layer. We successfully tested the LSP ping solution for EVPN VPWS point-to-point services and verified that the echo reply packets bear a return code set to 3, as the standard requires.

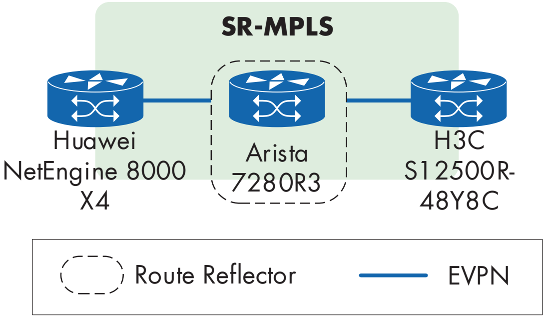

Figure 1: LSP ping mechanism for EVPN VPWS

Ethernet Segment Identifiers (ESI)

RFC 7432 defines six ways to generate an Ethernet Segment Identifier (ESI). We tested type 1, which uses the LACP system MAC and Port Key and allows automated generation of unique ESIs. This test verified that the ESI for the whole group remained identical and that RT-4 propagated throughout the entire Ethernet segment.

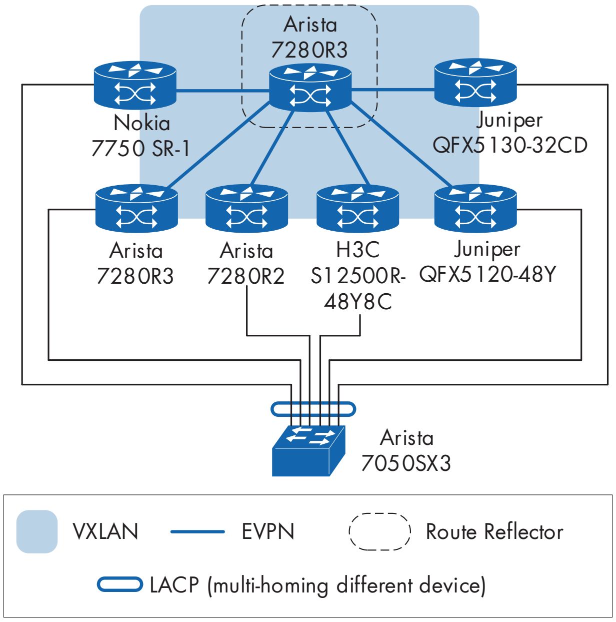

Figure 2: EVPN ESI Auto-Configuration

EVPN Interworking

EVPN-SR-MPLS and EVPN-VXLAN gateway interworking is a typical real-world scenario in which datacenter interconnect traffic passes through multiple WAN transport domains. We have covered it in EANTC testing many times—this year, we included the BGP Domain Path attribute in our test for the first time, which improves the loop prevention function.

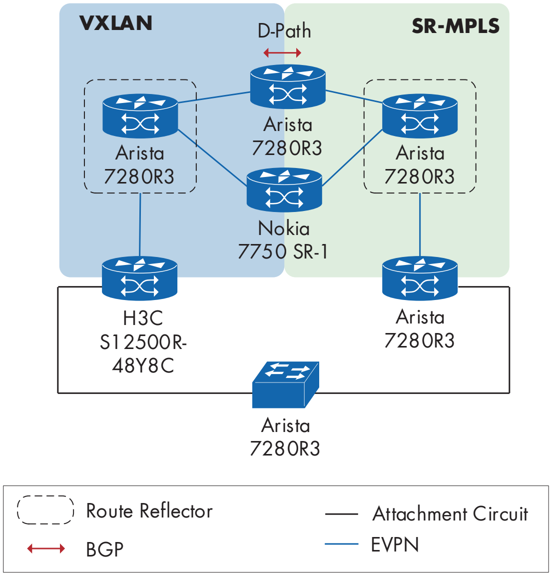

Figure 3: EVPN-SR-MPLS and EVPN VXLAN gateway interworking with D-path

Orchestration and Automation Test Results

Transport networks are becoming increasingly complex due to hybrid clouds, growing numbers of micro-segmented services, and differentiated quality and policy demands. By using centralized controllers for provisioning and monitoring network elements, operators can greatly improve efficiency, increase reliability, and simplify operations of network maintenance tasks.

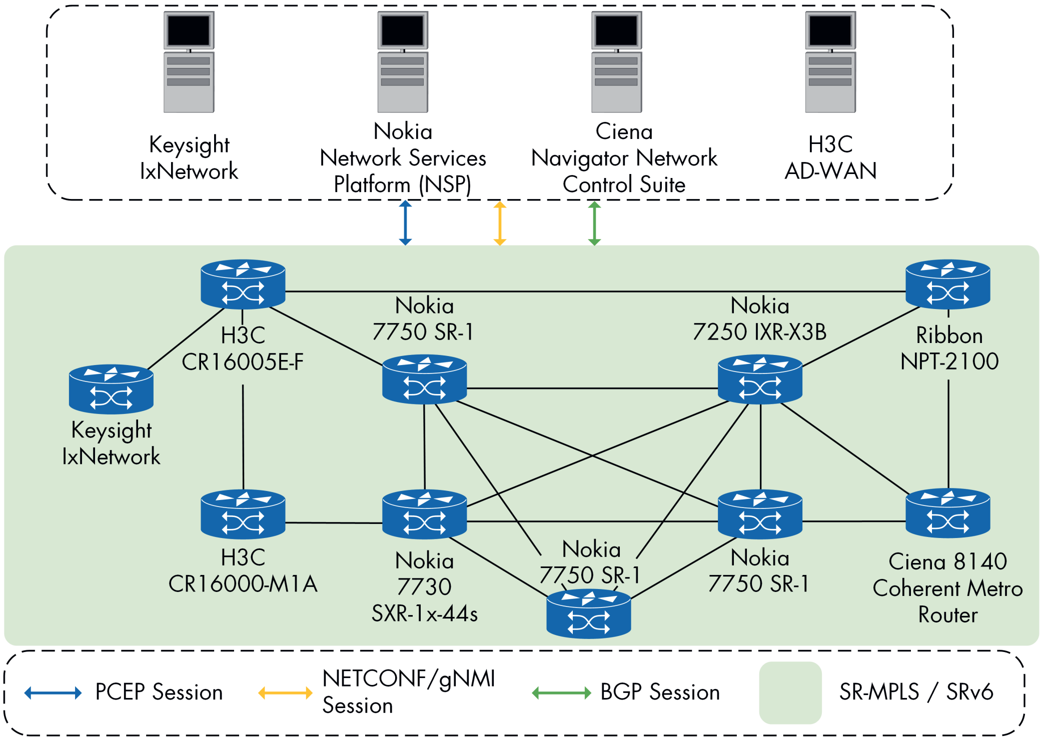

Figure 4: Orchestration and Automation Physical Test Setup

The participating network controllers and network elements (routers) utilized Path Computation Element Protocol (PCEP), Border Gateway Protocol (BGP), Network Configuration Protocol (NETCONF), and gRPC Network Management Interface (gNMI), to accomplish network automation and respond to real-time network changes. All tests were conducted in a multi-vendor environment, validating the interoperability of collaborating systems.

Nokia, Ciena, and H3C provided controllers and network nodes for the Orchestration and Automation test area. Ribbon supplied routers, Juniper participated in a test case with a router, and Keysight participated with IxNetwork, emulating controllers and routers alike.

Path Computation Element (PCEP) Tests

PCEP computes and signals paths to network nodes based on various constraints. As a baseline, continuing similar efforts in previous EANTC interop events, we conducted multiple tests on the SR-MPLS and SRv6 data planes (the latter with full SIDs and μSIDs).

We assessed PCEP's main functions, including the computation and signaling of traffic-engineered paths for IPv4/IPv6 SR-TE (SR-MPLS) and IPv4/IPv6 SR-Policy (SR-MPLS/SRv6). The scenarios included both PCE-Initiated and PCC-Initiated with combinations of reported-only paths and reported and delegated paths.

We tested dynamic path instantiation, where a PCC requests a path from the PCE in response to a routing protocol update, allowing the device to adapt to real-time network changes. We also conducted a latency-based path optimization test, where the PCE computes and optimizes the path in response to high latency, bypassing affected links. Furthermore, we validated the establishment of PCEP sessions using IPv6 addresses and the path computation using IPv6 endpoints in an IPv6 SR-MPLS and IPv6 SRv6 environment, ensuring seamless functionality between the PCE and PCC in IPv6 environments.

Bidirectional PCE Paths

This year, we validated bidirectional PCEP paths with strict symmetrical routing for the first time. In this case, the controller groups two opposite unidirectional paths, and if the end-to-end delay increases in one direction, the controller recomputes both unidirectional paths, even if the other direction is not impaired. This ensures paths in both directions always traverse the same links. The new bidirectional path is then signaled to the two headend PCCs.

Additionally, we validated PCE's capability to discover and visualize the μSID topology in an SRv6 environment. We also verified the PCE's ability to signal a binding-SID to a PCC and revisited the PCEP association groups test, covering both its parts: diversity path and policy association.

BGP in Software-Defined Networks

BGP plays an essential role in software-defined networks by enabling communication between the PCE and PCC. We conducted multiple tests focusing on BGP's capabilities, including BGP link-state for topology discovery

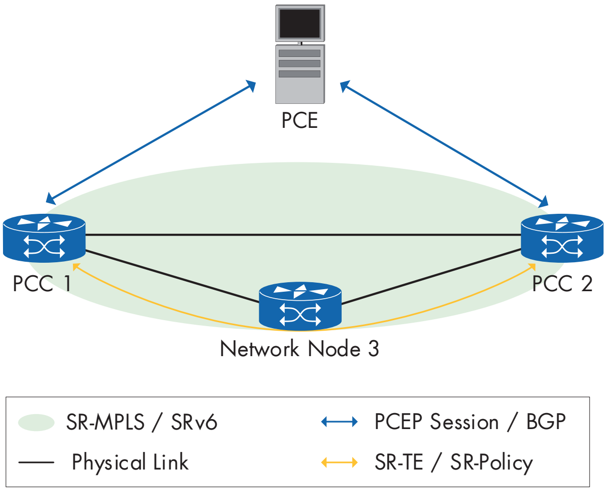

Figure 5: Path Computation Element Topology Overview

Physical Test

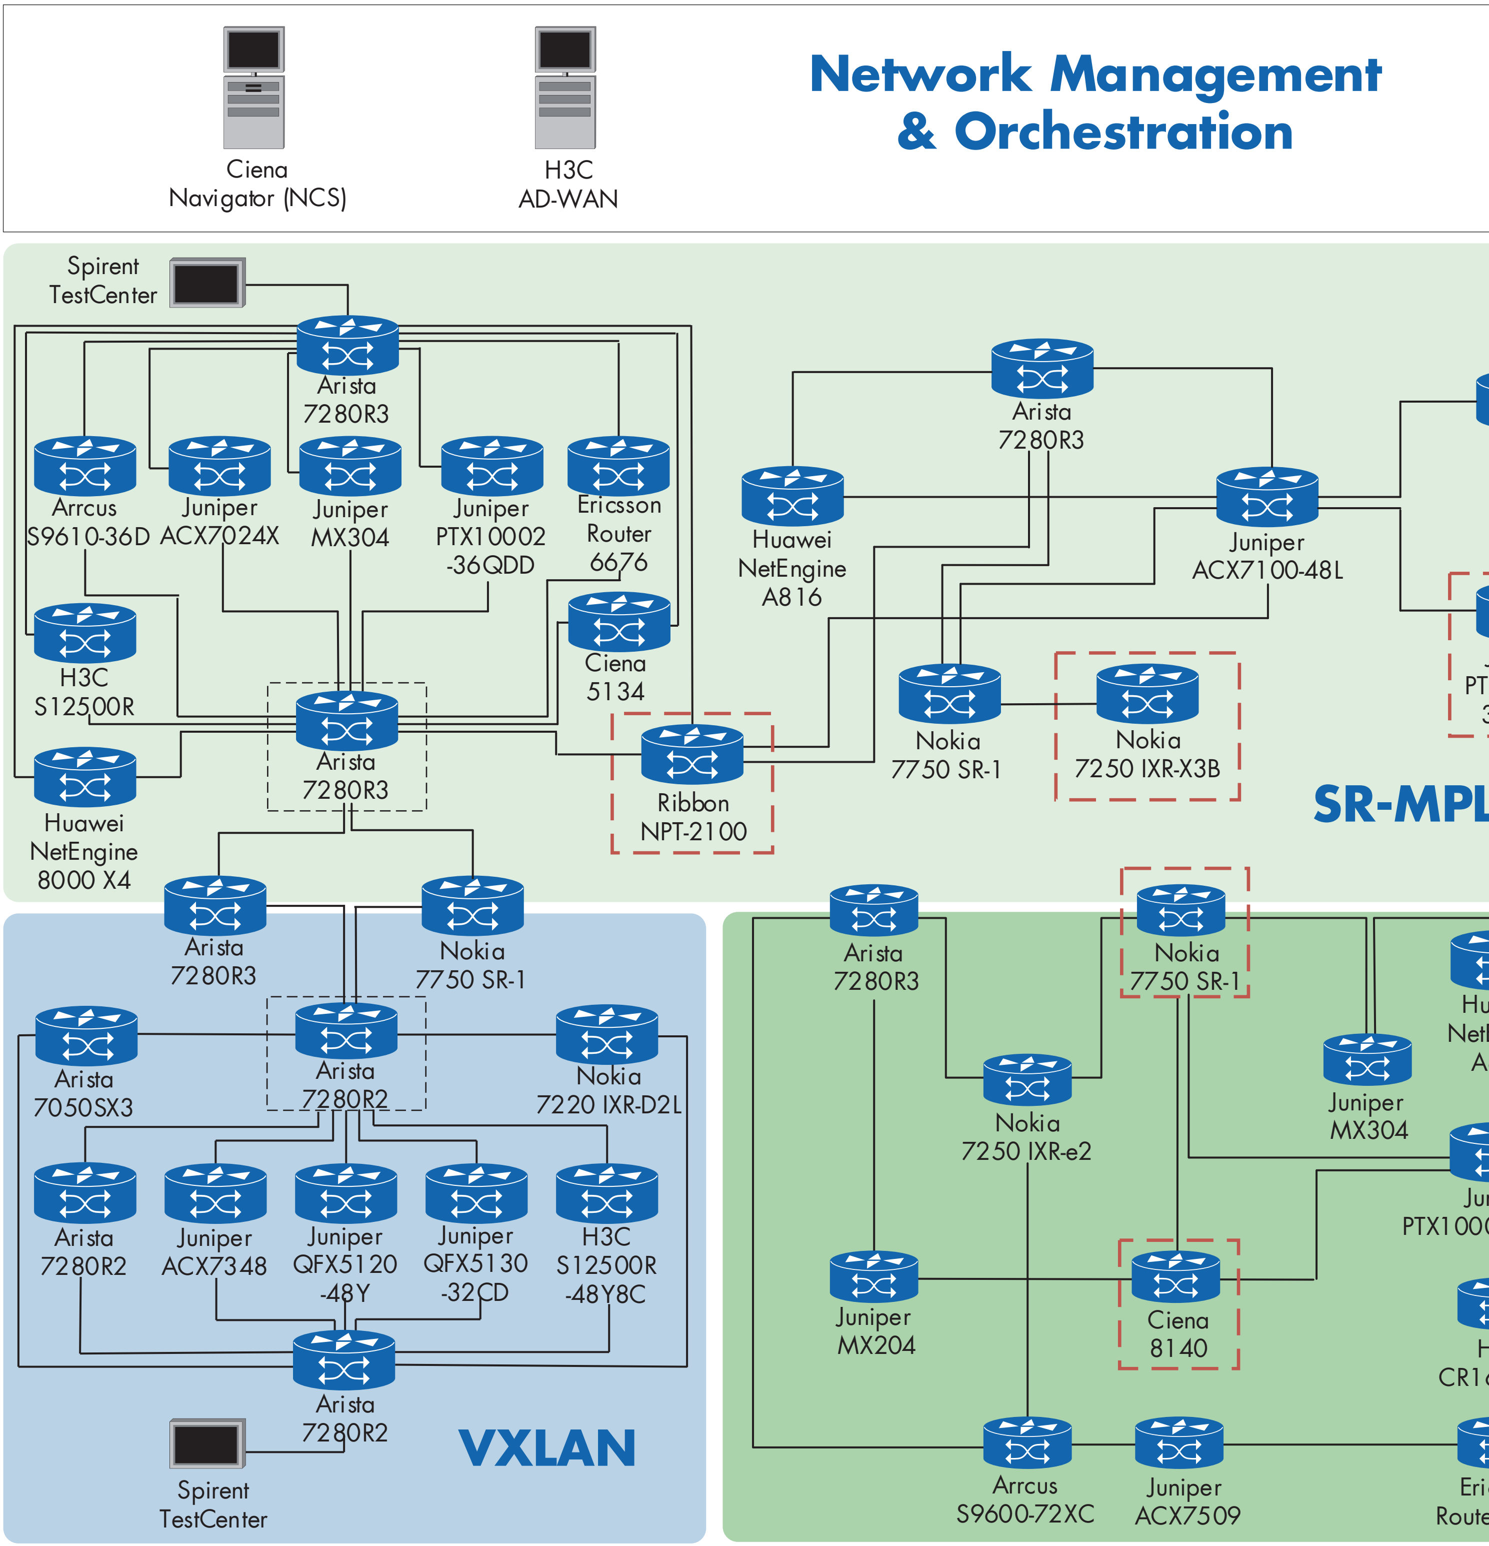

Router

Synchronization
Test Tool

Grandmaster
Clock

IP Tra
Gene

t Topology

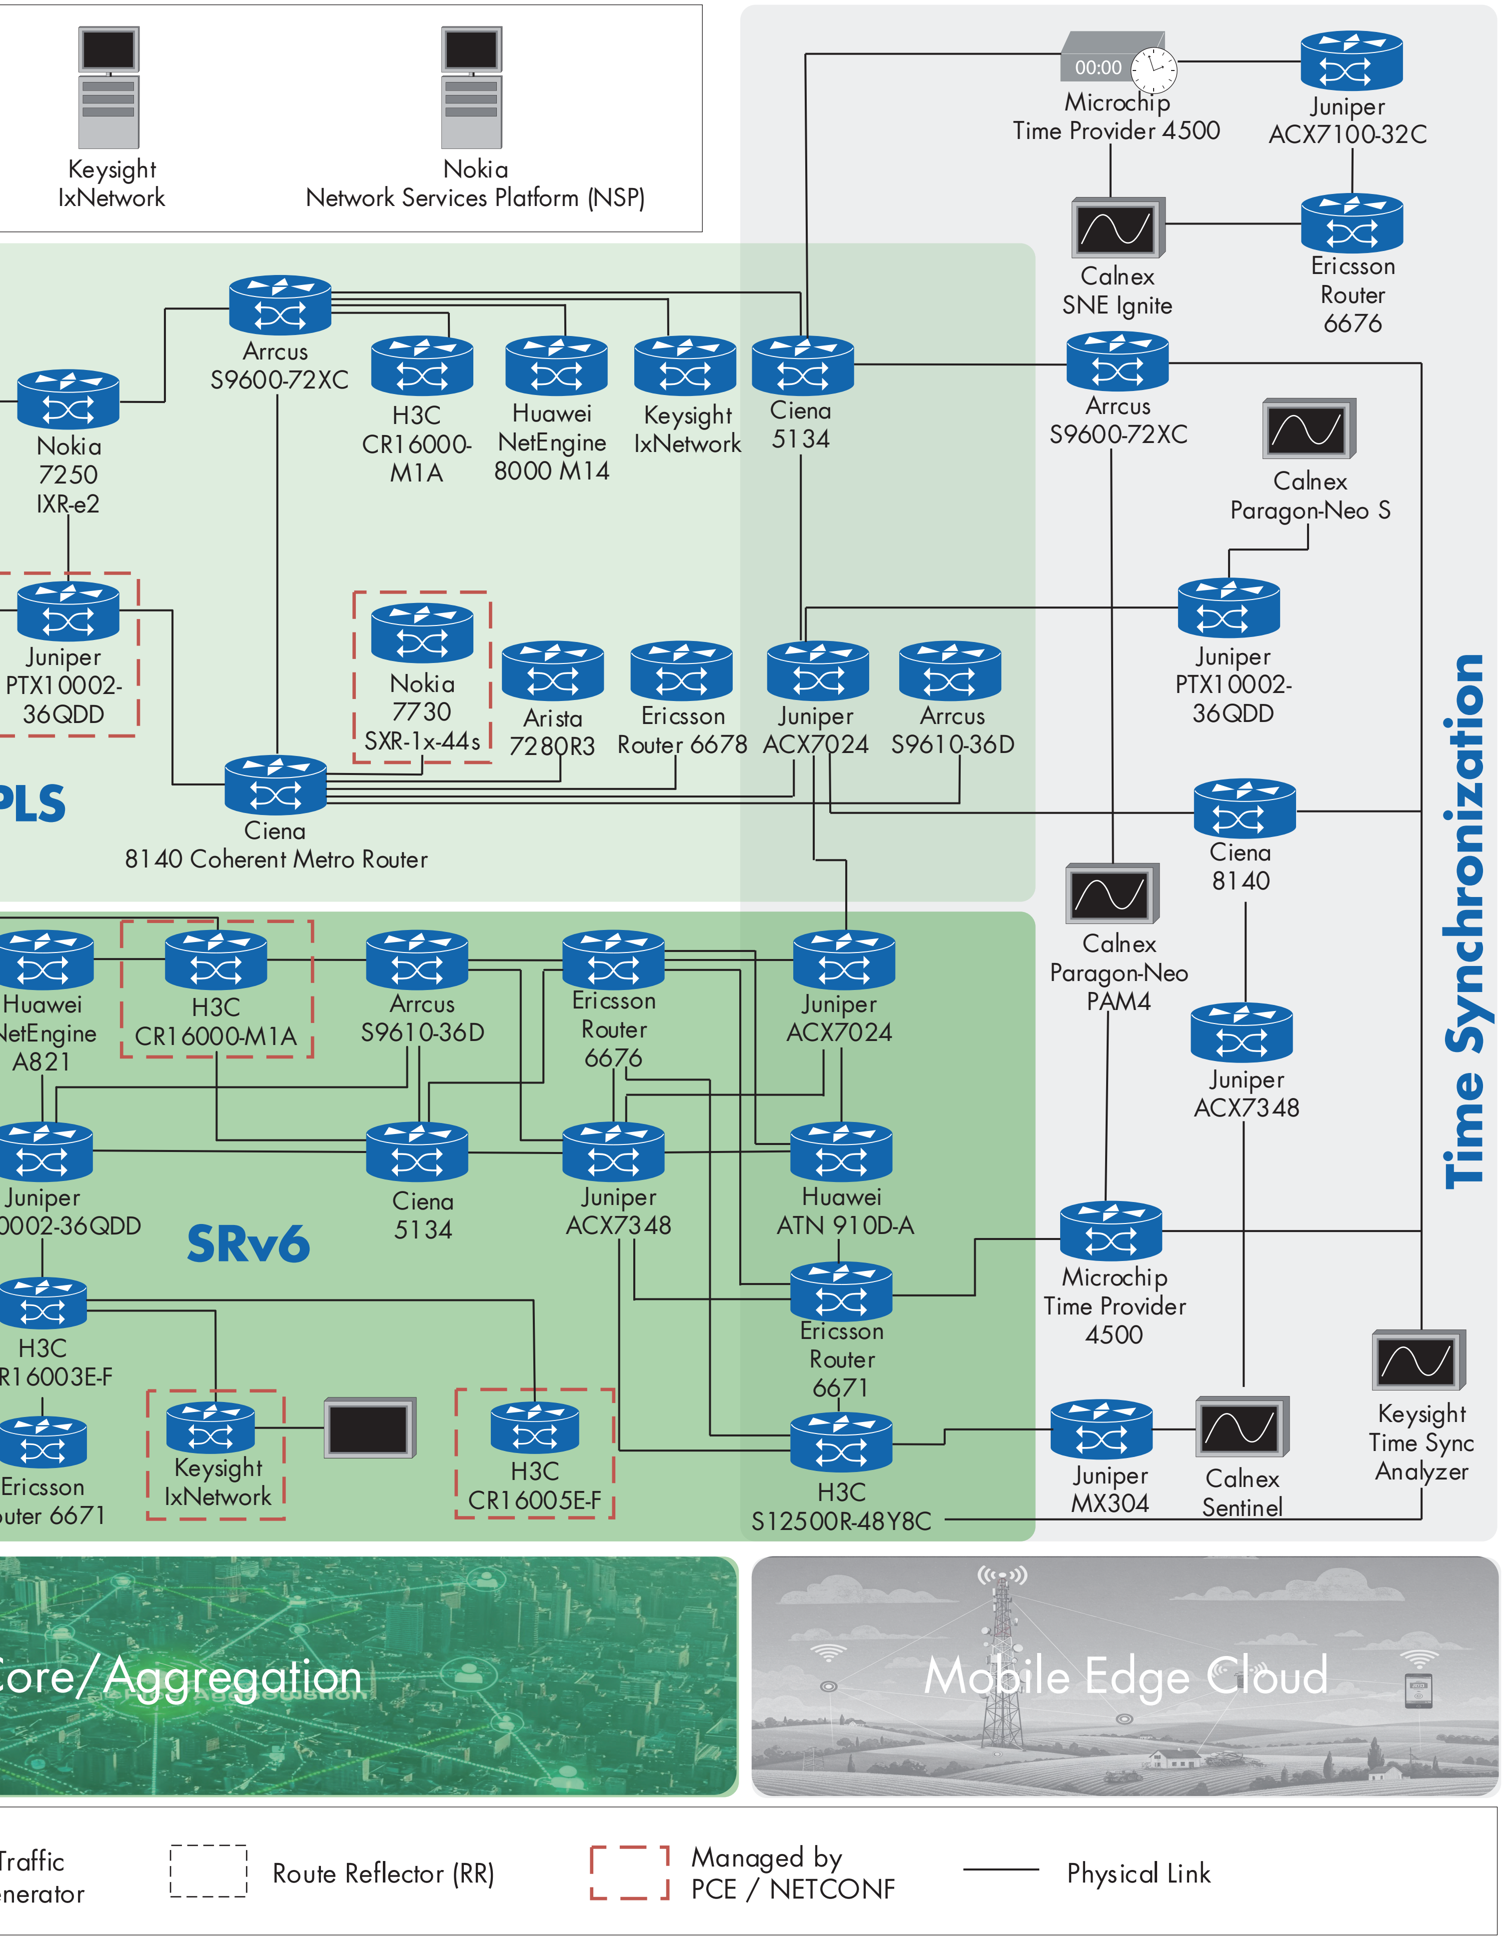
Keysight
IxNetwork
Nokia
Network Services Platform (NSP)
00:00
Microchip
Time Provider 4500
Juniper
ACX7100-32C
Calnex
SNE Ignite
Ericsson
Router
6676
Arrcus
S9600-72XC
Nokia
7250
IXR-e2
H3C
CR16000-
M1A
Huawei
NetEngine
8000 M14
Keysight
IxNetwork
Ciena
5134
Arrcus
S9600-72XC
Calnex
Paragon-Neo S
Juniper
PTX10002-
36QDD
Nokia
7730
SXR-1x-44s
Arista
7280R3
Ericsson
Router 6678
Juniper
ACX7024
Arrcus
S9610-36D
Juniper
PTX10002-
36QDD
PLS
Ciena
8140 Coherent Metro Router
Ciena
8140
Calnex
Paragon-Neo
PAM4
Huawei
etEngine
A821
H3C
CR16000-M1A
Arrcus
S9610-36D
Ericsson
Router
6676
Juniper
ACX7024
Juniper
ACX7348
Juniper
0002-36QDD
Ciena
5134
Juniper
ACX7348
Huawei
ATN 910D-A
SRv6
Microchip
Time Provider
4500
H3C
R16003E-F
Ericsson
Router
6671
Keysight
Time Sync
Analyzer
Ericsson
uter 6671
Keysight
IxNetwork
H3C
CR16005E-F
H3C
S12500R-48Y8C
Juniper
MX304
Calnex
Sentinel
Time Synchronization
ore/Aggregation
Mobile Edge Cloud
raffic
enerator
Route Reflector (RR)
Managed by
PCE / NETCONF
Physical Link

(SRv6 Full-SID/⌠SID and Flexible Algorithms), SR-Policy signaling for traffic engineering (TE), and BGP link-state for SR-Policy state reporting.

In the multipath information signaling test, the PCE successfully computed multiple paths between the headend and tail end, applying load balancing across these paths.

NETCONF/Yang Tests

Our testing efforts also focused on NETCONF. A new test executed successfully this year was the configuration of an optical 400G long-range (ZR) pluggable using NETCONF. Provisioning Ethernet- and IP-layer point-to-point and multipoint services is still an anchor test case for the vendors, and various test combinations with different controllers were executed. In addition, we tested the configuration of the NETCONF routing policies and the subscription to NETCONF notifications. We also revisited NETCONF basic operations to ensure interoperability on the protocol's foundational level.

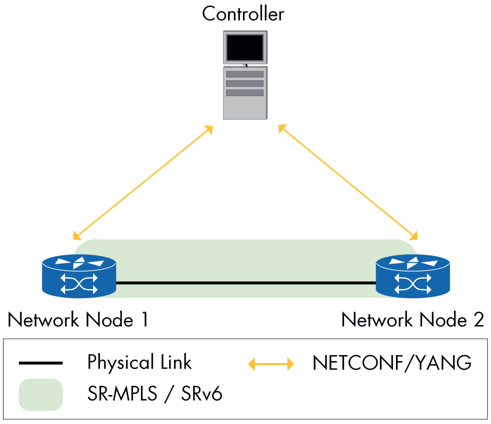

Figure 6: NETCONF General Topology

Telemetry Tests

gNMI was the main center of attention for telemetry, as we executed multiple test combinations between collectors and network nodes to gather device operational status. In one test run, we collected power consumption data from the device, taking a primary step toward creating more energy-efficient networks. We also verified interoperability at the protocol level by testing gNMI root operations, including various subscription types and notifications.

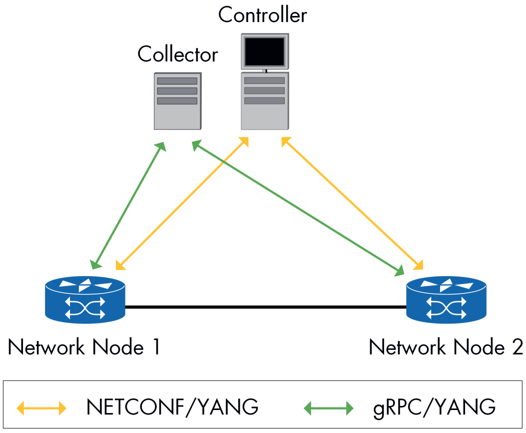

Figure 7: gNMI Telemetry Overview Topology

In addition to the positive results mentioned earlier, the tests discovered several areas where interoperability and feature support could be improved. For example, during the path computation test in a segment routing environment, we found that many vendors either support SR-TE paths as defined in RFC8664 or SR-Policy as defined in the internet draft "draft-ietf-pce-segment-routing-policy-cp." The two groups cannot interoperate with each other. Furthermore, we observed a lack of standardized implementation and support for point-to-multipoint SR policies and bidirectional SR paths, as described in the drafts "draft-ietf-pce-sr-p2mp-policy" and "draft-ietf-pce-sr-bidir-path."

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

More Orch/Mgmt Results
https://eantc.de/or&au

Segment Routing—SR-MPLS Results

Segment Routing over MPLS (SR-MPLS) is based on MPLS but uses segment identifiers (SIDs) instead of traditional MPLS label distribution protocols to create paths. SR-MPLS provides a more scalable and flexible network architecture, allowing operators to steer traffic dynamically. Traffic engineering policy support built into SR-MPLS helps accommodate various demands, like low latency or high bandwidth. At the same time, SR-MPLS runs on any equipment supporting MPLS without hardware chipset modifications, a significant difference from SRv6.

SR-MPLS in 5G x-Haul Networks

This year, our interoperability tests were built on the previous year's baseline and extended to use cases such as SR-MPLS for 5G xHaul networks.

We focused on the design of 5G xHaul. The goal was to evaluate SR-MPLS' readiness as a reliable transport for O-RAN components across the fronthaul, midhaul, and backhaul segments (see Figure 8). We measured the latency for L3VPN and VPWS for different vendor scenarios, where any best-effort traffic above 10Gbit/s was dropped, but high-priority traffic marked with DSCP/802.1p passed without issues. We tested PTP pass prioritization on access routers to ensure time synchronization packets were not dropped (see Figure 8 on the next page).

Policy- and FlexAlgo-based Traffic Steering

We evaluated traffic steering with SR-TE, testing how headend routers can direct flows according to ingress and egress colors/destinations.

Multi-Domain Routing

We also tested SR-MPLS in an Inter-AS environment to see how Anycast SIDs guide traffic to the closest routing domain border router (ASBR). We ensured the egress label stack maintained the expected sequence for the packets to reach their intended destinations appropriately (see Figure 9 on the next page).

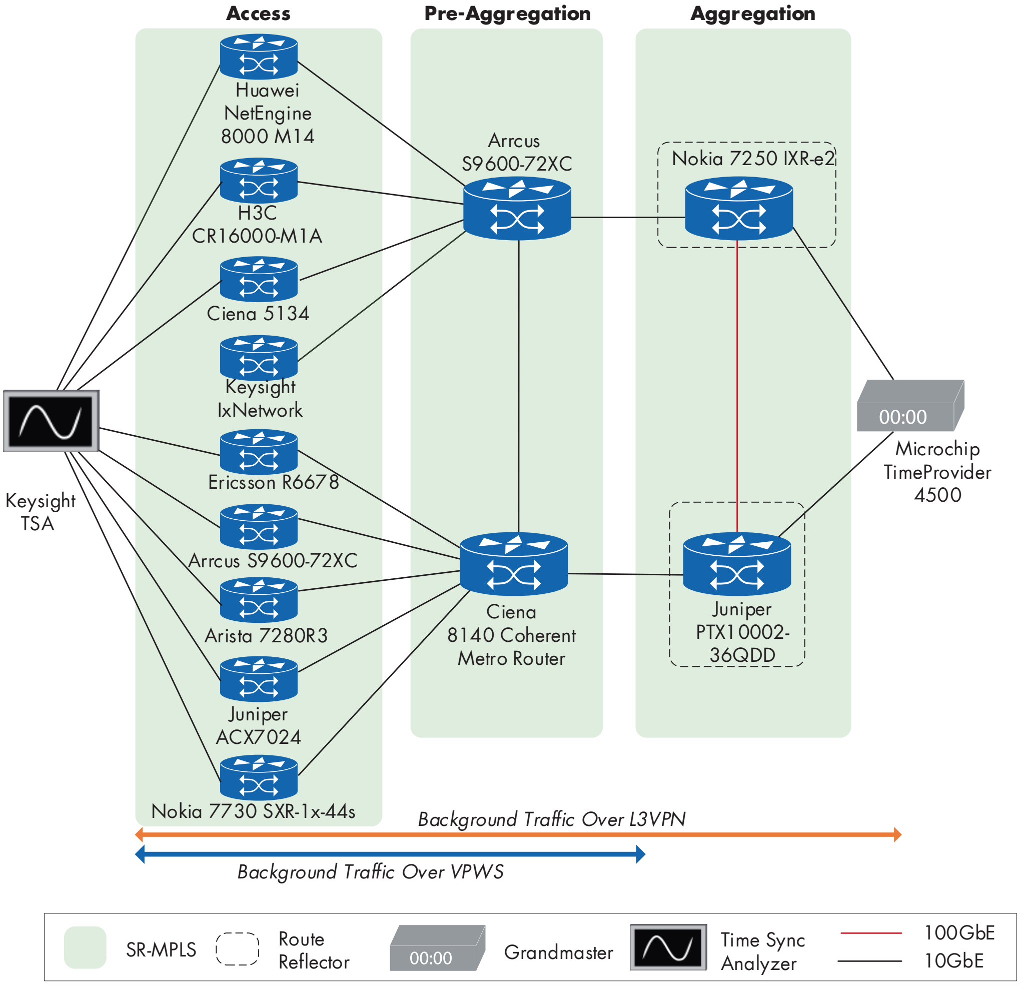

Figure 8: 5G xHaul over SR-MPLS

Continuing our FlexAlgo interoperability efforts from 2024, we validated the routers' capability to determine constraint-based paths (TE metric, link-delay and Admin Group). We added new types of constraints, such as excluding links with less than a specific capacity, exceeding maximum allowed delay, and admin-group inclusion/exclusion considered as a "reverse affinity" functionality that checks link attributes in the reverse direction.

We evaluated Seamless BFD (S-BFD) for path status monitoring. Using S-BFD, the devices under test could detect failures and swiftly failover to a backup path, based on SR-TE policies. Backup paths can be configured based on network performance parameters, such as bandwidth and delay. Such configuration options are helpful in situations where the standard IGP metric does not serve well.

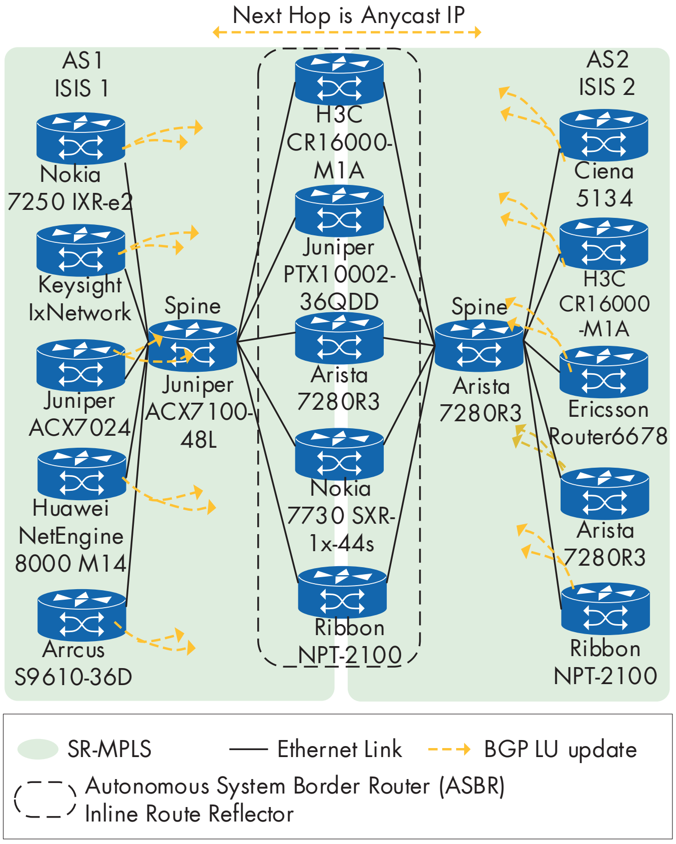

Figure 9: Inter-AS Option C using Anycast

Further SR-MPLS Tests

Other important tests included multi-vendor SR-MPLS connectivity via 400ZR/ZR+ pluggables, including support of topology-independent loop-free alternate (TI-LFA) with micro-loop avoidance for local and remote protection and Shared Risk Link Group (SRLG) scenarios. Not all vendors participated in every test case; however, many different platforms worked well together. These evaluations confirmed operational advantages of SR-MPLS in large-scale deployments.

Finally, we reconfirmed the multi-vendor interoperability of L3VPN services over SR-MPLS transport in new combinations; both IS-IS and OSPF were used to advertise SIDs.

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

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

Segment Routing—SRv6 Results

Segment Routing over IPv6 (SRv6) has advanced significantly in recent years, establishing itself as an entire transport network solution. An increasing number of SRv6 deployments worldwide reflects the advantages of the technology's pillars: end-to-end micro-segmentation, service scalability, SLA enforcement, and programmability.

2025 marks the seventh time we have covered SRv6 in the annual EANTC interoperability test event. We have progressed from fundamental protocol aspects to EVPN services and high-availability support over compressed segment IDs ("μSIDs")—all this with a growing number of supporting vendors. This year, we extended the test scope to include a 5G x-haul transport use case, assessing the readiness of participating SRv6 implementations to meet the strict 5G fronthaul requirements.

SRv6 Services Testing

To create a common baseline with all participants, we evaluated the interoperability of SRv6 regarding L3VPN, EVPN, SRv6 features, and SRv6/SR-MPLS interworking using the topology in Figure 10.

These tests included L3VPN over SRv6 using both compressed segment IDs (μSID) and Generalized SID (G-SID). For the second consecutive year, all participating vendors—Arista, Arrcus, Ciena, Ericsson, Juniper, H3C, Huawei, Keysight, and Nokia—supported μSID. However, only two H3C and Keysight implemented G-SID.

We also evaluated BGP IPv4/IPv6 Global Routing Tables over SRv6. The EVPN tests covered ELANs, EVPN RT5, and VPWS single-homed services. Other essential tests evaluated SRv6's Unreachable Prefix Announcements (UPA), Flexible Algorithms, and Locator Summarization.

SRV6 in 5G xHaul

5G fronthaul connections commonly use Ethernet protocol-based eCPRI (evolved Common Packet Radio Interface specification), typically via direct fiber links due to the strict performance requirements (< 150 μs latency and near-lossless transport). Because 5G networks grow in the number of cell sites and throughput, switched infrastructures promise efficiency gains. In our testing, we verified if a routed SRv6 solution can meet the strict fronthaul latency requirements, even when the links become congested.

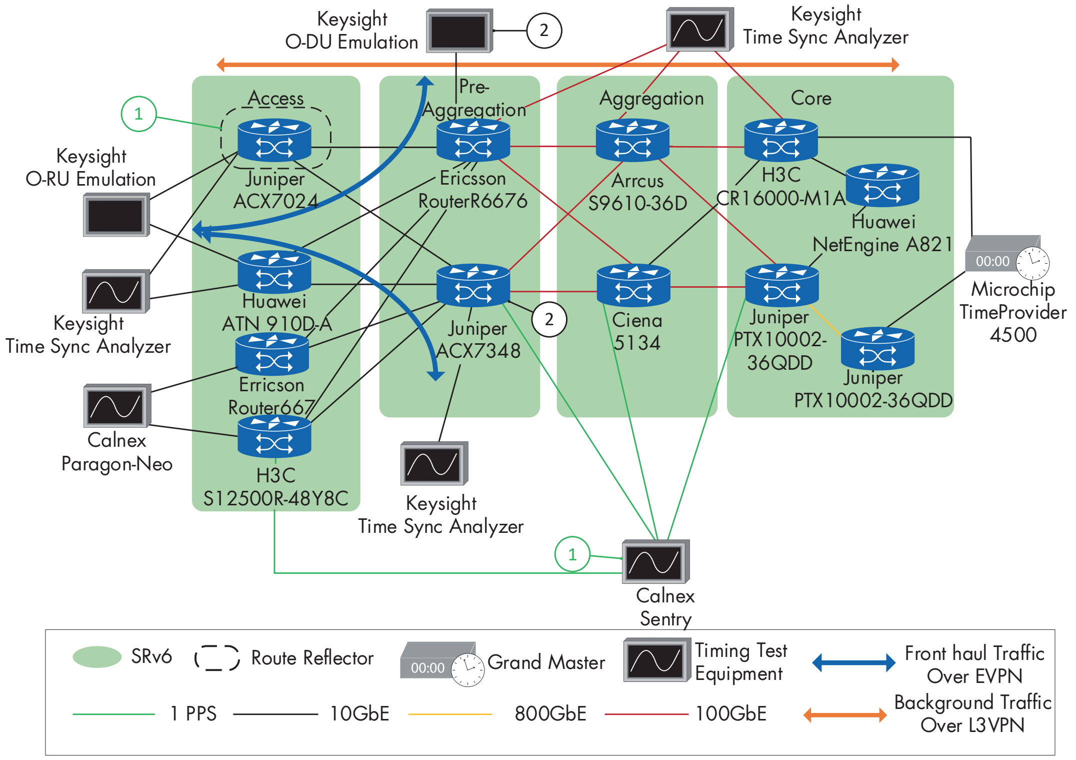

Figure 11: 5G xHaul over SRv6

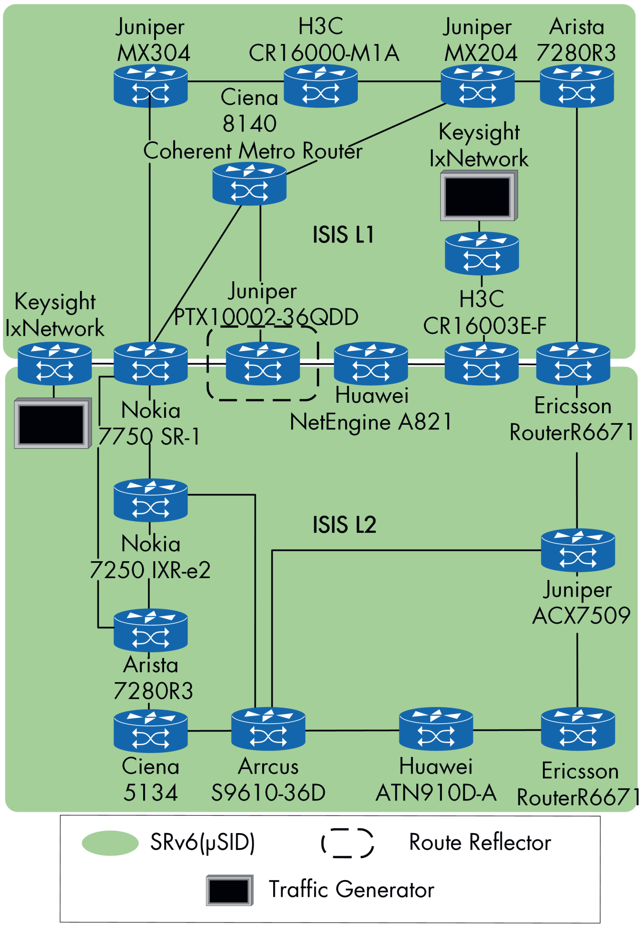

Figure 10: SRv6 Baseline Test Setup

Figure 11 depicts a typical O-RAN xhaul topology, consisting of the segments for access, pre-aggregation, aggregation, and transport core.

5G xHaul Latency Requirements/Classification

O-RAN Alliance WG9 defines transport latency thresholds, ensuring compliance with strict timing constraints required for fronthaul deployments. These latency budgets must be adhered to for optimal performance.

For instance, achieving standard NR performance requires less than 100 μs one-way latency between the O-RU (Radio Unit) and O-DU (Distributed Unit).

In these tests, we deployed the ITU-T G.8275.1 Telecom Profile for PTP synchronization and implemented packet classification using IPv6 Traffic Class (TC) and Ethernet Priority Code Point (PCP) bits. Additionally, we applied IEEE 802.1CM Time-Sensitive Networking (TSN) Profile A, specifically designed to optimize fronthaul transport performance.

Our tests confirmed that even under heavy network oversubscription, with multiple best-effort and eCPRI flows generated simultaneously, clock synchronization remained stable, ensuring reliable timing for fronthaul operations. Additionally, the one-way latency of eCPRI streams consistently met the required limits, staying within the 26 µs and 51 µs budgets defined by O-RAN WG9. These results validate the feasibility of SRv6-based transport in meeting the stringent latency requirements of 5G xHaul deployments.

SRv6 Slicing Support

Another core capability of SRv6 networks is the creation of network slices to partition the physical network resources into multiple and different logical networks, each slice dedicated to a specific service or customer.

Our test with H3C, Juniper, and Keysight verified that the ingress node could encode the slice ID in the source address and enforce a policy that guarantees bandwidth to each slice. Then, a transit P node could recognize the slice ID in each packet, apply traffic control mechanisms, and perform path selection based on the slice-specific forwarding policies.

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

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

Time Synchronization Test Results

Time synchronization requirements for 5G/6G mobile networks and other use cases are evolving. Synchronization implementations must keep pace with the increasing precision, complexity, scale, and security requirements. This year, we implemented new time synchronization test cases to verify the accuracy, stability, and interoperability of different time synchronization mechanisms and devices. All tests were based on ITU-T and IEEE standards and O-RAN Alliance specifications.

This year, nine vendors participated in the Time Synchronization test area: Arrcus, Calnex Solutions, Ciena, Ericsson, Huawei, Juniper Networks, Keysight Technologies, Microchip Technology, and H3C Technologies. They extensively tested and validated their time synchronization solutions.

Time Synchronization in SRv6-Enabled 5G xHaul

In traditional Centralized Radio Access Networks (C-RAN) architectures, fronthaul and mid-haul ("x-Haul") connections consisted of direct fiber cables running Ethernet framing without intermediate switches. However, modern deployments increasingly adopt packet-switched architectures, as supported by standards like ITU-T G.8271.1 and O-RAN.WG9.XPSAAS.0-R003-v08.00. The Open Radio Access Network (O-RAN) disaggregation enables fully switched/routed x-haul infrastructures, enabling better RAN transport scaling, more flexible component placement, and seamless transport network management.

For the first time at EANTC, we validated 5G x-Haul using a SRv6 architecture. This test was performed to see if the time synchronization performance in an O-RAN 5G network with SRv6 would still fulfill the requirements of 1100 ns (nanoseconds) maximum absolute time error from the Grandmaster to the output of the access nodes while having a relative time error of fewer than 130 ns between the access nodes' outputs (see Figure 12).

This test was performed with and without congested links. During this test, we had two types of traffic: best-effort background traffic and strict priority ORAN eCPRI traffic. When the links were congested, we observed packet drops;

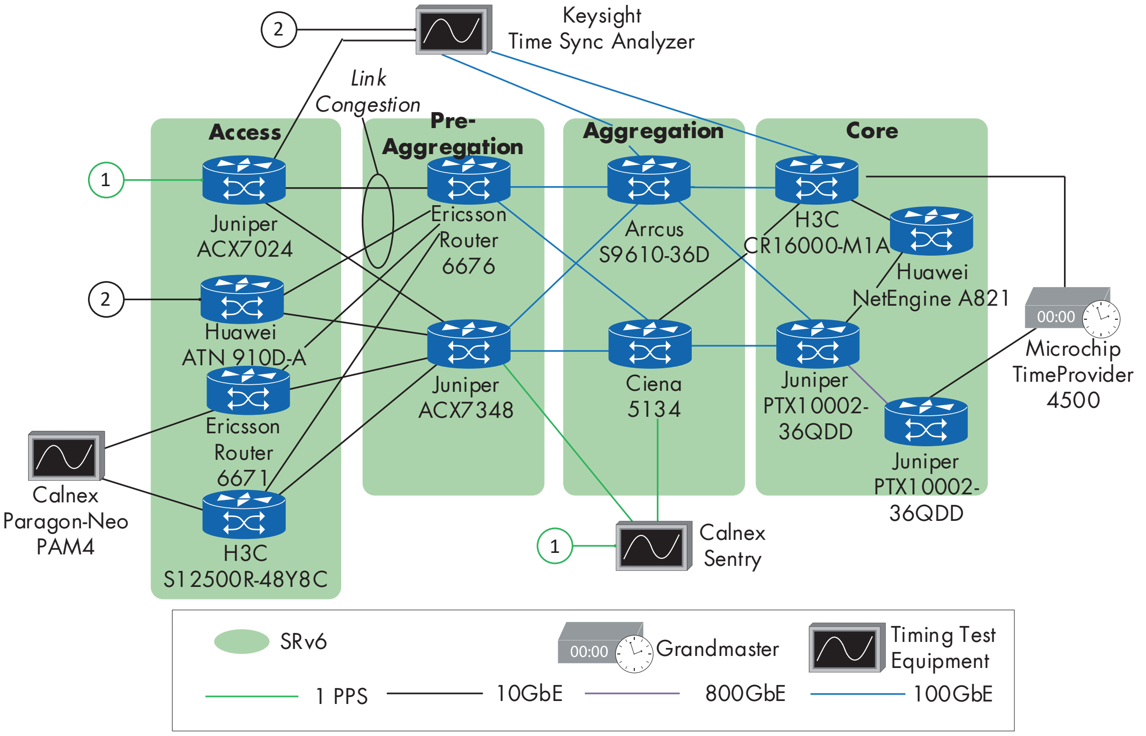

Figure 12: Time Synchronization Performance in 5G xHaul with SRv6 Architecture with Link Congestion

More TimeSync Results
https://eantc.de/timeSync

however, only the background traffic was dropped, not O-RAN eCPRI traffic, and no timing packets, proving that PTP prioritization worked correctly.

Secure PTP Transport

Precision Time Protocol (PTP) is a mission-critical technology targeted by denial of service and other attacks. So far, all PTP traffic worldwide runs unencrypted. MACsec, a technology that encrypts Ethernet links, has been around for a long time. PTP over MACsec, while not sounding necessarily complex, introduces performance challenges: PTP relies heavily on accurate packet timestamping to maintain synchronization within a network. However, MACsec requires insertion and removal of the 24-to-32-byte long MACsec header on all or some of the link frames, and the encryption and decryption can cause delay variations between the egress timestamping point and the physical link.

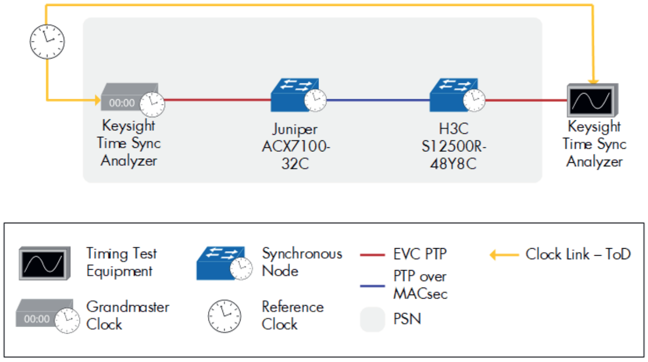

Figure 13: PTP over MACsec

This year marked our second attempt to test PTP over MACsec. Interoperability issues prevented successful completion last year. One multi-vendor combination was completed successfully this time, beating the challenges MACsec introduces to time synchronization.

In addition, we conducted PTP tests over DWDM interfaces, holdover/failover test scenarios, simulated GNSS disruption tests, delay asymmetry detection/measurement, interworking protocol and APTS tests. For the first time, we also tested Time Synchronization over 800G links, between the Juniper Networks PTX10002-36QDD and the Calnex Solutions Paragon-Neo S. Where applicable, the clocks involved complied at least with Class C requirements—often even with Class D.

With these and many other tests in the time synchronization area this year, we continued the efforts to validate and optimize multi-vendor interoperability of advanced time synchronization standards.

To access the full Test Report with additional information and details regarding the individual test cases and the vendors involved, please visit the link below or scan the QR code.

EANTC

Access the full Test Report
https://eantc.de/testreport25

This report is copyright © 2025 EANTC AG

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

EANTC AG, Salzufer 14, 10587 Berlin, Germany

info@eantc.de / www.eantc.de

V1.0 20250317