EANTC Independent Test Report

Huawei Intelligent Cloud-Network Solution Empowered by NetEngine 8000 Series, ATN Series and Network Cloud Engine

June 2021

Introduction

What constitutes a perfect transport network solution? Is it super-powerful hardware, extensive protocol support that leaves no wish unfilled, or very versatile management software? In fact, it has always been the combination of these three requirements—and most importantly, the integration of all products and software solutions towards a seamlessly working, scalable solution. These days, competition in the core and aggregation router market takes place based on integration aspects. Going even a step further, one of the main questions is whether the router manufacturer understands the requirements and use case scenarios customers have today and will have in the future. The proactive network architecture and design resulting from deep customer requirements analysis is probably of the most important success factors for a leading transport network solution.

With these considerations in mind, Huawei and EANTC have gradually moved our continued series of Huawei service provider routers from pure line card (LPU) or chassis-level throughput certifications towards much more advanced, solution-focused test reports. This report is a sequel to our January 2020 test publication about the NetEngine 8000 series and the ATN910 router family. It expands on the Huawei-focused Segment Routing over IPv6 (SRv6) architecture. Huawei is a major proponent of this architecture: Over the past years, Huawei has strongly contributed to Internet standards (maintained by the IETF) defining the SRv6 framework, and has implemented the solution throughout its family of routers for service providers and enterprises. In fact, the NetEngine 8000 series and ATN910 routers are offered to service providers, cloud providers, and large enterprises alike as a converged solution.

Huawei constructed a network with eight routers from the two families in the lab, which was used for most of the test cases. This setup with a core, a dual-router aggregation, access ring, and data center-like service termination proved useful to demonstrate more complex scenarios. From EANTC's point of view, the only drawback of this complex, access-to-data center scenario, was the focus on quick, functional demonstrations. The test configurations showed solutions working at small scale in a clean-cut lab environment; it is not possible to extrapolate the results to deployments of realistic scale. Specifically, the complex interdependencies between service deployment at scale, manageability, and efficient resiliency functions are difficult to predict when only 2-4 ports and a single digit number of tunnels are used on a flagship core router.

Test Highlights

- → Line rate throughput on ATN910D-A edge router on 10GE, 25GE, 100GE ports with packet sizes of 256 Bytes and larger
- → Optical long-haul distance support up to 120 km using 400GE QSFP-DD
- → Optical module mix-use function for 100GE/400GE ports
- → SRv6 header compression support
- → MPLS-to-SRv6 migration support by tunnel splicing
- → Multicast MVPN-to-BIERv6 migration support
- → FlexAlgo support for 5G slicing
- → TI-LFA failover time of less than 1.4 ms, zero-loss recovery
- → BNG support with up to 64,000 IPoE and 64,000 PPPoE tunnels on NetEngine 8000 M14 and NetEngine 8000-F1A
- → Network solution provisioning, management and closed-loop of DevOps with Agile Open Container (AOC) model using Huawei iMaster NCE-IP

From all the test results, it looks like Huawei would have nothing to fear; we look forward to an opportunity to validate more complex Huawei router performance scenarios in the future.

Other important design aspects demonstrated by Huawei in this engagement were: The uptake of 400GigabitEthernet; long-distance optics; support for BIER multicast; and options for smooth transition from legacy MPLS to SRv6 transport and EVPN services.

Additionally, this test report continues with our ongoing review of iMaster NCE-IP or in short NCE-IP, the Huawei controller for large transport networks. NCE-IP is part of Huawei's strategy to increase the level of automation for network provisioning, maintenance, and optimization. In this test engagement, Huawei demonstrated advanced SRv6 provisioning and automated policy management techniques, some of them involving performance monitoring. Telemetry solutions gathering live usage data and subsequently adapting policies are an innovative mechanism to optimize network usage and quality of experience.

All in all, Huawei selected a wide range of advanced transport and orchestration features, and demonstrated them to EANTC successfully in this engagement. The NetEngine 8000 router family and the NCE are definitely maturing; we found the expanded range of protocols and services comprehensive and reassuring specifically for SRv6 deployments.

Test Bed

To support a wide range of test cases covered in this report, Huawei built a test network (shown in Figure 1) consisting of eight devices under test (DUT) in total:

Six NetEngine 8000 family routers and two ATN family routers. These routers together formed a representative network in the lab with a core cluster, an aggregation network, and an access network. Each of the routers had different port speed capabilities as seen in Table 1. During the test, 1310nm single-mode optics were used for 10GE, 850nm single-mode optics were used for 25GE, four wavelengths of 1295.56nm, 1300.05nm, 1304.58nm, 1309.14nm single-mode optics were used for 100GE and 1310nm single-mode optics with the 8-lane solution were used for 400GE. For the tests, only two ports of the DUT were used; for some other, all the ports of DUT were used.

For this test campaign, Huawei chose to use an unpublished engineering software version (8.201) for the NetEngine 8000 family routers to benefit from the latest implementation features. Huawei confirmed that the features in this engineering software will be included in a future software release for ATN planned as V300R007C00, and a future software release for NetEngine 40E and NetEngine 8000 planned as V800R013C00.

The unpublished NCE management software features used in this test will be published in version V100R020C10, according to Huawei. Huawei mentioned that all software versions will probably be published until April 2021. To accelerate test preparation and execution, Huawei setup multiple copies of the test bed in different lab locations in Beijing, Nanjing, and Shenzhen (all in China). EANTC checked that the hardware and software versions were reported accordingly by the software. Since the software version was not published at the time of testing, it was not possible to download software from an official support server to validate the publication / support status. Huawei promised a software release available to customers by the end of April 2021.

Huawei configured separate test topologies for the ATN910 forwarding capacity test cases, for the broadband network gateway network test and the 400G end-to-end test. These topologies are described under the respective sections. During Covid-19 pandemic times, the test was conducted remotely. EANTC supervised all test preparations, executions, and documentation efforts live and had access to the test tools and router configurations. We also witnessed hardware reconfigurations—e.g. for the optical link tests—via video transmitted from the Huawei lab.

In the tests, we used a Huawei-supplied Ixia AresONE traffic generator with a QSFP-DD-400GE 8-port interface card and software version IxOS 9.05.1920.8. Additionally, an Ixia XGS12 test tool with 8x100GE and 16x10GE ports and software version IxOS 9.00.1900.10 was used.

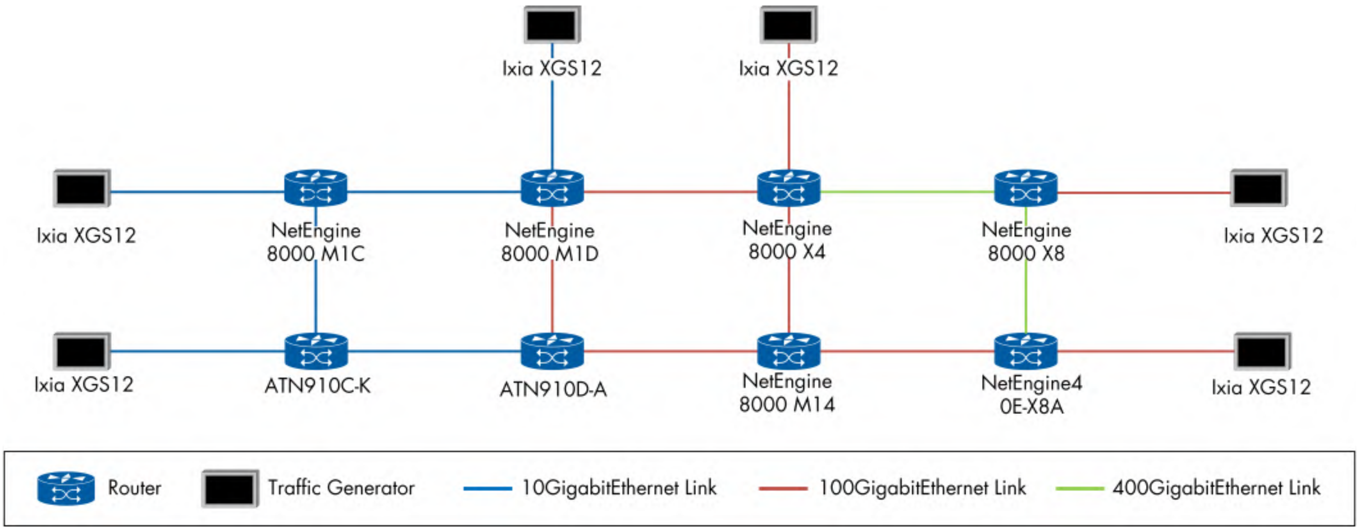

Figure 1: Test Network

Short Name	Rputer Type	Number of Ports In Test Configuration					Com-ponent	Hardware Version	Boot-ROM	CXP/MPU Software
		400GE	100GE	25GE	10GE	GE				
DUT1	NetEngine 8000 X8	8	48	0	0	0	MPU	GPMPUB REV C	19.07	8.201 (V800R013C00) Version 84
							LPU	GPL4TA REV CPIC0 PCB: GPLI40CQES1 REV A	19.07	Version 112
							LPU	GPL4TB REV DPIC0 PCB: GPLI8KQ8CQES1 REV D	19.07	Version 112
DUT2	NetEngine 40E X8A	2	8	0	0	0	MPU	CR57RPUH REV B	08.15	8.201 (V800R013C00)
							LPU	CR57LPUF1TQ REV BPIC0 PCB: CR57E8VBA REV APIC1 PCB: CR57E1KBA REV A	30.1	
DUT3	NetEngine 8000 X4	8	48	0	0	0	MPU	GPMPUB REV C	19.07	8.201 (V800R013C00) Version 122
							LPU	GPL4TA REV CPIC0 PCB: GPLI40CQES1 REV A	19.07	Version 112
							LPU	GPL4TB REV DPIC0 PCB: GPLI8KQ8CQES1 REV D	19.07	Version 112
DUT4	NetEngine 8000 M14	2	4	0	30	0	IPU	DP51CPUA REV B	08.15	8.201 (V800R013C00)
DUT5	NetEngine 8000 M1D	0	4	16	8	0	CXP	CR81IPU800AS REV A	08.99	8.201 (V800R013C00)
DUT6	ATN910D-A	0	4	16	8	0	CXP	ANG2CXPA REV A	08.99	8.201 (V300R007C00)
DUT7	NetEngine 8000 M1C	0	0	0	16	12	CXP	CR81IPU160CS REV A	08.99	8.201 (V800R013C00)
DUT8	ATN910C-K	0	0	0	16	12	CXP	ANG1CXPPS REV APIC3: ANG3MO1 REV A	08.99	8.201 (V300R007C00)

Table 1: Devices Under Test

Test Results - Performance

In this session, Huawei demonstrated the performance of a new type of ATN910D-A chassis and of new 400GE and mixed 400GE/100GE optical modules for the NetEngine 8000 series routers.

ATN910 Forwarding Capacity

With the common use of 5G and increasing bandwidth use in consumer scenarios for 4K video, gaming, etc., network traffic the access area is constantly increasing. The ATN910D-A is a 1U fixed box with a nominal maximum throughput of 880 Gbit/s and abundant features, according to Huawei. Huawei positions this router in the access area for both mobile and fixed network use case scenarios. EANTC verified Huawei's data plane performance claim for maximum throughput of the ATN910D-A.

The ATN910D-A chassis has 4x100GE, 16x25GE, and 8x10GE ports with a total bandwidth capacity of 880 Gbit/s. This scenario was traditionally tested with all ports utilized, in a partial-mesh scenario with native IP routing. Huawei configured the router with three groups of same speed (i.e. 10GE, 100GE, and 25GE) for port-to-port VPNs. All router ports were directly connected to the Ixia XGS12 test tool. Huawei configured a partial-mesh scenario where IPv4 and IPv6 traffic was sent from each tester port towards all the router ports in the same speed port group. This configuration created 12 port-to-port combinations on 100GE ports (4x3), 240 combinations on the 25GE ports (16x15), and 56 combinations on the 10GE ports (8x7). This way, the router could be tested for full utilization across all 28 ports.

The forwarding capacity test was carried out following RFC 2544 benchmarking methodology. This RFC specifies multiple test runs with individual packet sizes with different individual packet sizes. RFC 2544 tests allow qualification of the maximum forwarding performance well because the number of packets per second varies, inversely proportional to the packet size chosen for the individual test run. Huawei directed EANTC to use frame sizes of 256 bytes and larger for the test. Traffic was sent across the partial mesh configuration, and a test duration of 120 seconds per frame size was chosen.

The throughput test results are shown in Table 2. The router fulfilled Huawei's throughput claims with up to 879 Gbit/s. We tested it with up to 398.55 Million packets per second (Mpps).

The latency observations during the test were as shown in Figure 3, 4, and 5. The latency on 100GE ports met our expectations. Maximum latencies on the 25GE ports for some packet sizes, and on the 10GE ports for all packet sizes were higher than expected (10 times than minimum latency) - no buffering was expected. We had ruled out unsolicited control plane packets as a potential source.

Huawei explained that "the ATN910D-A router has a 16 MB default built-in buffer size per port queue; the buffer size can be configured up to 1.9 GB. With this buffer size, maximum latency per 10GE port up to 1.52 s and per 100GE up to 152 ms will not cause any packet loss."

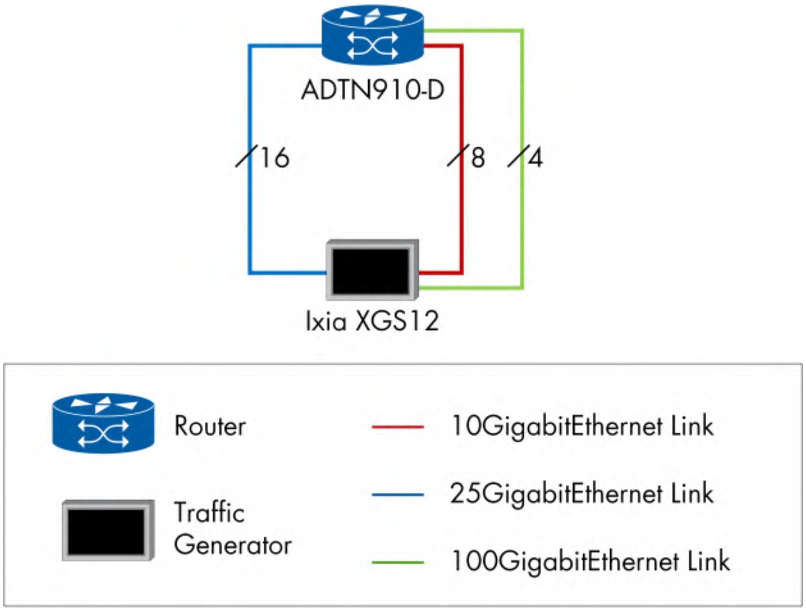

Figure 2: ATN910-D Test Topology

Packet Size	Throughput per 25GE Port, per Direction	Throughput per 10GE Port, per Direction	Throughput per 100GE Port, per Direction	Packet Loss	Throughput per DUT	Packet Forwarding Performance per DUT
256	24.9 Gbit/s	9.9 Gbit/s	99.9 Gbit/s	0	879.99 Gbit/s	398.55 Mpps
512	24.9 Gbit/s	10 Gbit/s	99.9 Gbit/s	0	879.90 Gbit/s	206.75 Mpps
1024	24.9 Gbit/s	9.9 Gbit/s	99.9 Gbit/s	0	879.88 Gbit/s	105.35 Mpps
1280	24.9 Gbit/s	9.9 Gbit/s	99.9 Gbit/s	0	877.76 Gbit/s	84.40 Mpps
1500	24.9 Gbit/s	9.9 Gbit/s	100 Gbit/s	0	879.70 Gbit/s	72.35 Mpps
2048	25.0 Gbit/s	9.9 Gbit/s	99.9 Gbit/s	0	879.64 Gbit/s	53.17 Mpps
4096	25.0 Gbit/s	9.9 Gbit/s	100 Gbit/s	0	879.17 Gbit/s	26.70 Mpps
9000	24.9 Gbit/s	10 Gbit/s	99.9 Gbit/s	0	878.90 Gbit/s	12.18 Mpps

Table 2: ATN910D-A Traffic Forwarding Performance Results

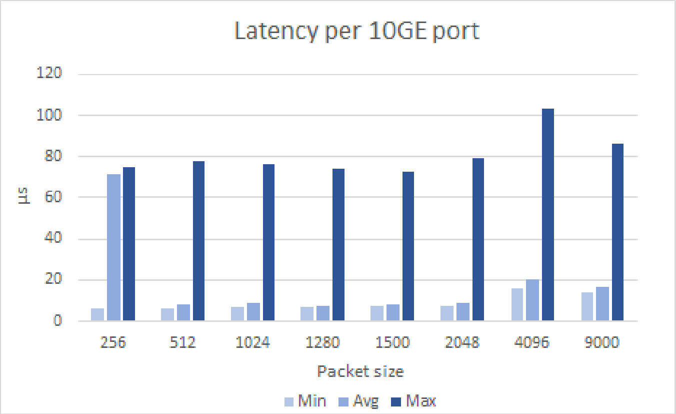

Figure 3: ATN910D-A Traffic Forwarding Latency Results per 10GE Port

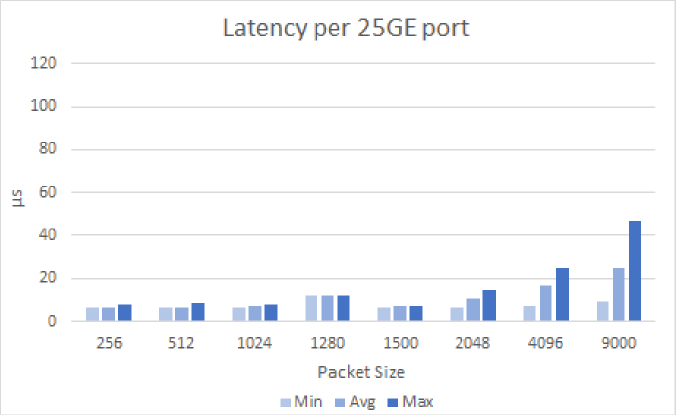

Figure 4: ATN910D-A Traffic Forwarding Latency Results per 25GE Port

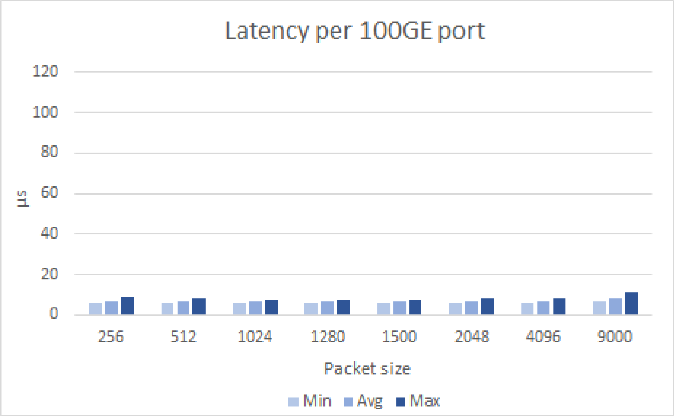

Figure 5: ATN910D-A Traffic Forwarding Latency Results per 100GE Port

400GE Optical Long-Haul Module

Long-distance optical modules are an important requirement for some large countries, for example Russia, Canada, and India. Traditionally, routers have been supplemented by optical transport network (OTN) equipment for this purpose, handling the long-distance transport. This solution is cumbersome and expensive as it involves an additional network layer. In some cases for the latest high-speed Ethernet interfaces such as 400GE, OTN equipment may not be widely available at all. As long as the distances between active router equipment remain below 120 km, it is possible to use long-distance optics directly in routers if supported. In this test case, we verified Huawei's claim to support up to 120 km distance with Huawei 400GE QSFP-DD optical modules. The QSFP-DD form factor is used in the line processing unit (LPU) under test; it is much smaller than the traditional CFP-2 optics. Due to its small size, an additional external optical amplifier is required to implement the actual optical conditioning functions for the transmitter and receiver. The QSFP-DD implements the important electronic signal recovery functions such as forward error correction (FEC).

In our test, Huawei connected two 400GE ports of a Net-Engine 8000 X8 router and two 400GE ports of a NetEngine 8000 X4 router with single-mode 400GE QSFP-DD optical modules, using a low-loss fiber cable. Long-haul distance fiber coils of 10 km, 40 km, 80 km, and 120 km were inserted into the link in the lab. For 80 km and 120 km distances, Huawei deployed the optical amplifier TN13OAU1 in a Huawei OptiX OSN 9800 WDM rack. The TN13OAU1 is an Optical Access Unit (OAU) used for internal amplification of the optical signal. We transmitted bidirectional near-line speed traffic using an Ixia AresONE with QSFP-DD-400GE ports on the local network side.

For all the long-haul distances used in the test, we observed zero traffic loss at 99% link speed throughput and latency of less than one millisecond as shown in Table 3. The Huawei 400GE QSFP-DD fulfilled the claims, providing lossless forwarding up to 99% link speed within the expected latencies.

Long-haul Distance	**Bi-Directional Throughput Tested**	**Packet Loss**	**Minimum Theoretical One-Way Latency**	**Measured One-Way Latency (min/avg/max)**
10 km	779.93 Gbit/s	0	50 µs	67/73/86 µs
40 km	779.93 Gbit/s	0	200 µs	215/221/233 µs
80 km	779.93 Gbit/s	0	400 µs	420/427/451 µs
120 km	779.93 Gbit/s	0	600 µs	619/626/643 µs

Table 3: 400G Optical Module long-haul Performance Test Results

400GE/100GE Optical Module Mix-Use Function

As the industry starts transitioning core and aggregation connections from 100GE to 400GE, a smooth migration strategy is needed. Initially, it is beneficial if 400GE ports support 100GE as well so that existing 100GE connections can be moved to new line cards smoothly—without requiring immediate speed upgrades to 400GE. Huawei provides a mixed-use capacity optical module to enable operators to keep using 100GE links for now. As traffic grows, service providers can then change link speeds to 400GE without changing the board; just the optics need to be exchanged. We tested Huawei's claims for smooth migration from 100GE to 400GE. Two 100GE/400GE mixed-speed ports of NetEngine 8000 X4 were equipped with 100GE QSFPs and connected directly to the Ixia AresONE. For the test, Huawei configured the native IPv4 addresses on the test ports to forward packets. In this scenario, the Ixia test tool was configured to transmit bidirectional IPv4 iMIX traffic for five minutes as shown in Table 4 with 97.5 % of link speed.

As expected, there was no packet loss. In the next step, the router's two 100GE/400GE mixed-speed ports were reconfigured to 400G mode, and the QSFPs were exchanged with 400GE optics. Another throughput test run was conducted for five minutes. For both the 100G and 400G speeds configured on DUT ports in the test, we observed no traffic loss, as shown in Table 5 in detail.

Frame Size (Bytes)	Weight IPv4 IMIX
64	3
78	-
100	26
373	6
570	5
1300	6
1518	16
9000	1

Table 4: Custom IMIX

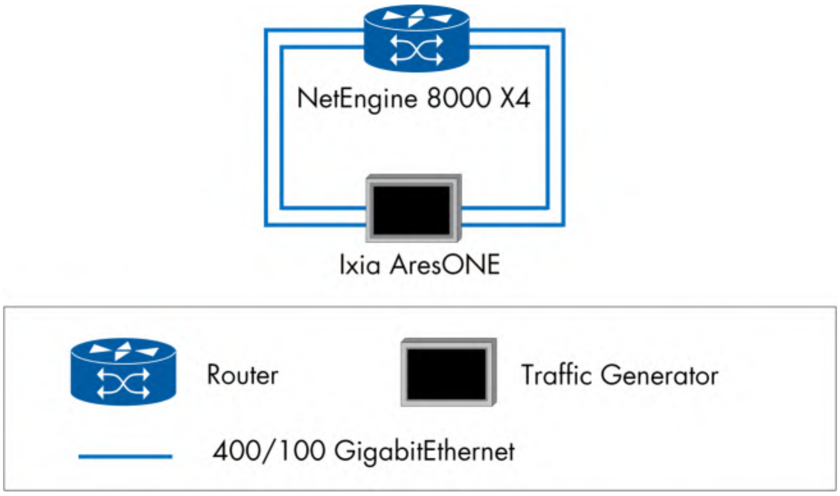

Figure 6: 400GE/100GE Optical Module Mix-Use Test Topology

Port Speed	Optics Used	Bi-Directional Throughput Tested	Packet Loss %	Latency (min/avg/max)
100GE	QSFP28 100GBASE-LR4	194.98 Gbit/s	0	8/14/15 μs
400GE	QSFP-DD 400GBASE-LR8	779.93 Gbit/s	0	9/12/21 μs

Table 5: 400G/100G Optical Module Mix-Use Test Results

Test Results - Determined Experience

SRv6 is a technology standardized in the Internet Engineering Task Force (IETF) since 2016. It provides disruptive innovation, eliminating the underlying MPLS or VXLAN transport required for legacy Segment Routing, and enabling seamless end-to-end connections from data centers through wide-area networks to other data centers. Huawei strongly supports this technology, and it has been a topic of testing in previous EANTC reports for the NetEngine 8000 and NCE. We have already previously validated the provisioning of SRv6 tunnels and the creation of EVPN services over SRv6 in single-vendor and multi-vendor scenarios. We have witnessed the configuration of SRv6 Traffic Engineering tunnels, operations support for SRv6 ping and trace route, and resiliency failover using Fast Reroute (FRR) and Topology-Independent Loop-Free Alternate (TI-LFA). In this test campaign, Huawei presented additional advanced SRv6 functions to the EANTC team. These functions are complementing the operational efficiency and other aspects of the SRv6 implementation.

SRv6 Compression

As part of the deployment of SRv6, ordered segment ID (SID) lists with multiple SIDs are required to steer packets along explicit paths. This function is used specifically when traffic engineering is enabled. One of the concerns of SRv6 is its overhead on the network throughput, caused by the segment headers containing long IPv6 addresses. The header length per segment is 128 bits (16 bytes). Especially when the original packet size is small and the SRv6 path has more than five hops, the forwarding efficiency will be decreased considering the segment routing header (SRH) length. This aspect of SRv6 architecture impacts the maximum usable maximum link bandwidth. To reduce the header length, Huawei has co-authored an IETF draft-cl-spring-generalized-srv6-np-02 (work in progress). This technology, in short called G-SRV6 for SRv6 compression helps to increase the forwarding efficiency in some cases. EANTC witnessed the solution at work. Using the main topology (Figure 1), Huawei configured an SRv6 policy on DUT1 for a path through DUT2, DUT4, DUT6, DUT8, to DUT7. This path contained a total of five (5) SRV6 path segments and 1 VPN segment. The path segments had the following SIDs:

- 2020:8000:8000:a2::222
- 2020:8000:8000:a2::444
- 2020:8000:8000:a2::666
- 2020:8000:8000:a2::888
- 2020:8000:8000:a2::777

We sent IPv6 traffic with small packet sizes (76 bytes) from the Ixia test tool into DUT1, from where traffic got forwarded through the SRv6 tunnel as an EVPN service, encapsulating the original IP packet. We monitored the packets on the way, analyzing the packet captures (PCAPs). The resulting packet size was 216 bytes, an overhead of 144 bytes for the SRv6 header (the VLAN ID was not forwarded across the service provider network). Next, Huawei configured SRv6 header compression on all routers involved in the test. When the Ixia traffic generation test was rerun, the SRv6 header size in the transport network dropped to 96 bytes. The results show that the service header length for SRv6 tunnels with five segments is reduced by 33 %. The total packet size was reduced from 216 to 168 bytes, by 22 %. Obviously, the forwarding efficiency is most affected for small data packets. With compression enabled, the network throughput with very long segment lists and very small 78-byte sized packets could be increased by 25 % without upgrading the network links. The compression optimizes throughput best when the packet sizes are small. Huawei confirmed that not all scenarios benefit from SRv6 header compression as much. For example, if the average packet size would be 777 bytes, the throughput could be increased by 5.2 % without upgrading the network links. Other aspects of generalized SRv6 as defined in the IETF draft mentioned above were not demonstrated in this test campaign.

Connecting EVPN Multi-Point Services with Legacy VPLS

Gradually, metro networks are evolving towards EVPN. One of the major challenges in this evolution is to transition existing MPLS networks with a large number of aggregation devices towards SDN. It is not practical to switch over from MPLS/VPLS to SRv6/EVPN at once. Instead, there is typically an extended transition phase during which conventional L3VPN, VPWS, and VPLS services continue to be used in some parts of the network. In this case, splicing between the EVPN and existing network services is required to support end-to-end service connectivity. Huawei demonstrated that this functionality is supported by the NetEngine family of routers.

In the main test topology (Figure 1), Huawei configured legacy MPLS on the link between DUT7 and DUT5 with a VPLS service between these two routers. On the rest of the path (DUT5 to DUT1 via DUT3), SRv6 was configured with an EVPN VPLS service. BGP was used both on the EVPN side and the MPLS side to exchange VPN routes.

We verified that traffic could be forwarded along the hybrid path from legacy to next-generation service, confirming Huawei's claim. It is possible to manually create end-to-end service connections spanning both network generations MPLS and SRv6.

Connecting Routed EVPN Services with Legacy L3VPNs

In a similar fashion as the Ethernet multi-point services in the previous test case, Huawei demonstrated that MPLS-based L3VPN services can be spliced with Layer 3 EVPNs over SRv6. Huawei configured a classic L3VPN over MPLS between DUT5 and DUT1 with transit hop DUT3. Additionally, a routed EVPN service was created between DUT7 and DUT5, transported over SRv6. All the necessary legacy configurations were done on the MPLS side, and BGP peering was used to exchange the service information. BGP peering was also utilized on the SRv6 side to exchange service information. The EANTC team witnessed that an end-to-end service could be created, forwarding traffic in both directions. This test was only functional and no high performance, scalability, or resiliency tests were performed.

SRv6 Multicast

Multicast has always been a complex transport topic. It continues being an important special application for service provider and enterprise networks - for video/audio broadcasting services and financial information and trading alike. A technology developed in the past few years—Bit-Indexed Explicit Replication (BIER)—is slowly gaining market share. Its implementation scales much better and is radically different from the previous generation of protocol-independent multicast (PIM) routing and multicast VPNs (MVPNs).

EANTC had witnessed functional demonstrations of Huawei's BIER implementation on the NetEngine 8000 family of routers before. This time, Huawei configured a number of add-on services to showcase the versatility of Huawei's implementation. One of the main advantages of BIER is service scale in the core and aggregation network. BIERv6 does not require intermediate routers to maintain any per-flow state. Its main differentiator is to use IPV6 as the forwarding plane. In the edge, migration of multicast services often takes longer and is of lower importance because there might be lower scalability requirements. Terminating classic Multicast VPN (MVPNv4) services at the service edge allows gradual transition. Huawei demonstrated the ability of the NetEngine 8000 software to encapsulate MVPNv4 services over BIERv6.

To this end, DUT1 was setup as a multicast sender, and DUT7 plus DUT8 were configured as multicast receivers. All three endpoints used MVPN and had the appropriate classic protocols configured between them (MP-iBGP; BGP MVPN; PIM). However, the actual transport of multicast traffic throughout the network was carried out over BIERv6.

The NetEngine 8000 routers implement the IETF draft (work in progress) draft-xie-bier-ipv6-encapsulation-09. Once everything was setup and BIER tunnels (I-PMSI, S-PMSI) got established, multicast traffic could be forwarded between the edge routers, confirming Huawei's claim of a functioning solution for the multicast service migration.

Additionally, we verified whether the MVPN to BIERv6 migration would work in more complex routing scenarios. In a first step, we split the test network (topology as in Figure 1) into two IS-IS domains: DUT1, DUT2, DUT3, and DUT4 were configured to be part of an IS-IS level 2 process 1; DUT7 and DUT8 were configured into IS-IS level 1 process 1. DUT5 became the routing interconnection point and got configured into IS-IS level 1-2 process 1. In this scenario, we again configured a multicast MVPN between DUT1, DUT7, and DUT8. This time, the internal BIERv6 tunnels had to span the different IS-IS domains. By sending traffic through the test network as before, we verified that the inter-AS scenario was supported.

In a final step, we verified whether the MVPN to BIERv6 migration would work across multiple BGP autonomous systems (AS). For this purpose, Huawei simplified the test setup by temporarily disconnecting DUT2, DUT4, DUT6, and DUT8 from the infrastructure. This left us with four routers. Huawei configured DUT1 and DUT3 with AS 200; DUT 5 and DUT7 as AS 100. DUT3 and DUT5 were setup as autonomous system border routers (ASBR). Between these two ASBRs, an MP-EBGP peering was setup. Three routers (DUT1, DUT5, DUT7) were configured with BIERv6 and IS-ISv6; DUT3 was setup to emulate the router not supporting BIERv6 to demonstrate the migration scenario.

To notify the ingress (transmit) endpoint of the desired multicast routing via ASBR (DUT5), the command "protocol static-bift" was issued on DUT1. We verified that this setup worked: The expected BIERv6 tunnels were setup, multicast routes were learned in the MVPN instance. When multicast traffic was sent by the tester port 1 connected to DUT1, it was received on tester port 2 connected to DUT7.

Slicing Support with FlexAlgo

In mobile networks, specifically in 5G standalone solutions, traffic flows need to be routed separately following types of constraints. Some flows require low latency, some high bandwidth, some low packet loss, some economic paths. Traditionally, there exist only two types of routing algorithms in Segment Routing networks: The default shortest-path first (SPF) algorithm and a strict path alternative where the operator (or a network management solution) defines transit nodes manually.

With FlexAlgo, it is possible to auto-steer traffic via any topology and path based on operator-defined criteria. FlexAlgo enriches IGP path computation capabilities, including algorithms and constraints, to achieve traffic engineering capabilities. In this test session, Huawei demonstrated FlexAlgo support both for Segment Routing MPLS networks and for SRv6 networks. Huawei defined two different logical topologies for different services (IPv4 and IPv6) in scenario one as seen in Figure 7. The topology with red links alone refer to the FlexAlgo 128 path and with the blue links included refer to the FlexAlgo 129 path. Huawei configured the FlexAlgo 128 and FlexAlgo 129 as shown in Table 6.

In this case, two EVPN tunnels were created end-to-end between DUT1 and DUT7. FlexAlgo algorithm 128 was configured to carry red colored traffic and FlexAlgo algorithm 129 was configured for traffic colored differently across the network under test. The blue colored links of FlexAlgo algorithm 129 carry the traffic with different colors. We verified that the routing tables showed distinct routes for each of the tunnels. Figures 7 and 8 show the logical topologies from the point of view of the different FlexAlgo algorithms:

The test results show that SR and SRv6 best effort can forward traffic not only based on the minimize routing (IGP) cost. Traffic can be routed based on minimizing TE metric as well, according to Huawei—a validation of the TE metric support was included in this test case. Additionally, SR/SRv6 BE can forward traffic in specific topology with include or exclude affinity. This enables disjunctive service routing.

Parameter	FlexAlgo 128	FlexAlgo 129
EVPN	L3VPNV4	L3VPNV6
Metric type	IGP	TE-Metric
Calculation-Type	Shortest path	Shortest path
Constraint	Include all red	None

Table 6: Configuration of FlexAlgo Algorithms

	Segment Routing	Service Layer	IGP
Scenario A	SR MPLS BE (best effort)	EVPN (IPv4)	ISIS
Scenario B	SR MPLS BE	EVPN (IPv6)	ISIS
Scenario C	SRv6	EVPN (IPv4)	ISISv6
Scenario D	SRv6	EVPN (IPv6)	ISISv6

Table 7: Slicing Test Scenarios

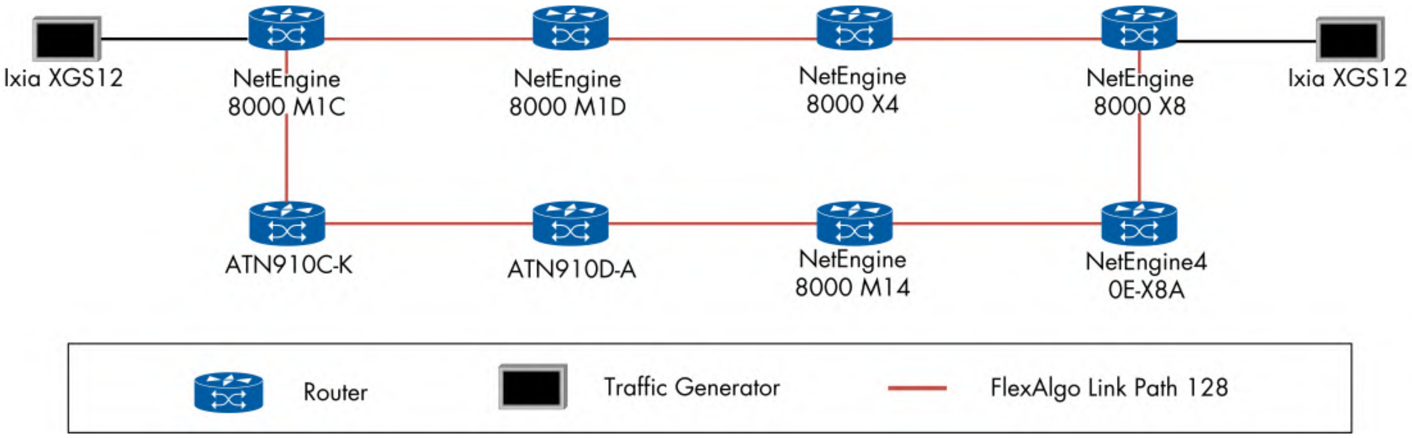

Figure 7: FlexAlgo Red Coloured Traffic Test Topology

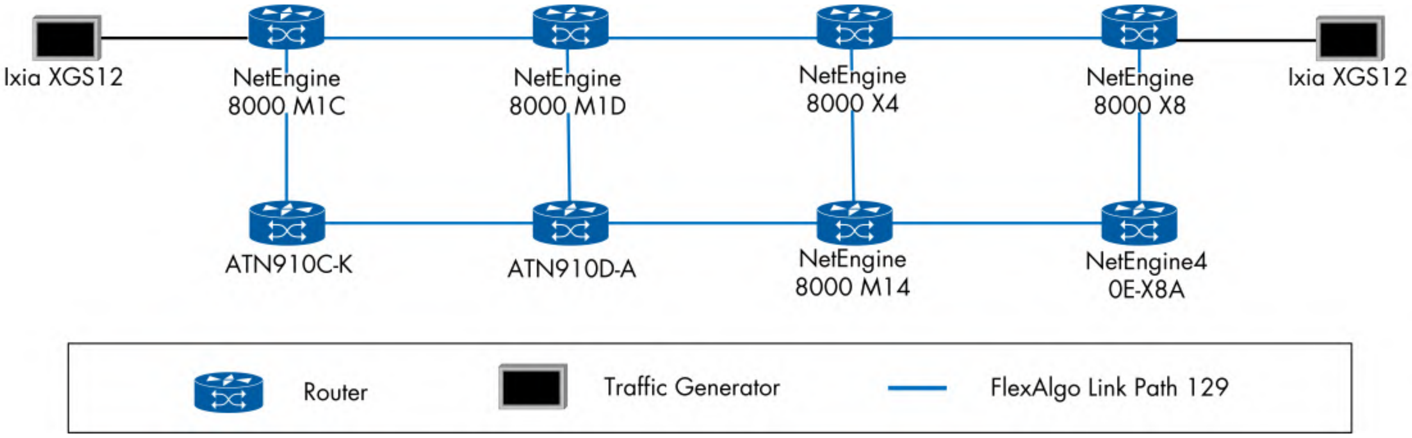

Figure 8: FlexAlgo Blue Coloured Traffic Test Topology

Tunnel Failover and Micro-Loop Avoidance

In both SR-MPLS and SRv6 FlexAlgo scenarios as described above, we tested failover with TI-LFA technology to accelerate the failover independent of the test topology, and we verified that micro-loop avoidance was configured on the devices under test.

To this end, Huawei enabled both TI-LFA and micro-loop avoidance on all DUTs. Using the same service configuration with tunnels using FlexAlgo 128 routing, we failed the link between DUT1 and DUT3 by decision of Huawei. Other links were not disconnected during this test. This test run also demonstrates that the fast rerouting backup path (FRR) can be calculated based on user-defined FlexAlgo constraints.

To measure the failover time, Huawei configured the Ixia test tool to generate traffic with a fixed, well-known packet rate of 512 bytes at 8 Gbit/s. Traffic was sent for one flow unidirectionally from DUT1 to DUT7. We measured the number of packets lost during the failover and derived the out-of-service time accordingly by dividing the number of lost packets by the packet rate. Afterwards, we reconnected the link and measured an out-of-service time during recovery (if any). Each test case was conducted three times to gauge the reproducibility of out-of-service times.

Micro loops can appear during IP network failover scenarios because each node converges its routing table independently. Nodes that are fewer hops away from a failure will naturally converge earlier. During this network-wide procedure, one node might have converged and might send traffic towards an alternate path, while another node on this alternate path would bounce this traffic towards the original path thus creating a loop. These loops typically exist for the duration of the route update timer, multiplied by the number of affected hops. They can blackhole a considerable amount of traffic and should thus be avoided. This problem of legacy Internet routing can be resolved by Segment routing's extensions for micro loop avoidance.

Scenario	Out-of-Service Time Failover	Out-of-Service Time Recovery
SR-MPLS	0.87-1.41 ms	no packets lost
SRv6	0.86-0.94 ms	no packets lost

Table 8: TI-LFA Service Failover and Recovery Times

On one hand, the failover time measured above already proves that micro loop avoidance was working. As an additional verification, Huawei configured very high values (10 seconds) for the routing update timers, which would cause micro loops for at least 10 seconds, probably much longer, during the reconvergence phase in case micro loop avoidance would not work. We did not see any such reconvergence delays in the test. Finally, we captured the control traffic between DUT1 and DUT3. The IS-IS protocol on this link included the expected extensions.

The results confirm Huawei's claims that TI-LFA is supported in a FlexAlgo scenario, and that micro loop avoidance is implemented in the NetEngine 8000 software under test.

End-to-End 400GigabitEthernet Forwarding

With the increasing demand of service providers to upgrade the aggregation networks to 400GigabitEthernet, this test scenario was defined to validate the end-to-end 400GbE near-line rate performance on Huawei NetEngine 8000 routers and NetEngine 40E-X8 routers. A subset of the test bed (Table 1 above) was used in this scenario: The NetEngine 8000 X8, NetEngine 8000 X4, NetEngine 40E-X8A, and NetEngine 8000 M14. For the test, two 400GE Ixia AresONE tester ports were connected to NetEngine 8000 X8, and NetEngine 8000 X4 respectively, as shown in Figure 11.

All links used 400GigabitEthernet. In this test, Huawei configured an SRv6 policy across the DUTs. EANTC validated the end-to-end 400G bidirectional throughput performance following RFC 2544 benchmarking methodology. During the test, IPv4 traffic of different frame sizes 256, 512, 1024, 1280, 1500, 2048, 4096, 9000 were simulated. The results described in Table 9 include an extra 40 bytes for the SRv6 header and 20 bytes for the Inter Frame Gap.

The results confirm that 400GbE near-line rate speeds are supported across the NetEngine 8000 X4, NetEngine 8000 X8, NetEngine 8000 M14, and the NetEngine 40E-X8A.

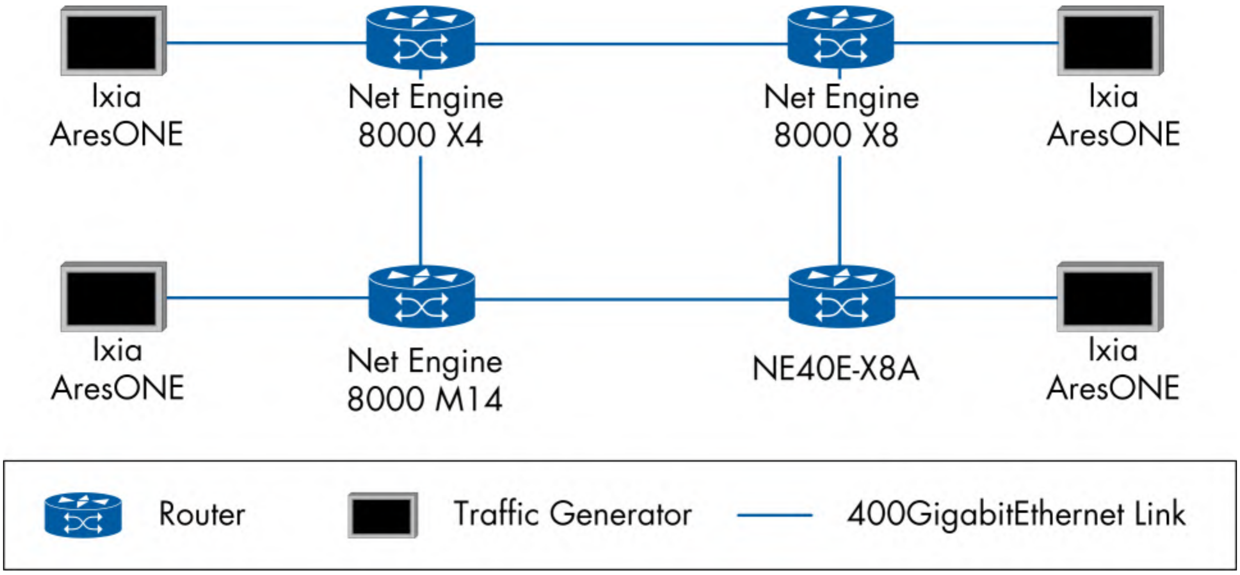

Figure 9: 400G End-to-End Speed Test Topology

Packet Size	Bi-Directional Throughput per 400GE Port	Latency (min/avg/max)	Bi-Directional Forwarding Performance per 400GE Port
256	798.84 Gbit/s	48/1351/1357 µs	316.44 Mpps
512	799.88 Gbit/s	47/1340/1353 µs	174.82 Mpps
1024	799.55 Gbit/s	47/1341/1344 µs	92.24 Mpps
1280	799.71 Gbit/s	47/1340/1346 µs	74.62 Mpps
1500	799.96 Gbit/s	47/1357/1362 µs	64.10 Mpps
2048	799.35 Gbit/s	48/1349/1362 µs	47.42 Mpps
4096	799.35 Gbit/s	48/1340/1344 µs	24.06 Mpps
9000	797.28 Gbit/s	50/1340/1346 µs	11.02 Mpps

Table 9: End-to-end 400G Throughput Performance Results

SRv6 Policy Master/Standby Protection

Service providers always aim to achieve the highest possible service availability. Typically, there are two types of network services: First, there are services that can be handled as bulk traffic and protected with general network policies and overall hop-by-hop protection schemes (such as TI-LFA in SR). Second, there are high priority services that require special proactive end-to-end protection. Huawei demonstrated the functionality of SRv6 policy-based end-to-end master/standby protection in this test campaign.

In the generic test bed (see Figure 1), an end-to-end protected SRv6 tunnel was configured between DUT1 and DUT7. First, to avoid conflicts with TI-LFA, Huawei disabled the crosslinks between DUT5/DUT6 and DUT3/DUT4, leaving exactly two disjunct end-to-end physical paths available in the test bed network. Huawei then deployed SRv6 policies with two disjunct paths on DUT1 and DUT7, respectively (for the two directions). On DUT1, for instance, the two paths DUT1 - DUT3 - DUT5 - DUT7 and DUT1 - DUT2 - DUT4 - DUT6 - DUT8 - DUT7 were configured as primary and backup path for a protected SRv6 tunnel. Path failures were discovered by BFD which was enabled for the tunnel on DUT1 and DUT7.

Next, the usual network service configurations were undertaken by Huawei: End-to-end BGPv4 service peering was setup, an L3VPN was configured over the SRv6 policy between DUT1 and DUT7. Huawei chose to configure the Ixia traffic generators to transmit 100 Byte frames with a load of 0.2 Gbit/s traffic unidirectionally, only in the direction from DUT1 to DUT7. There was no traffic sent in the other direction. In this optimized lab setup, the functional test was started by pulling out the cable between DUT1 and DUT3. BFD was configured but was not utilized in this configuration, as the direct connection to the ingress router maintaining the policy was disconnected. As a result, the failover times measured in this test do not include failure detection delays.

Across three runs, we measured failover times between 26 ms and 29 ms and no packet loss at all during recovery when the cable was plugged in again. These results confirm Huawei's claim that the failover should happen with an out-of-service time less than 50 ms in this best-case lab setup.

Network Slicing Based on Slice ID

Huawei explained that new service types may require connectivity services with advanced characteristics comparing to traditional VPNs, such as strict isolation with guaranteed bandwidth, with the introduction of 5G. Huawei supports FlexE technology as defined by the OIF Flex Ethernet Implementation Agreement (2016). FlexE provides strict isolation below the Ethernet service layer. Huawei claimed that their FlexE implementation supports up to 20 slices. 5G networks might require more than 20 slices in the future; a mechanism is required to increase the number of slices per link. Huawei's legacy method was to enable channelized sub interfaces based on VLANs, which can support isolated VI (virtual interface) resources, allowing hundreds of slices. However, this number might still be insufficient in the future.

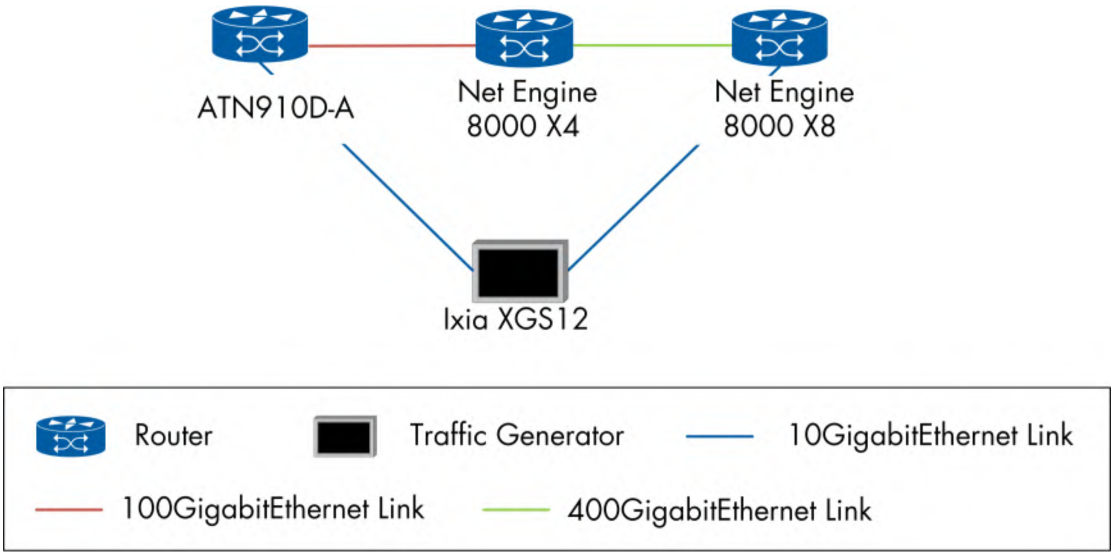

Figure 10: Slicing Based on Slice ID Test Topology

In addition, FlexE- and VLAN-based slices are not directly visible on the IP/SRv6 layer, making slice-aware routing difficult.

To solve this challenge in SRv6, Huawei has co-authored an Internet draft (draft-dong-6man-enhanced-vpn-vtn-id-02, work in progress) introducing a "SliceID" in IPv6 extension headers. "SliceID" is a Huawei term; in the IETF draft, it is referenced as Virtual Transport Network identifier (VTN ID). The SliceID identifies the physical layer (FlexE) slice and makes it available at the SRv6 layer. With this new technique, Huawei claims that the NetEngine 8000 routers can support up to 1000 slices. Besides the high capacity, the new solution also simplifies the slicing configuration. All slices in the same physical interface share the same IP address and IGP configurations. It is only needed to configure the SliceID number and reserved bandwidth for each slice. The SliceID is stored in the IPV6 Hop-by-hop header. IP routing lookups take the End.X.SID and the SliceID in the packet header into account.

Huawei demonstrated that 1,000 slices are supported by ATN-910D-A and by the NetEngine 8000 X4 and NetEngine 8000 X8 using the SRv6-based SliceID technology. For this test case, the Huawei team constructed a separate test bed using only three routers (see Figure 10) connected by 10GbE, 100GbE and 400GbE interfaces.

Once everything got connected and configured, including IS-IS, SRv6-TE policies, and L3VPNs, we transmitted a 10 Gbit/s bidirectional traffic across the three DUTs to confirm the functionality. Huawei's claim was confirmed that 1,000 slices are supported across the NetEngine 8000 X4, the NetEngine 8000 X8 and ATN910D-A with their 10GbE, 100GbE and 400GbE interfaces.

Test Results - Broadband Network Gateway (BNG)

Broadband network gateways are essential for consumer broadband Internet services and can be part of any fixed or converged network topology. With the continued increase in IPTV, gaming, and VR/AR services, the number of consumer Internet endpoints grows further. BNGs provide the support to control large-scale subscriber infrastructures in the service provider network. Specifically, some new services are more sensitive to latency.

The NetEngine 8000 family of routers supports BNG functions as well, some of which Huawei demonstrated to EANTC in this test session. NetEngine 8000 F1A and M14 are compact small routers; Huawei said that they are often used in the aggregation layer in metro networks. Huawei claimed that these F1A and M14 routers can support large-scale BNG subscribers to support the service downward and improving user experience.

In this group of tests, Huawei provided two single-router test scenarios with the F1A and M14 router, respectively (see Figure 11). After the subscriber management configuration over the DUT interface, a User Network Route (UNR) was generated in the routing table. All the subscriber traffic was egressed via the UNR. Huawei stated to support a maximum of 64,000 active subscriber sessions in total per protocol type and per port. Due to limitations of the Ixia XGS12 emulator used, we were able to verify the total number of 64,000 subscribers per protocol type but only a maximum of 32,000 subscribers per port and protocol type.

In the test bed, four 10GE ports of the NetEngine 8000 M14 and two 10GE ports of the NetEngine 8000 F1A were directly connected to the Ixia traffic simulator. In the test, 64,000 IPoE and 64,000 PPPoE subscriber sessions were verified on the NetEngine 8000 M14 with two bidirectional traffic streams, one with IPoE traffic and one with PPPoE traffic. For DUT NetEngine 8000 F1A, 32,000 IPoE and 32,000 PPPoE subscriber sessions were verified with IPoE and PPPoE bidirectional traffic streams. During the test, Ixia XGS12 was used to simulate the subscriber sessions and we began the test by verifying the single IPoE subscriber. The test was repeated for different Local Area Network (LAN) groups as mentioned in Table 10 and 11. Also, for the tests, all the requested subscriber sessions were successfully established and zero traffic loss was observed.

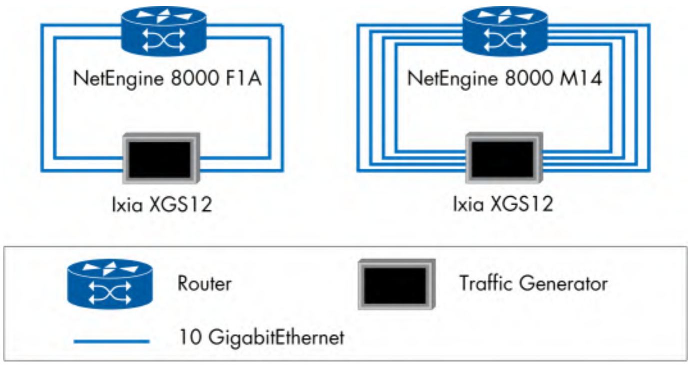

Figure 11: BNG Test Topology

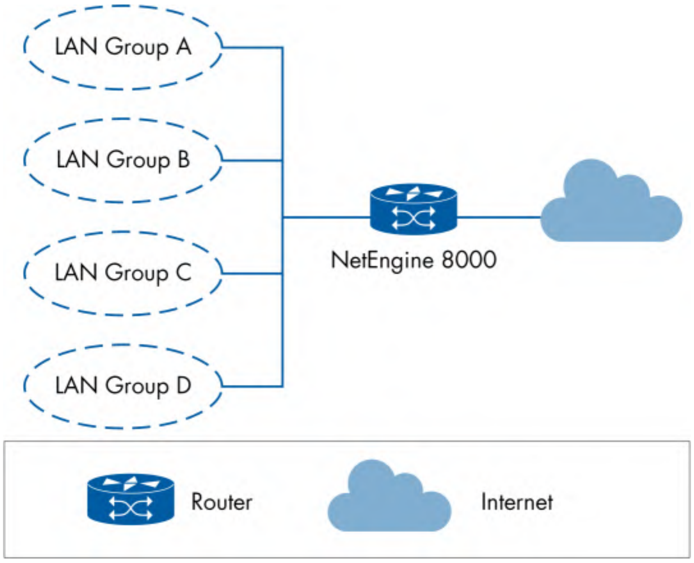

Figure 12: BNG Logical Test Topology

Group Name	IPoE Subscribers	PPPoE Subscribers	IPoE Throughput	PPPoE Throughput	Latency (min/avg/max)
Group A	1	0	18.2 Gbit/s	0	9/10/12 µs
Group B	0	1	0	17.7 Gbit/s	10/10/12 µs
Group C	64,000	64,000	18.2 Gbit/s	17.7 Gbit/s	11/13/29 µs

Table 10: BNG Scalability Results on NetEngine 8000 M14

Group Name	IPoE Subscribers	PPPoE Subscribers	IPoE Throughput	PPPoE Throughput	IPOE & PPPoE Throughput	Latency (min/avg/max)
Group A	1	0	18.2 Gbit/s	0	N/A	7/7/12 µs
Group B	0	1	0	17.7 Gbit/s	N/A	7/8/41 µs
Group C	32,000	32,000	N/A	N/A	18.5 Gbit/s	8/10/12 µs

Table 11: BNG Scalability Results on NetEngine 8000 F1A

Test Results - Intelligent Operations & Maintenance

As part of our series of EANTC reports on Huawei core and aggregation routing solutions, we have investigated multiple aspects of the network provisioning and management solution, the Huawei iMaster NCE-IP. This solution has developed into a powerful provisioning and operations tool. Previously, we had validated the segment routing tunnel and EVPN service provisioning, the network optimization support, and the network monitoring functions in our test published in March 2020.

Huawei explained that NCE-IP is committed to enabling network autonomous driving and working with TMF to innovate the future autonomous driving network architecture. In this test, Huawei demonstrated closed-loop automation to EANTC. The tests were conducted remotely under Covid-19 conditions as well.

The NCE team used the same master test bed as shown in Figure 1. The line cards and connections remained the same as before, only the software configurations were changed to meet the goals of the respective evaluations. All routers were managed by NCE. With all of this having been configured, we went through a number of test scenarios.

Note that the evaluations are functional demonstrations, despite the fact that the topic was performance monitoring-related. Huawei always setup only one EVPN service and frequently changed the configuration of the maximum bandwidth reservation on NCE, to force specific policy decisions.

This method is very appropriate to validate the functional correctness of the NCE's decision algorithms for a single service. We cannot extrapolate the results for service provider networks with more than one EVPN service. We recommend to evaluate the NCE's ability to manage a realistically scaled network with multiple services and a much higher load of network monitoring data as part of any service provider proof of concept test.

Closed-Loop for On-Demand Traffic (SRv6 Policy Path Navigation)

This test was set up to validate Huawei's claims of closed-loop for on-demand traffic in NCE. "Closed-loop" describes a solution that continually monitors the network situation, analyzes it, and makes decisions based on the situation. In the context of SRv6 and on-demand traffic, a closed-loop solution requires on-demand traffic optimization to avoid network congestion, which was the main topic of our evaluation. In this section we simulated service traffic that increased periodically, NCE was able to detect traffic threshold-crossing and calculated other available paths on the entire network. NCE also implemented network optimization by doing load balancing and automatic traffic splitting to avoid traffic congestion.

Huawei configured IGP link performance monitoring with a 5-minute interval on all links. Then, using the NCE, Huawei created a new SRv6 policy named Policy 1 between DUT1 and DUT7 and set UCMP to the policy with a minimum path to 2 as shown in Figure 13.

To show that even non-SRv6 intermediate routers would not break this scenario, Huawei set DUT3 as a device that would not support SRv6 and created a Virtual IGP link on NCE manually.

NCE computed two SRv6 end-to-end paths as shown in Figures 14 and 15. Huawei then added a new IP-EVPN service between DUT1 and DUT7; newly created sub-interfaces were connected to the Ixia tester on both DUTs, and were setup in this EVPN service as access nodes. 5 Gbit/s bidirectional traffic was sent from the Ixia traffic generator with maximum bandwidth reservation set to 100% on NCE. We monitored the links and noticed that the traffic was distributed equally between the two Equal Cost Multi Paths (ECMP) computed by NCE as expected in Figure 16.

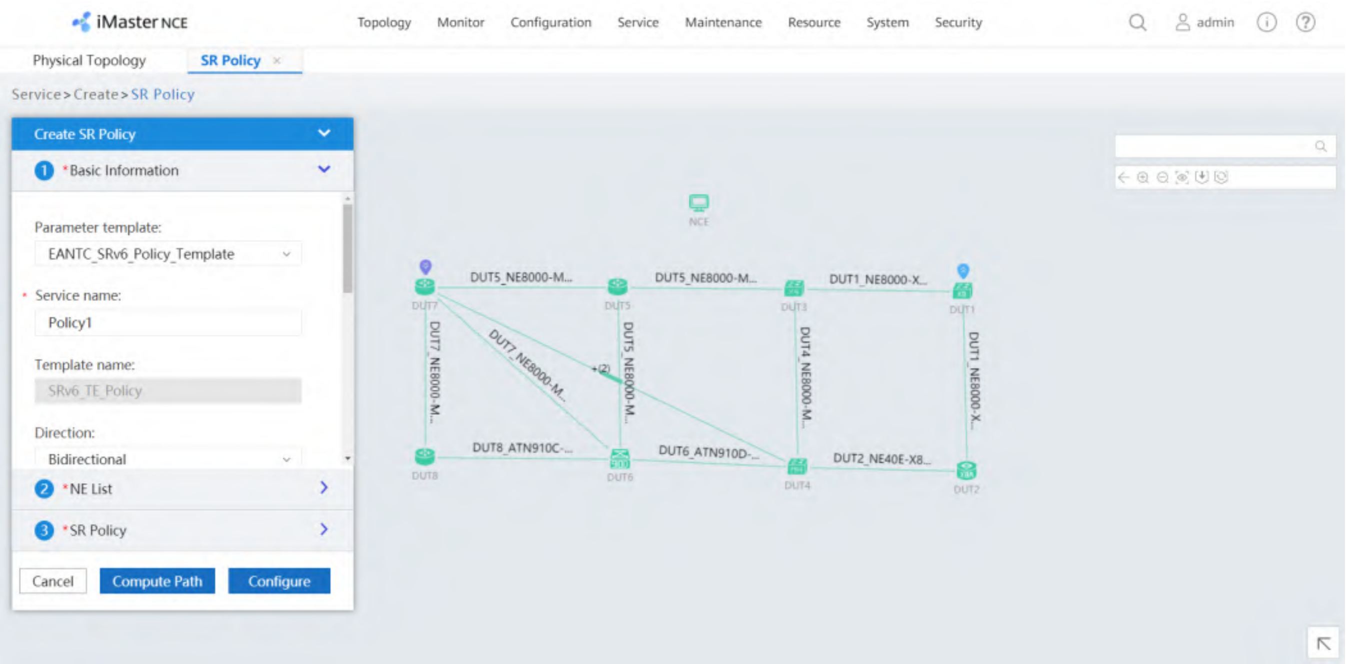

Parameter template: EANTC_SRv6_Policy_Template

* Service name: Policy1

Template name: SRv6_TE_Policy

Direction: Bidirectional

Figure 13: SRv6 Policy Creation from NCE

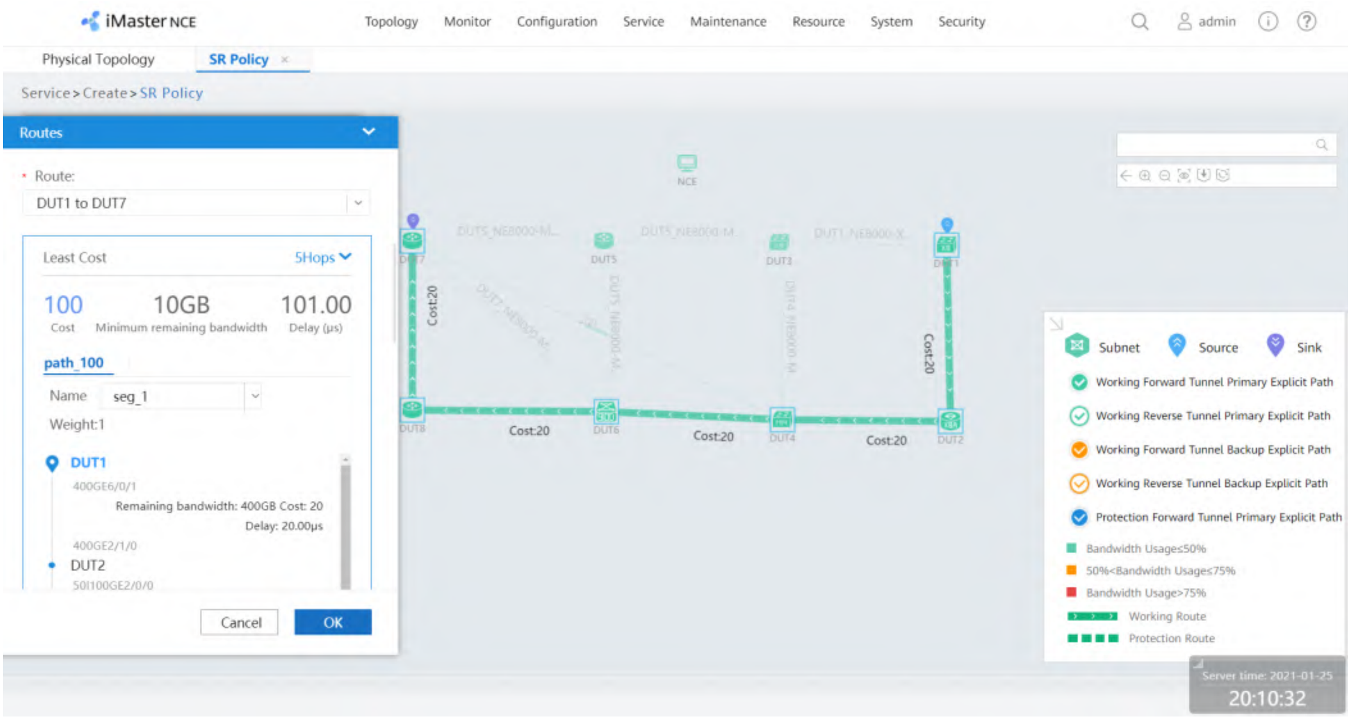

* Route: DUT1 to DUT7

Least Cost 5Hops

100	10GB	101.00
Cost	Minimum remaining bandwidth	Delay (µs)

path_100

Name seg_1

Weight:1

DUT1
400GE6/0/1
Remaining bandwidth: 400GB Cost: 20
Delay: 20.00µs
400GE2/1/0
DUT2
50|100GE2/0/0

Server time: 2021-01-25
20:10:32

Figure 14: SRv6 Policy Path 1 with Weight 1

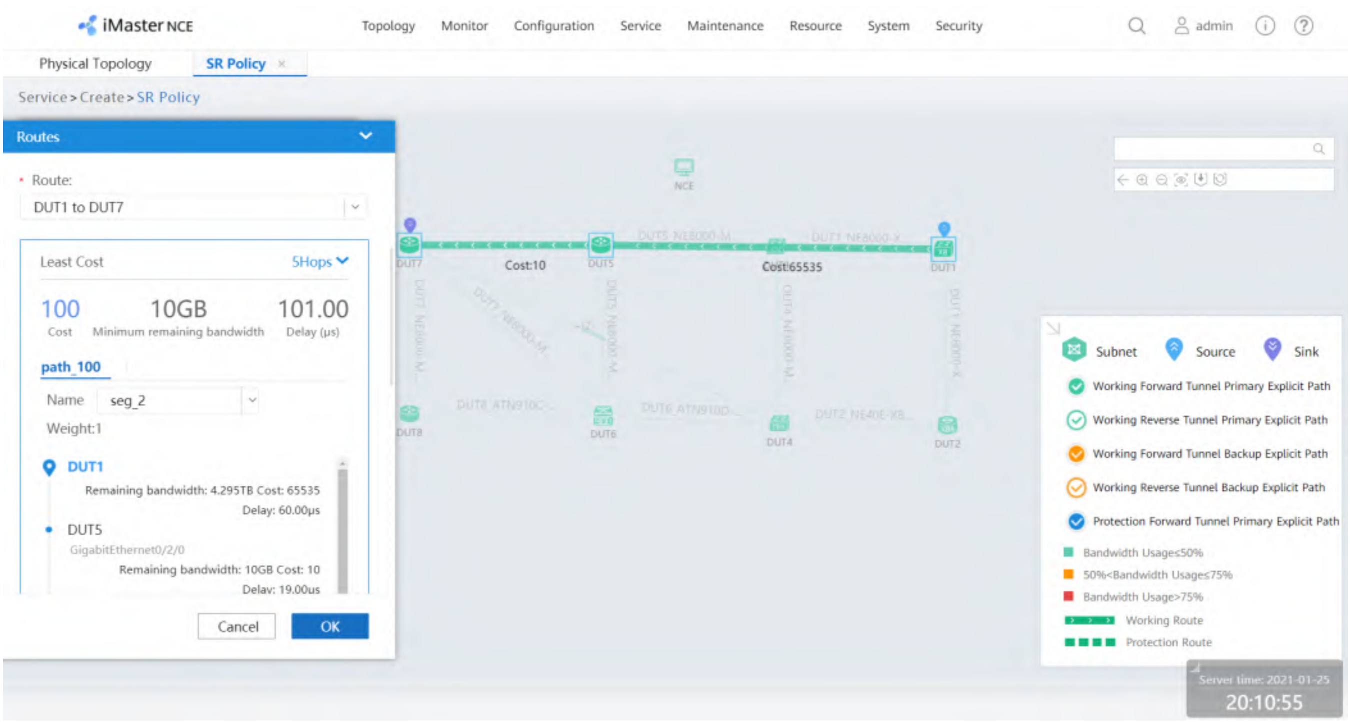

Route: DUT1 to DUT7

Least Cost 5Hops

100	10GB	101.00
Cost	Minimum remaining bandwidth	Delay (μs)

path_100

Name seg_2

Weight:1

DUT1
Remaining bandwidth: 4.295TB Cost: 65535
Delay: 60.00μs

DUT5
GigabitEthernet0/2/0
Remaining bandwidth: 10GB Cost: 10
Delay: 19.00μs

Server time: 2021-01-25 20:10:55

Figure 15: SRv6 Policy Path 2 with Weight 1

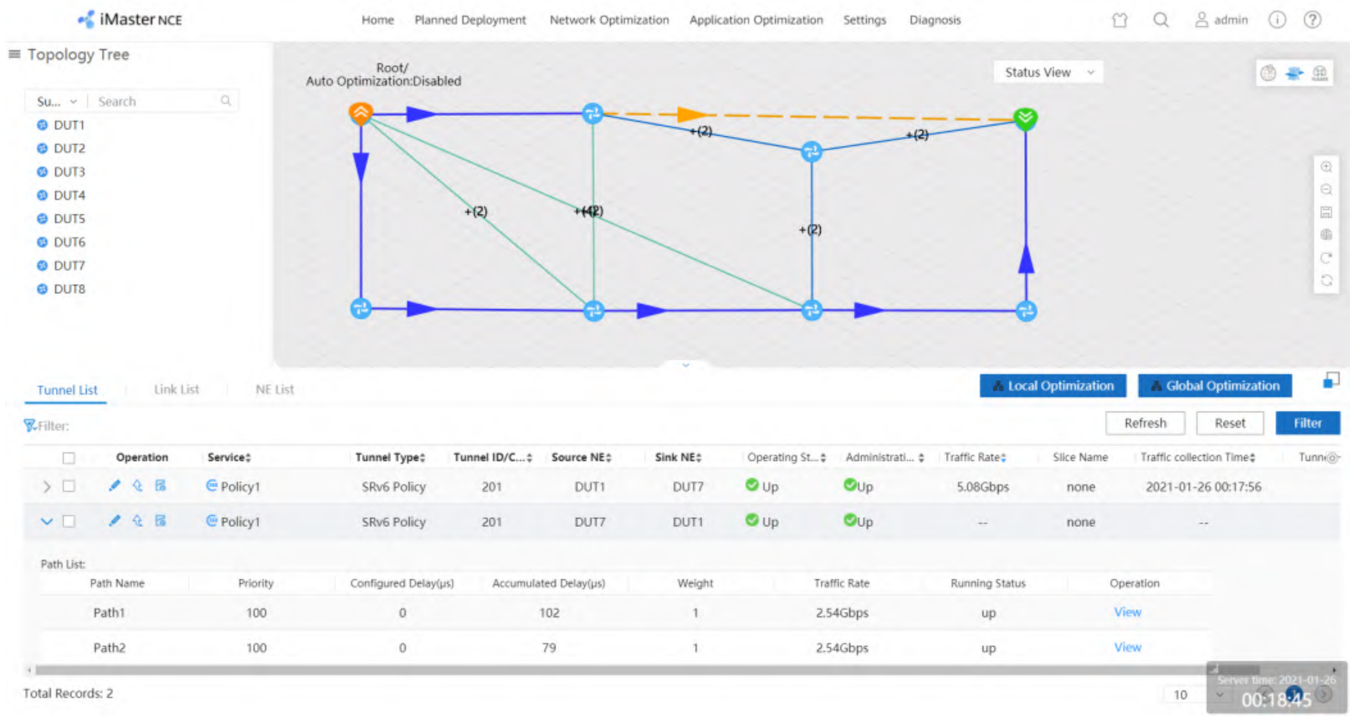

	Operation	Service	Tunnel Type	Tunnel ID/C...	Source NE	Sink NE	Operating St...	Administrati...	Traffic Rate	Slice Name	Traffic collection Time	Tunn
☐		Policy1	SRv6 Policy	201	DUT1	DUT7	Up	Up	5.08Gbps	none	2021-01-26 00:17:56	
☐		Policy1	SRv6 Policy	201	DUT7	DUT1	Up	Up	--	none	--	

Path List:

Path Name	Priority	Configured Delay(μs)	Accumulated Delay(μs)	Weight	Traffic Rate	Running Status	Operation
Path1	100	0	102	1	2.54Gbps	up	View
Path2	100	0	79	1	2.54Gbps	up	View

Total Records: 2

Server time: 2021-01-26 00:18:45

Figure 16: Traffic Load Shared Equally Between Two Computed SRv6 Policies

In the next step, the test was repeated with different values for Max Bandwidth Reservation. With Max Bandwidth Reservation was set to 30% as shown in Figure 17, the NCE computed a 3rd path for Policy 1, and the traffic was allocated to the 3 paths based on their cost as seen in the Figure 18.

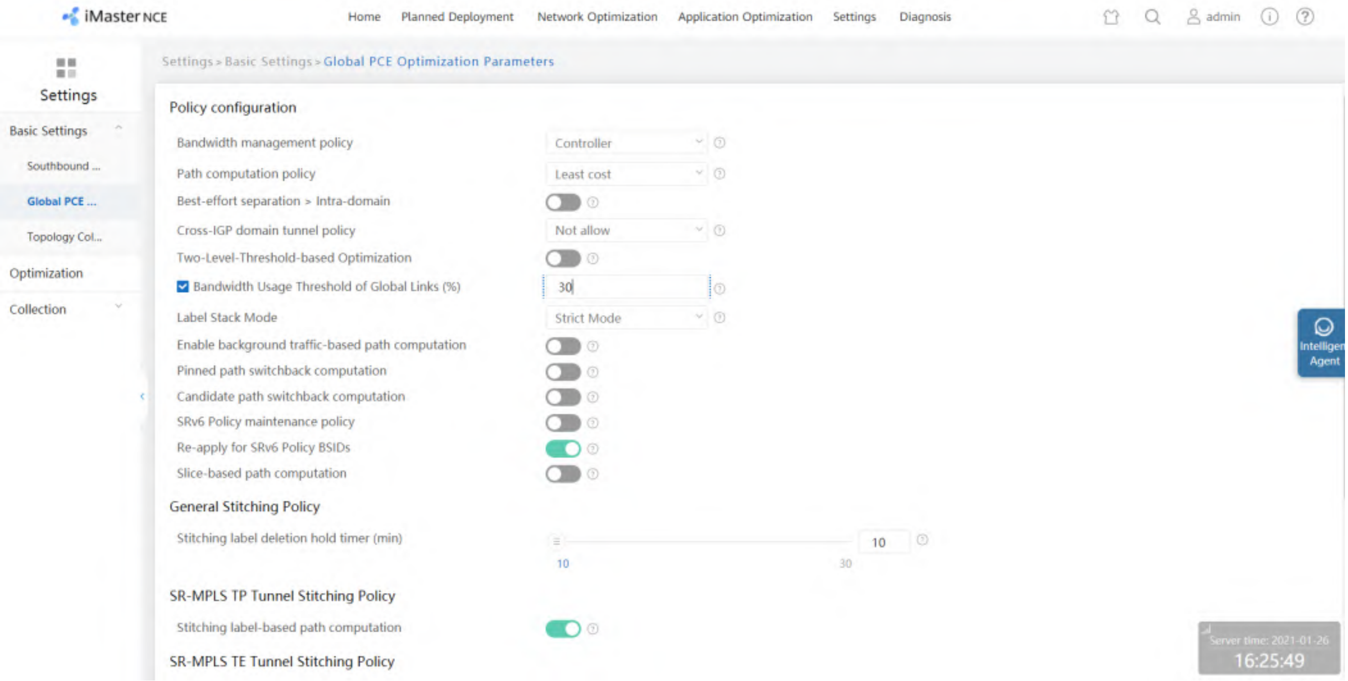

Settings > Basic Settings > Global PCE Optimization Parameters

Policy configuration

Bandwidth management policy: Controller
Path computation policy: Least cost
Best-effort separation > Intra-domain: ○
Cross-IGP domain tunnel policy: Not allow
Two-Level-Threshold-based Optimization: ○
☑ Bandwidth Usage Threshold of Global Links (%): 30
Label Stack Mode: Strict Mode
Enable background traffic-based path computation: ○
Pinned path switchback computation: ○
Candidate path switchback computation: ○
SRv6 Policy maintenance policy: ○
Re-apply for SRv6 Policy BSIDs: ●
Slice-based path computation: ○

General Stitching Policy

Stitching label deletion hold timer (min): 10
10 30

SR-MPLS TP Tunnel Stitching Policy

Stitching label-based path computation: ●

SR-MPLS TE Tunnel Stitching Policy

Figure 17: Bandwidth Reservation Configuration

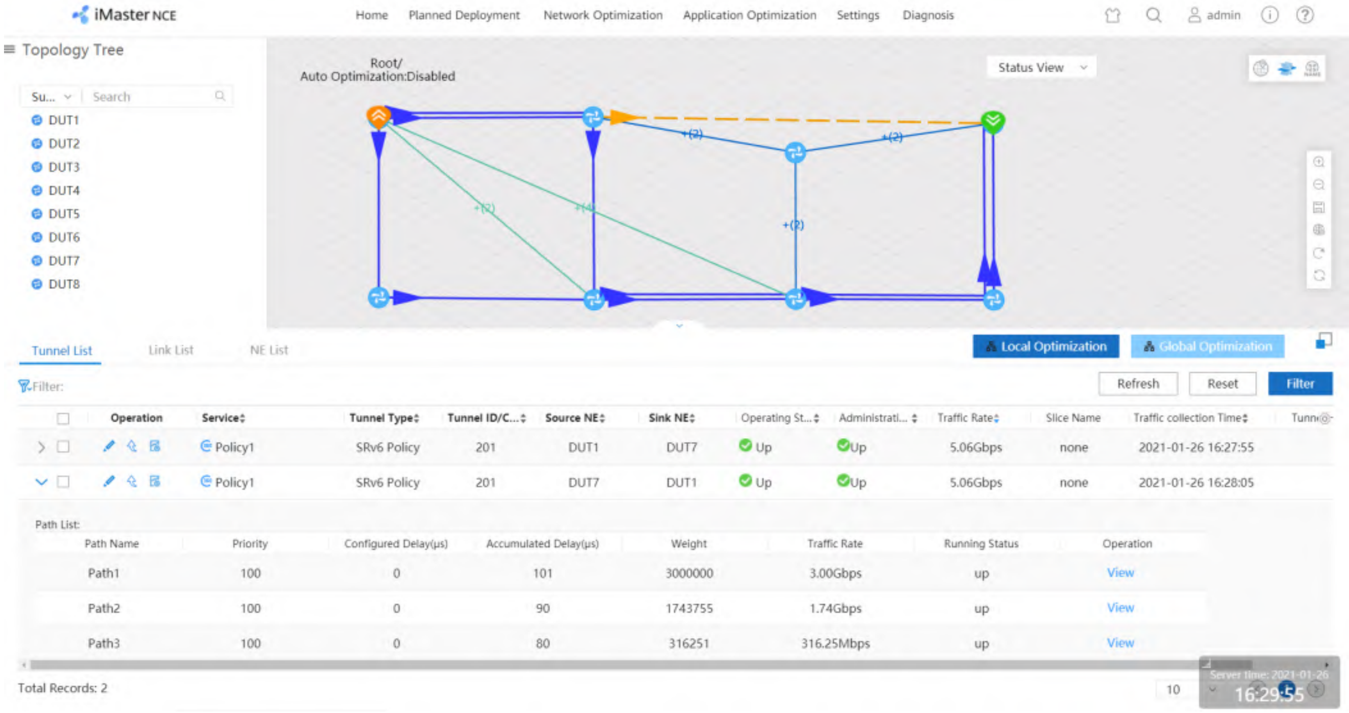

	Operation	Service	Tunnel Type	Tunnel ID/C...	Source NE	Sink NE	Operating St...	Administrati...	Traffic Rate	Slice Name	Traffic collection Time
☐		Policy1	SRv6 Policy	201	DUT1	DUT7	Up	Up	5.06Gbps	none	2021-01-26 16:27:55
☐		Policy1	SRv6 Policy	201	DUT7	DUT1	Up	Up	5.06Gbps	none	2021-01-26 16:28:05

Path List:

Path Name	Priority	Configured Delay(μs)	Accumulated Delay(μs)	Weight	Traffic Rate	Running Status	Operation
Path1	100	0	101	3000000	3.00Gbps	up	View
Path2	100	0	90	1743755	1.74Gbps	up	View
Path3	100	0	80	316251	316.25Mbps	up	View

Total Records: 2

Figure 18: Traffic Load Shared on Three Paths

To check the NCE algorithms, we further reduced the Max Bandwidth Reservation to 20%. In this scenario, NCE computed a fourth and fifth path for Policy 1, and the traffic was allocated to the five paths based on their cost as seen in Figure 19. Finally, when we set Max Bandwidth Reservation to 15%, NCE reported a path computation failure as shown in Figure 20, because there were no more available paths for the policy.

With Maximum Bandwidth Reservation reset to 100%, the NCE computed the initial path 1 and 2 for the policy and traffic was shared between the two Unequal Cost Multi Path (UCMP) paths based on the weights as shown in Figure 21.

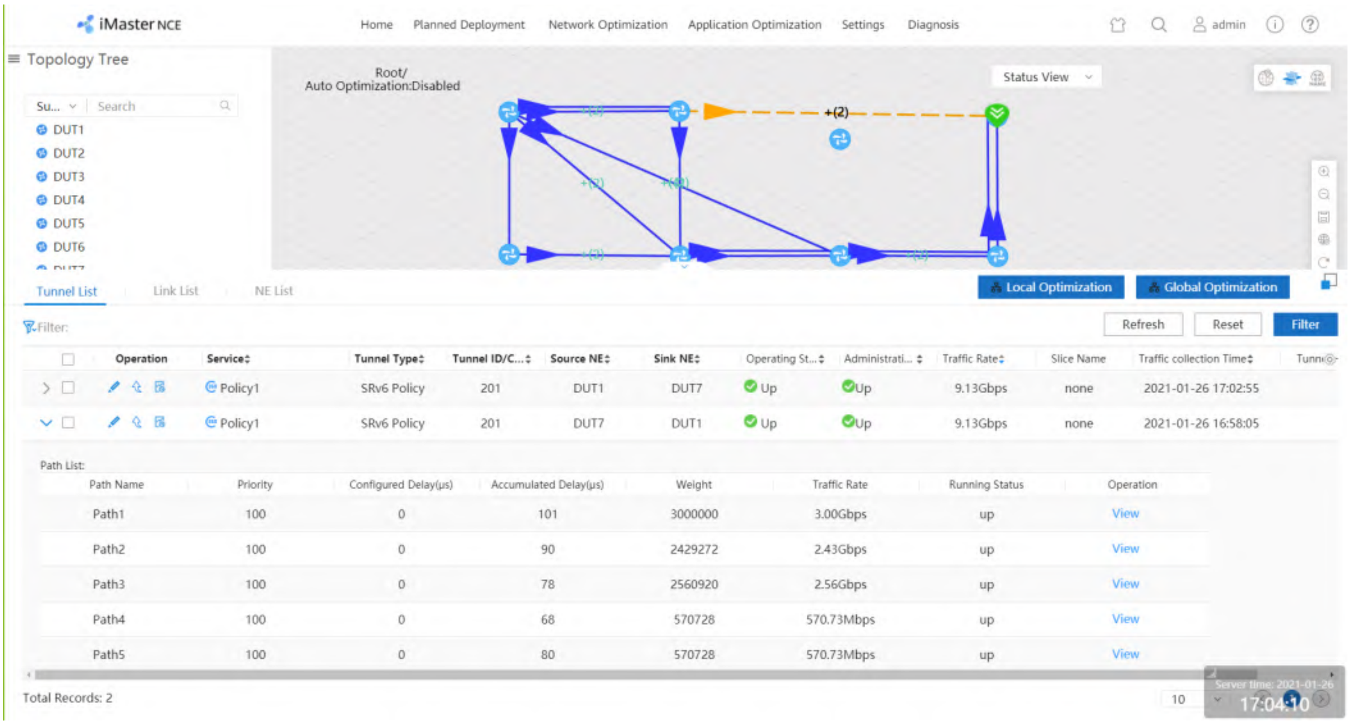

☐	Operation	Service	Tunnel Type	Tunnel ID/C...	Source NE	Sink NE	Operating St...	Administrati...	Traffic Rate	Slice Name	Traffic collection Time
☐		Policy1	SRv6 Policy	201	DUT1	DUT7	Up	Up	9.13Gbps	none	2021-01-26 17:02:55
☐		Policy1	SRv6 Policy	201	DUT7	DUT1	Up	Up	9.13Gbps	none	2021-01-26 16:58:05

Path List:

Path Name	Priority	Configured Delay(μs)	Accumulated Delay(μs)	Weight	Traffic Rate	Running Status	Operation
Path1	100	0	101	3000000	3.00Gbps	up	View
Path2	100	0	90	2429272	2.43Gbps	up	View
Path3	100	0	78	2560920	2.56Gbps	up	View
Path4	100	0	68	570728	570.73Mbps	up	View
Path5	100	0	80	570728	570.73Mbps	up	View

Total Records: 2

Figure 19: Traffic Load Shared on Five Paths

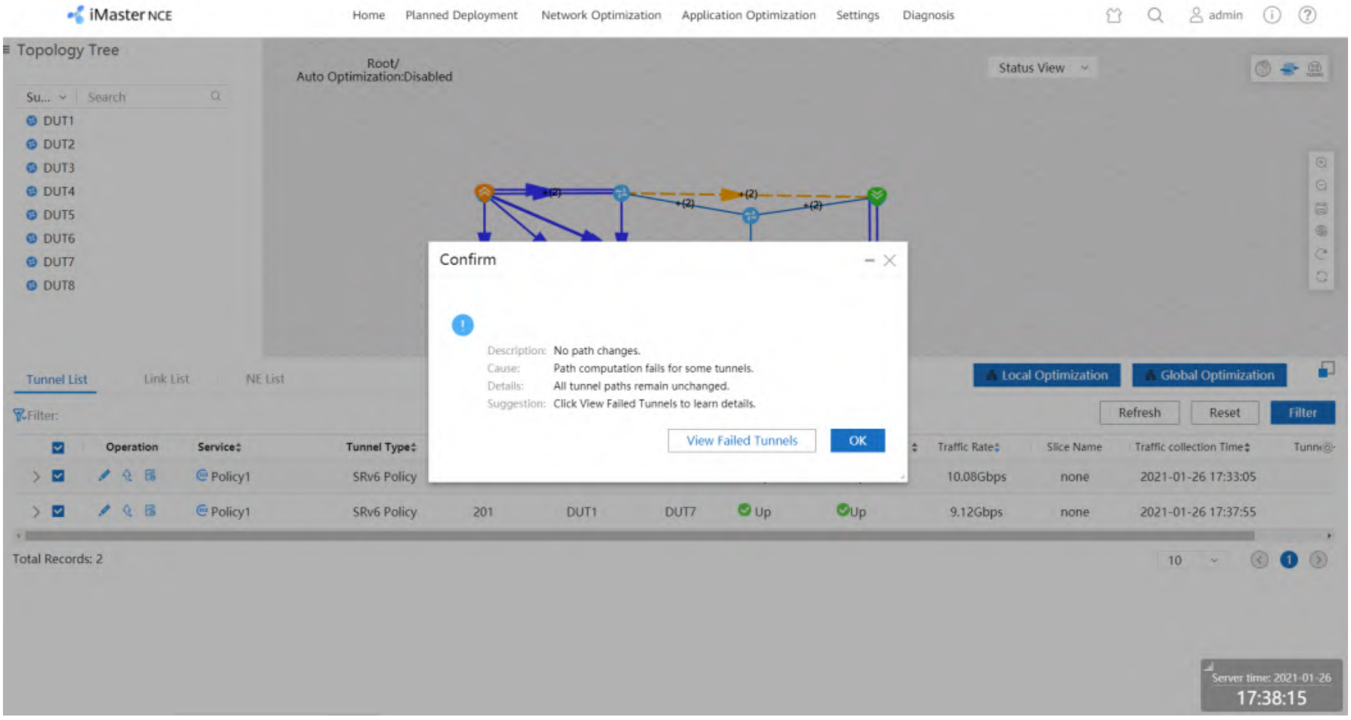

Confirm

Description: No path changes.
Cause: Path computation fails for some tunnels.
Details: All tunnel paths remain unchanged.
Suggestion: Click View Failed Tunnels to learn details.

View Failed Tunnels | OK

☑	Operation	Service	Tunnel Type						Traffic Rate	Slice Name	Traffic collection Time
☑		Policy1	SRv6 Policy						10.08Gbps	none	2021-01-26 17:33:05
☑		Policy1	SRv6 Policy	201	DUT1	DUT7	Up	Up	9.12Gbps	none	2021-01-26 17:37:55

Total Records: 2

Figure 20: Path Computation Failure

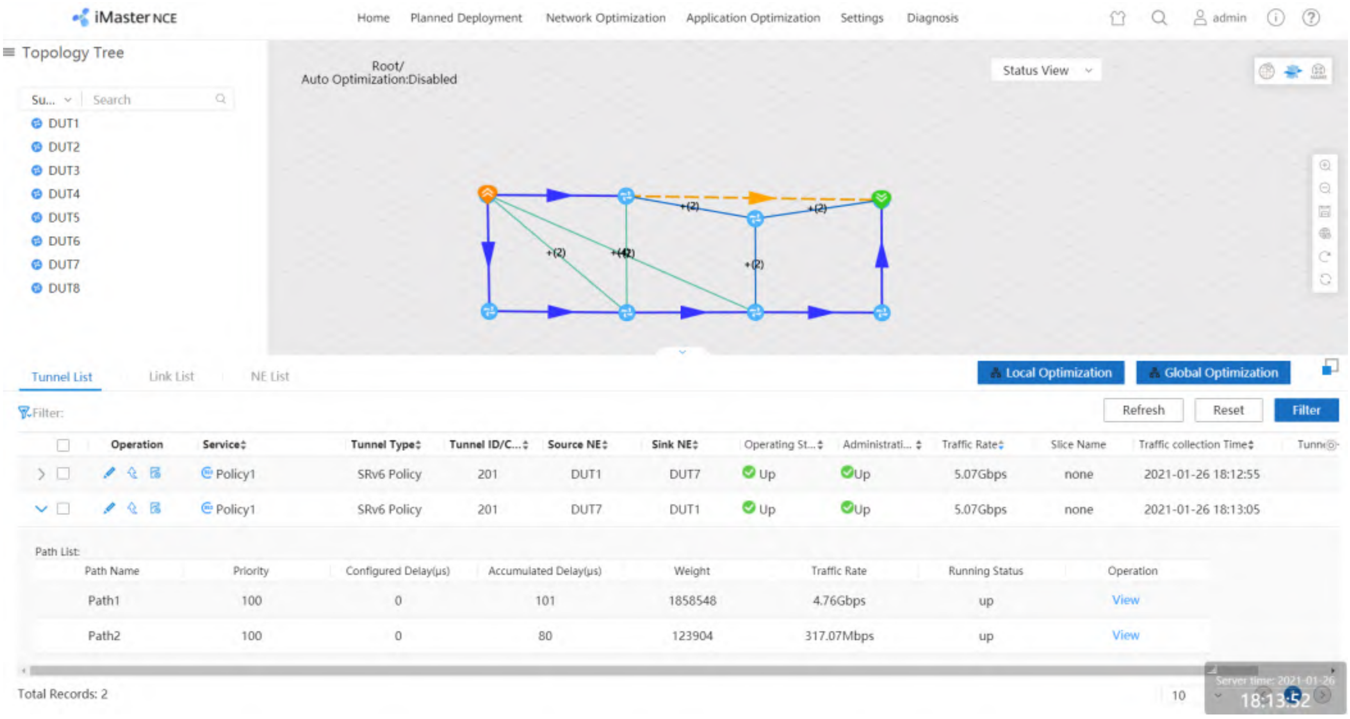

Root/
Auto Optimization:Disabled

☐	Operation	Service	Tunnel Type	Tunnel ID/C...	Source NE	Sink NE	Operating St...	Administrati...	Traffic Rate	Slice Name	Traffic collection Time	Tunn
☐		Policy1	SRv6 Policy	201	DUT1	DUT7	✓ Up	✓Up	5.07Gbps	none	2021-01-26 18:12:55	
☐		Policy1	SRv6 Policy	201	DUT7	DUT1	✓ Up	✓Up	5.07Gbps	none	2021-01-26 18:13:05	

Path List:

Path Name	Priority	Configured Delay(μs)	Accumulated Delay(μs)	Weight	Traffic Rate	Running Status	Operation
Path1	100	0	101	1858548	4.76Gbps	up	View
Path2	100	0	80	123904	317.07Mbps	up	View

Total Records: 2

Server time: 2021-01-26
18:13:52

Figure 21: Traffic Shared Between Two Paths

Closed-Loop For Poor SLA Based on TMF ADN Architecture

In the context of L3EVPN Service over SRv6 policy, a closed-loop solution requires SLA-based optimization to avoid service poor quality, which was the main focus in the test evaluation. Huawei NCE-IP supports the iFit feature which enables service providers to monitor the topology hop-by-hop.

In this case, an end-to-end L3EVPN Service over SRv6 policy was deployed between DUTs. The service was monitored by iFit, and a hop-by-hop monitor view was used to find the quality deterioration position in Network. For the test, traffic shaping was set on multiple DUTs to observe the traffic loss on different points of the topology from the iFit.

During the test, port shaping was set to 2Gbit/s on DUT6's 10GE0/2/0, which would cause traffic deterioration as seen in Figure 22. In the test, from the NCE select the path optimization, NCE will optimize new paths for the policy. The new paths bypassed the links which have deteriorated as shown in Figure 23.

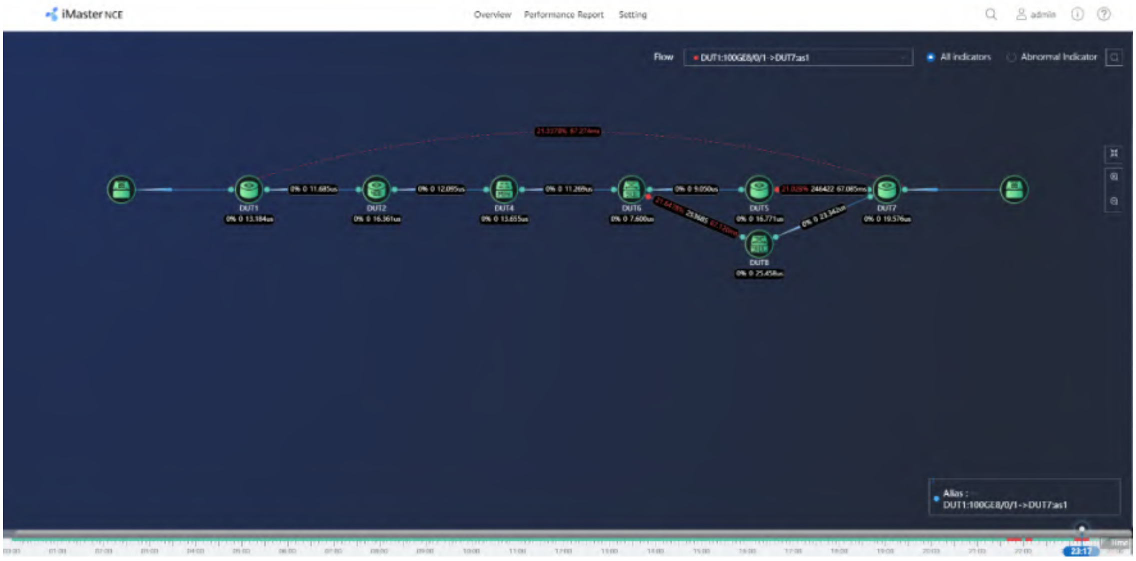

Figure 22: iFit Displaying Traffic Loss on DUT5

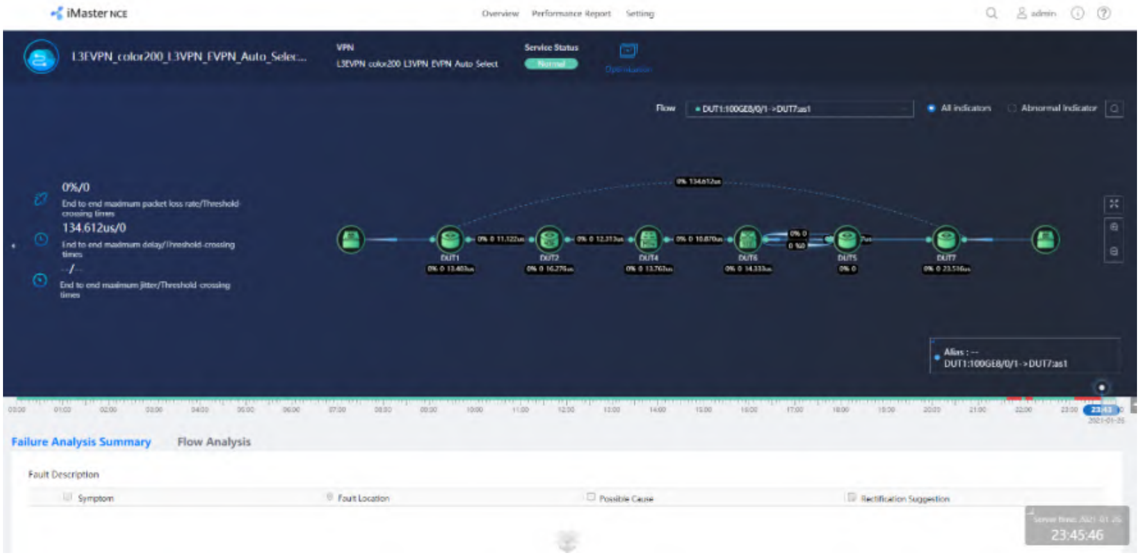

Figure 23: iFit Displaying No Traffic Loss After Path Optimization

Closed-Loop For DevOps Based on TMF ADN Architecture

NCE-IP provides AOC model to support closed-loop of DevOps. It is an open programmable platform, helping customer to customized their personalized service model and accelerating service rollout and adjustment. Huawei supports the DevOps using the Specific Network Element Driver (SND) packages and the Specific Service Plug-in (SSP), where SND provides a data model for the interaction between the OPS and NE, and SSP provides user-define service model, could be any standard of YANG model. SSP can make conversion of YANG model between service YANG and device YANG.

During the test, on the NCE DevOps view, SND packages for all DUTs were uploaded. All DUTs were added to NCE by NETCONF and data synchronization was performed on NCE. For the test, L3EVPN service templates were used to configure DUT1 as a VPN Hub node and DUT5 to DUT8 as VPN spoke nodes. In the test, specific VPN access IP address to interfaces were configured. Later in the test, the VPN service, and port shaping on the DUT's was modified using the service templates. Traffic was verified on the DUTs with the configured shaping profile.

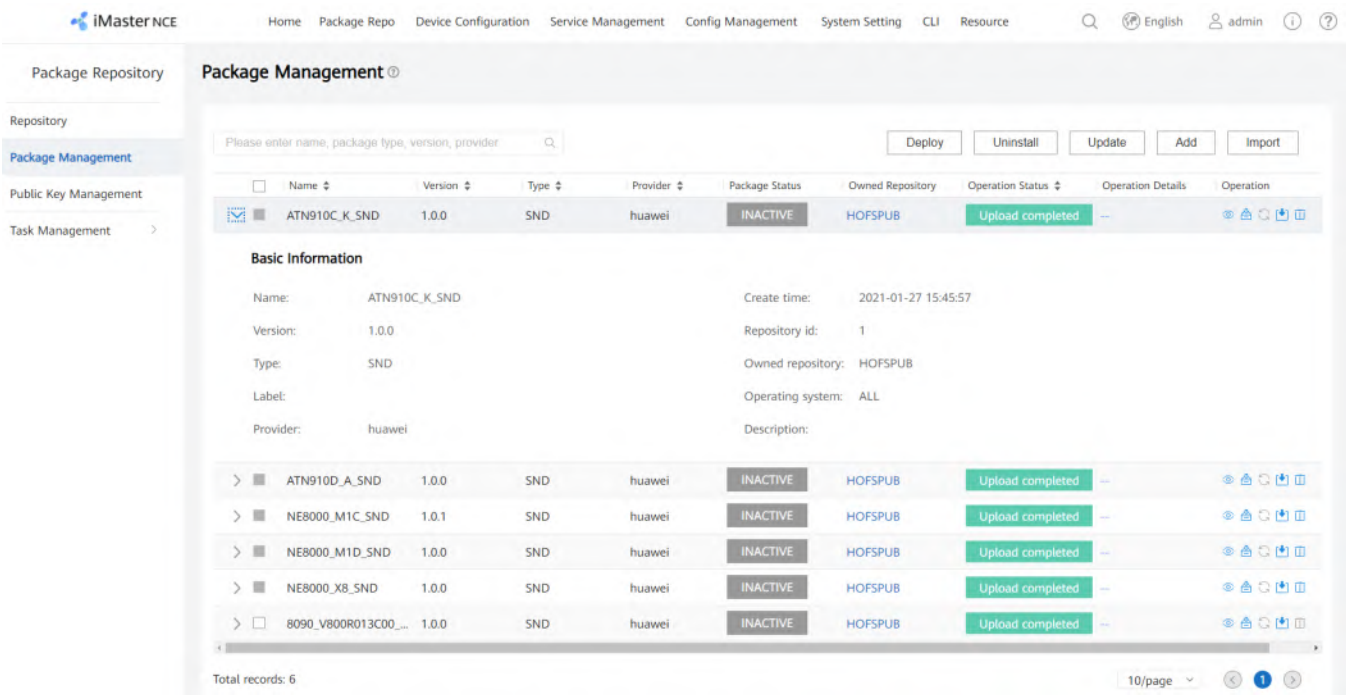

Package Management

☐	Name	Version	Type	Provider	Package Status	Owned Repository	Operation Status	Operation Details
☑	ATN910C_K_SND	1.0.0	SND	huawei	INACTIVE	HOFSPUB	Upload completed	--

Basic Information

Name: ATN910C_K_SND
Version: 1.0.0
Type: SND
Label:
Provider: huawei
Create time: 2021-01-27 15:45:57
Repository id: 1
Owned repository: HOFSPUB
Operating system: ALL
Description:

☐	Name	Version	Type	Provider	Package Status	Owned Repository	Operation Status	Operation Details
■	ATN910D_A_SND	1.0.0	SND	huawei	INACTIVE	HOFSPUB	Upload completed	--
■	NE8000_M1C_SND	1.0.1	SND	huawei	INACTIVE	HOFSPUB	Upload completed	--
■	NE8000_M1D_SND	1.0.0	SND	huawei	INACTIVE	HOFSPUB	Upload completed	--
■	NE8000_X8_SND	1.0.0	SND	huawei	INACTIVE	HOFSPUB	Upload completed	--
☐	8090_V800R013C00_...	1.0.0	SND	huawei	INACTIVE	HOFSPUB	Upload completed	--

Total records: 6

Figure 24: DUT Package Management

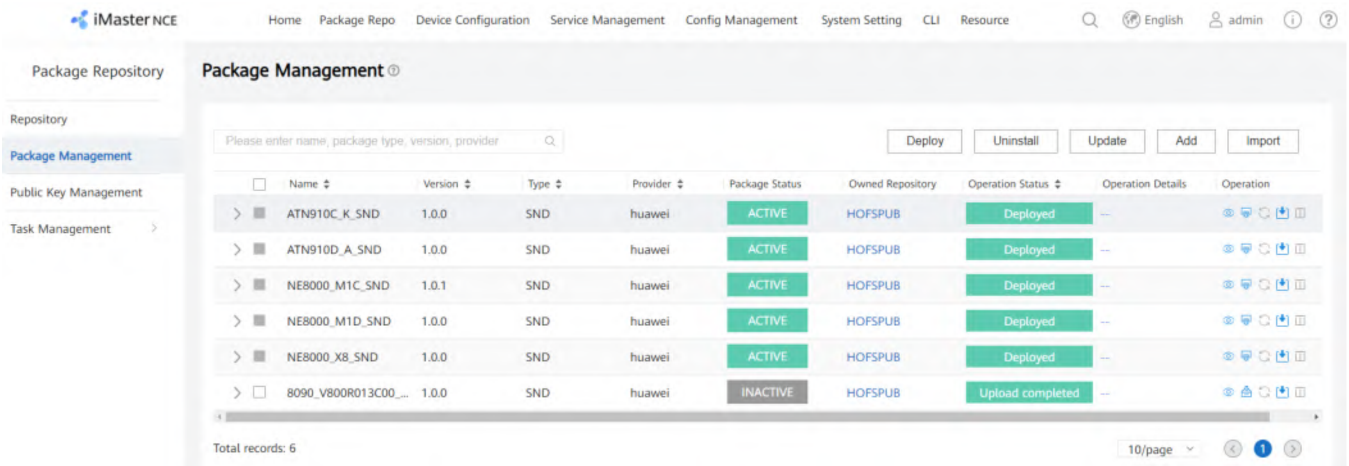

Package Management

☐	Name	Version	Type	Provider	Package Status	Owned Repository	Operation Status	Operation Details
■	ATN910C_K_SND	1.0.0	SND	huawei	ACTIVE	HOFSPUB	Deployed	--
■	ATN910D_A_SND	1.0.0	SND	huawei	ACTIVE	HOFSPUB	Deployed	--
■	NE8000_M1C_SND	1.0.1	SND	huawei	ACTIVE	HOFSPUB	Deployed	--
■	NE8000_M1D_SND	1.0.0	SND	huawei	ACTIVE	HOFSPUB	Deployed	--
■	NE8000_X8_SND	1.0.0	SND	huawei	ACTIVE	HOFSPUB	Deployed	--
☐	8090_V800R013C00_...	1.0.0	SND	huawei	INACTIVE	HOFSPUB	Upload completed	--

Total records: 6

Figure 25: DUT Package Upload Status

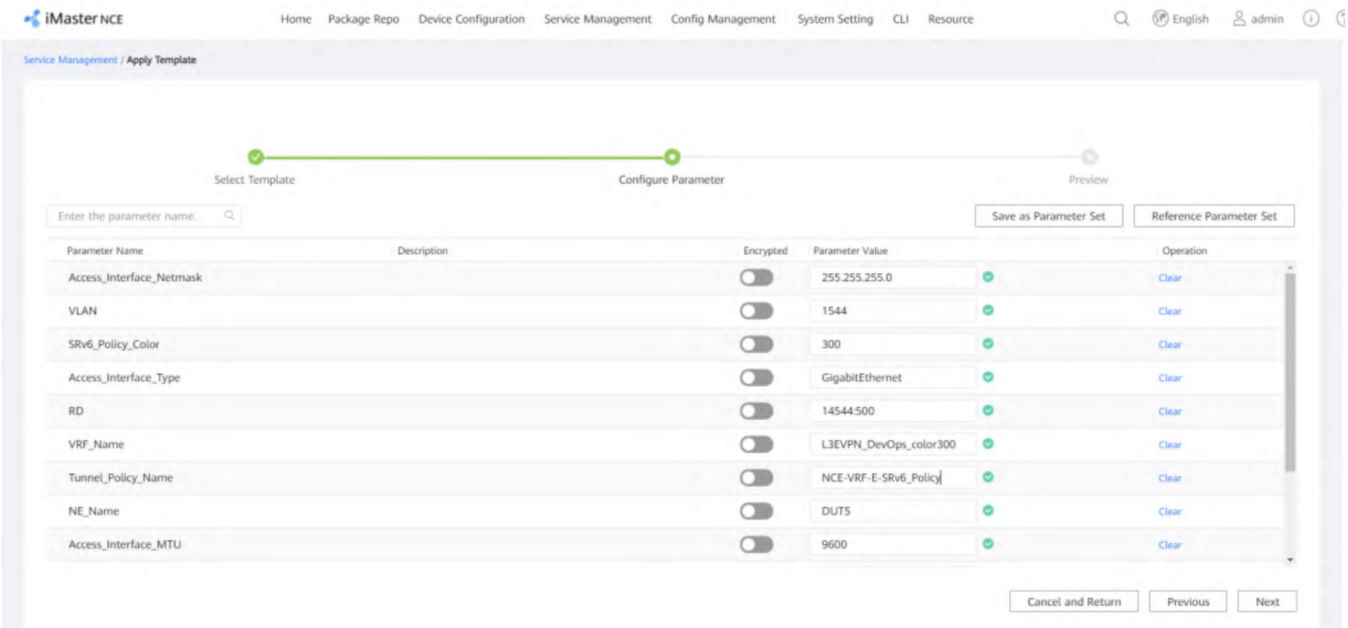

Parameter Name	Description	Encrypted	Parameter Value	Operation
Access_Interface_Netmask		○	255.255.255.0 ✓	Clear
VLAN		○	1544 ✓	Clear
SRv6_Policy_Color		○	300 ✓	Clear
Access_Interface_Type		○	GigabitEthernet ✓	Clear
RD		○	14544:500 ✓	Clear
VRF_Name		○	L3EVPN_DevOps_color300 ✓	Clear
Tunnel_Policy_Name		○	NCE-VRF-E-SRv6_Policy ✓	Clear
NE_Name		○	DUT5 ✓	Clear
Access_Interface_MTU		○	9600 ✓	Clear

Figure 26: SR Policy Configuration over DUTs

Conclusion

Huawei demonstrated an impressive range of new functions supported across the family of NetEngine 8000 routers and the new ATN910D-A chassis. The physical layer extensions for 400GbE interfaces and long-haul connection support are well noted. On many fronts, the Huawei success story of SRv6 is supported for more use case scenarios - with expanded failover techniques, header compression, and extended 5G slicing support.

Most notably, the enhanced functionality of the Huawei iMaster NCE-IP management solution was demonstrated coherently. While all of these demos were functional only, and the NCE's performance in realistic-scale service provider networks cannot be extrapolated from our results, the solution nevertheless showed an impressive closed-loop support. One of the biggest challenges for service providers is to manage complex technologies in their networks. NCE, by enabling the simplification of provisioning, maintenance, and network monitoring tasks, could play a vital role to accelerate the deployment of complex SRv6 technologies.

About EANTC

EANTC (European Advanced Networking Test Center) is internationally recognized as one of the world's leading independent test centers for telecommunication technologies. Based in Berlin, the company offers vendor-neutral consultancy and realistic, reproducible high-quality testing services since 1991. Customers include leading network equipment manufacturers, tier 1 service providers, large enterprises and governments worldwide. EANTC's Proof of Concept, acceptance tests and network audits cover established and next-generation fixed and mobile network technologies.

EANTC AG
Salzufer 14, 10587 Berlin, Germany
info@eantc.de, https://www.eantc.de/
[v1.0 20210616]