EANTC Independent Test Report

Huawei Intelligent IP Network Solution Empowered by NetEngine 8000 Series and Network Cloud Engine

April 2020

EANTC

HUAWEI

Introduction

In the beginning of 2020, EANTC carried out an independent functionality, performance, availability and manageability test commissioned by Huawei. We validated advanced aspects of the NetEngine 8000 family of routers, including new line card types, together with the Network Cloud Engine (NCE), and with ATN routers. At Huawei's Shenzhen headquarters, we conducted extensive tests with Huawei's team in January 2020.

The project was quite unusual: Usually, our manufacturer customers ask us to conduct tests with carrier products for a service provider audience, or alternatively with enterprise products for an enterprise audience. Usually, not only the products but also the network technologies under test are quite different between these target audiences: A carrier audience typically looks for MPLS, SDN, and large-scale WAN scenarios supported by complex network management solutions; whereas an enterprise audience is interested in data center interconnection, high throughput and port density in the data center, straightforward network management and simple network protocols. For the first time in EANTC's history, we got commissioned to run a single test for both service provider and enterprise audiences. Huawei was confident that the brand-new product families would be suitable for both scenarios, specifically due to the introduction of Segment Routing over IPv6, in short SRv6, and due to the latest updates in the NCE management solution.

Overall, Huawei focused on three goals for the evaluation: A simplified network, intelligent connectivity, and high availability.

Huawei presented a comprehensive transport network and management solution for the EANTC test. SRv6 constituted the heart of the architecture. This is a new, revolutionary transport technology, enabling the seamless interconnection of wide-area and data center networks. It eliminates the need to learn, configure or troubleshoot MPLS – a technology, that many enterprises had dreaded to deploy due to its perceived complexity. SRv6 by itself is not a trivial solution, if only because it involves IPv6 - that said, it is the only transport solution required in the end-to-end network once fully deployed. This simplicity is one of the great promises of SRv6.

An important aspect for enterprises is a clean, simple to operate network provisioning and monitoring solution. Huawei provided NCE to achieve this goal, and EANTC tested a range of functionalities successfully. Additionally, we evaluated the performance of new hardware including new high-density line cards for 100GigabitEthernet, and a 400GigabitEthernet line card.

Related to all three goals of the simplified network, intelligent connectivity, and high availability, SRv6 was a major focus area of our test. We evaluated a range of Huawei's router SRv6 implementation features including Layer 2 Ethernet VPN (L2 EVPN) and Layer 3 IP EVPN services. One of SRv6's strengths is the ability to create seamless segment routing tunnels across multiple network areas from the data center to the wide-area network (WAN). Huawei proved that the resulting end-to-end tunnels can be efficiently managed: We created tunnels with traffic engineering constraints, checked path calculation based on bandwidth, latency, and link cost. The EANTC team verified the resiliency as well: SRv6 path protection was tested independent of the network architecture (TI-LFA FRR), showing less than 11 milliseconds detection and rerouting time for single node failure. Additionally, SRv6 path egress protection was verified by utilizing the mirror SID and the function (END.M) to protect the egress node (tail node) of an SRv6 path. The results showed a fast response in 2 milliseconds to detect and switchover the traffic path between the egress routers.

Huawei presented new line cards to us for performance benchmarking: The 4T line card for the NetEngine 8000 series supporting high-density 100GigabitEthernet ports, and the brand-new 400GigabitEthernet line card with eight 400GE ports and another eight 100GE ports. In most of these scenarios except the 400GE prototype card, we conducted standard RFC2544 throughput and latency benchmarks and evaluated the power efficiency.

We also confirmed the support and attenuation budget of special optical modules with 80 km range for 50GigabitEthernet and 100GigabitEthernet, and an optical module with 40 km range for BiDir single fiber 50GigabitEthernet.

Finally, we also witnessed a complete configuration of SRv6 tunnels and features through the NCE's SDN controller functions. Huawei demonstrated that it is possible to control SRv6 provisioning, fault management and performance monitoring functions from NCE's graphical user interface. Specifically, the straightforward configuration of complex functionality is enabled through NCE. Using NCE, we verified SRv6 tunnel reoptimization without any packet loss based on live network utilization data. Moreover, Huawei's implementation of the draft standard iFIT was verified to monitor the loss and latency of selected types of real service traffic on the application layer.

Overall, our tests showed that the NetEngine 8000 family is designed for enterprise and carrier markets likewise. The EANTC team focused on network simplification, intelligent connectivity, and high availability functions, which are specifically important for large enterprise scenarios. We hope the detailed description of each test area below will provide the reader with insight into SRv6 test methodology applied to the latest generation of Huawei unified routers and management solutions.

Segment Routing IPv6

Testbed Description

DUT Code	Hardware Platform	Software Version
DUT1	NetEngine 8000 M8	V800R012C00
DUT2	NetEngine 8000 M8	V800R012C00
DUT3	NetEngine 8000 X4	V800R012C00
DUT4	NetEngine 8000 X8	V800R012C00
DUT5	NetEngine 8000 M14	V800R012C00
DUT6	NetEngine 8000 M14	V800R012C00
DUT7	NetEngine 8000 F1A	V800R012C00
DUT8	NetEngine 8000 M8	V800R012C00
DUT9	ATN980C	V300R006C00
DUT10	ATN910C-G	V300R006C00
DUT11	NetEngine 8000 M1A	V800R012C00
DUT12	NetEngine 8000 M6	V800R012C00
SDN Controller	NCE (IP Domain)	V100R019C00SPC600

Table 1: Testbed Components

SRv6 Service Provisioning

The power of SRv6 sources from the simplicity of the operations to create transport tunnels, layer 3 VPN (L3VPN), or layer 2 VPN (L2VPN) services. Those services are organically inherited from the MPLS technology but with more straight forward deployment steps and less protocol stack to manage and operate the services.

We began this test by verifying the functional capabilities of the Huawei selected routers as listed in Table 1 to compute the traffic-engineered (TE) paths, provisioning L3VPN/L2VPN services.

EVPN L3VPN, VPWS and L2VPN Services over SRv6

The unified deployment of EVPN and SRv6 in the transport network brings an easier way to create and provision the classical MPLS-based VPN services such as L3VPN, E-Line (VPWS), and L2VPN E-LAN. EVPN is a unified control plane protocol that supports many VPN services over a single MP-BGP instance. Complementary to that, SRv6 provides a unified transport protocol to encapsulate and route the traffic efficiently in the transport network.

The purpose of this testing section is to demonstrate the functionality of the SRv6 data plane to transport and forward multiple EVPN-signaled services. Huawei configured the routers in the testbed for the following services:

1. EVPN L3VPN over SRv6 Best Effort Tunnel
2. EVPN L2VPN over SRv6 Best Effort Tunnel
3. EVPN VPWS over SRv6 Traffic Engineering Tunnel

For the L3VPN service, Huawei created ten different VPN instances on the DUT1, DUT4, DUT8, DUT9. Each VPN instance support both IPv4 and IPv6 customer traffic. Also, Huawei configured ten different bride-domains on the same DUTs for L2VPN service. Huawei enabled EVPN on the DUT routers to exchange the customer IP routes for L3VPN service and the customer MAC addresses for L2VPN and VPWS services.

The Internet draft [I-D.filsfils-spring-SRv6-network-programming] defines multiple SRv6 functions that can be programmed on an SRv6-capable router. For example, the L3VPN service, the function code END.DT represents "cross-connect to a VRF" or END.DX represents "cross-connect to a next-hop" functions. SRv6 Service SID refers to an SRv6 SID that may be associated with one of the service-specific functions.

EVPN Services over SRv6 data plane requires the advertisement of the SRv6 Service SID in an EVPN route-type 1,2,3 and 5. The SRv6 Service SID is advertised in SRv6 Service TLV, as described in [draft-dawra-idr-SRv6-vpn-05]. The two objectives of exchanging SRv6 Service SID are to indicate the reachability of the egress router via SRv6 data plane, and the second goal is to signal the value of the VPN SID.

Huawei configured the data plane of the L3VPN and L2VPN services using SRv6 best effort (BE) tunnels, as shown in Figure 1. The established SRv6-BE tunnels between the PE's follow the calculated IS-IS shortest path (lowest metric). The VPWS service is an emulation of L2 point-to-point circuits. The practical use cases of VPWS require some traffic engineering considerations, like include some transit nodes to the explicit path between the head and the tail of the tunnel. For this reason, Huawei created SRv6-TE tunnels between DUT3-DUT7 and DUT2-DUT10 as an SRv6 data plane for VPWS services, as shown in Figure 2.

To confirm the proper functionality and operation of each VPN service, we started by verifying the control plane of the testbed through the EVPN operations. We checked the established EVPN peering between the DUTs, the successful installation of the remote MAC addresses in the matched bridge domain instance for L2VPN service, and the remote L3 prefixes in the matched VRF instance for L3VPN service. To evaluate the operation of the SRv6 data plane, we checked the "Local -SID End.DT4 and End.DT6 Forwarding Table" in each DUT.

Each L3VPN instance must allocate a SID (or VPN SID), and this SID will be used by the remote DUTs to reach a specific local VPN instance. Figure 3 shows the output of the Local-SID End.DT6 forwarding table for DUT9. Figure 4 shows the destination IPv6 address on the outer IPv6 header, which matches the VPN SID of the VPN ID 11.

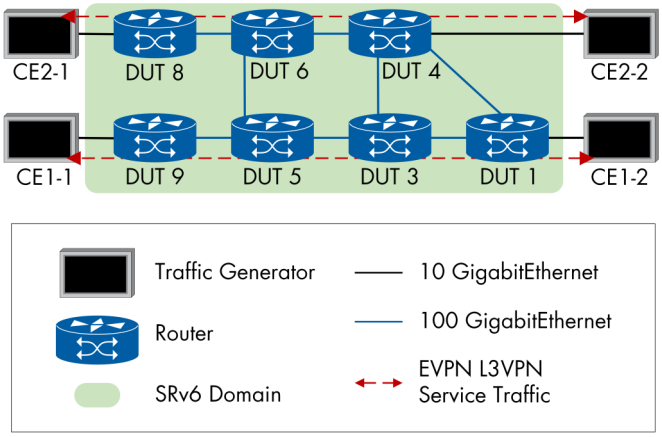

Figure 1: EVPN L3VPN over SRv6-BE

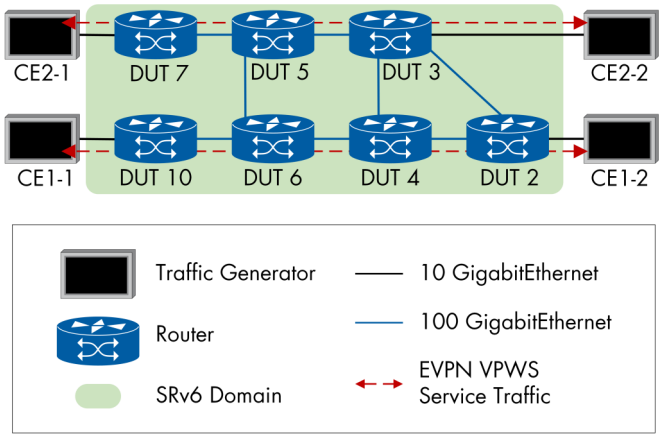

Figure 2: EVPN VPWS over SRv6-TE

```
<M1A-DUT9>display segment-routing ipv6 local-sid end-dt6 forwarding

                  My Local-SID End.DT6 Forwarding Table
                  -------------------------------------

SID         : 111::1:0:51/128                          FuncType : End.DT6
VPN Name    : 230                                      VPN ID   : 11
LocatorName: as                                        LocatorID: 1

SID         : 111::1:0:53/128                          FuncType : End.DT6
VPN Name    : 231                                      VPN ID   : 12
LocatorName: as                                        LocatorID: 1

SID         : 111::1:0:55/128                          FuncType : End.DT6
VPN Name    : 232                                      VPN ID   : 13
LocatorName: as                                        LocatorID: 1

SID         : 111::1:0:57/128                          FuncType : End.DT6
VPN Name    : 233                                      VPN ID   : 14
```

Figure 3: L3VPN SID Allocation

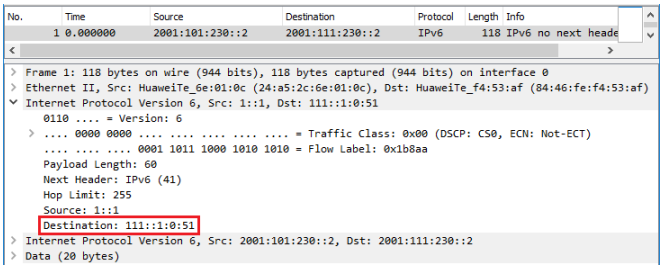

Figure 4: L3VPN Packet Capture

EVPN and SRv6 formulate the basis of the next-generation transport networks. In this test section, we verified the functional readiness of the Huawei NetEngine 8000 platforms to support successfully EVPN-signaled L3VPN, E-LAN L2VPN, and E-Line VPWS services using SRv6 data plane. For the next level of testing, EANTC recommends evaluating the scalability and the performance of Huawei NetEngine 8000 platforms because the service providers consider scalability and performance testing as necessary as functional testing once it comes to introduce new technology to their networks.

SRv6 Service Resiliency

Service continuity is one of the leading design aspects, which was considered during SRv6 development. Any link or node failure in the network should be protected to fulfill the tight SLA requirements for the next-generation services (i.e., access to the public cloud, UHD video streaming, or autonomous robotics control). SRv6 adopts a lot of robust service protection mechanisms to fulfill different levels of protection in the transport network. For example, SRv6 Topology-Independent Loop-Free Alternate Fast Reroute (TI-LFA FRR) handle the transit node failures with restoration time less than 50 ms. Moreover, SRv6 Path Egress Protection provides another level of protection for the dual-homed CE sites. In the following section of the test cases, we verified the capability of Huawei's solution to achieve the expected results of each protection mechanism.

SRv6 TI-LFA FRR

TI-LFA provides fast convergence in less than 50 ms in case of node failure. TI-LFA uses a backup path that pretends no dependencies on topology constraints and offers a more reliable FRR. Segment routing can encode the FRR backup path's entries to the segment-list. This enforces the packet to follow a loop-free path through the backup path. We started the test by setting the topology, as shown in Figure 5. TI-LFA and FRR were enabled under the IS-IS IPv6 process. DUT1 was selected to be the ingress node, and DUT11 was the egress node for ten L3VPN instances. We send IPv4 bidirectional traffic by the rate of 100,000 frames per second between DUT1 and DUT11. Then, we requested the Huawei test engineer to reboot the DUT5. After that action immediately, we observed frame loss in both directions of the traffic flow.

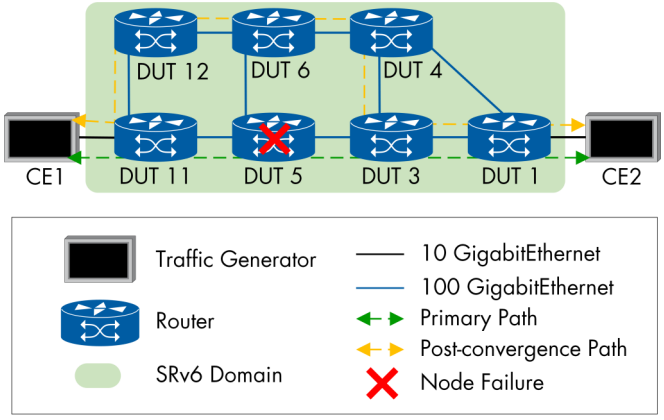

Figure 5: SRv6 TI-LFA FRR Topology

We repeated the test three times, measuring the frame loss during failover each time. The maximum loss was measured as 1085 frames. Based on the transmit rate, we calculated 10.85 ms maximum out-of-service time. This indicates that the out-of-service time was 10.85 ms. To check the restoration behavior, we asked the Huawei team to bring up DUT5 again. Subsequent to the network convergence period, the optimal SRv6 tunnel was established which crosses through DUT5. The traffic was restored with 0 frame loss.

SRv6 Path Egress Protection

The IETF Internet draft (draft-hu-rtgwg-SRv6-egress-protection-00) describes the required protocol extensions and procedures to protect the egress node (tail node) of an SRv6 path. The general idea of path egress protection is to use a mirror SID, with the function End.M, for protecting a VPN SID. The mirror SID (End.M) must always be the penultimate SID. Also, the Internet draft defines the required extensions for the IGP (IS-IS and OSPF) to support the advertisement of the mirror SID (End.M). In this test, Huawei configured the routers with IS-IS IPv6, which support new sub-TLV called "IS-IS SRv6 End.M SID". Figure 8 depicts the topology of the egress protection test case. DUT1 was CE router and dual-homed to DUT3 and DUT4. Huawei configured DUT3 as the primary tail router of the SRv6-TE tunnel, which transported the L3VPN service between DUT7 and DUT3 and 4. We verified the configuration by checking the END.M function is configured on DUT4 and mapped with DUT3 local SID. Then we requested the Huawei team to reboot DUT3 while the traffic was flowing in rate 100.00 frames per second to emulate node failure or out of service. The maximum lost frames were 201, which indicates that the required time to protect an egress node using the SRv6 data plane is 2 ms.

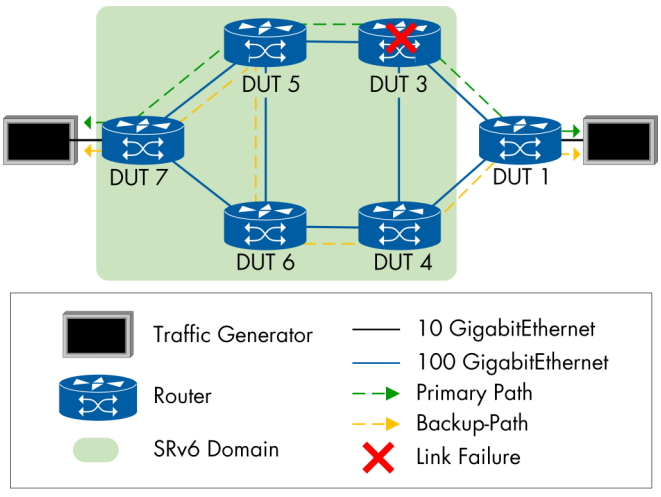

Figure 6: SRv6 Path Egress Protection Topology

In this section, we tested three main mechanisms that are commonly employed by the service providers for the service protection in the packet network. The achieved restoration time (2~3 ms) was significantly less than the typical value (30~50ms). In the more broaden networks with more complicated topologies, the restoration time could be higher, especially if we considered multiple links or nodes failure scenarios.

For future service resiliency tests, we are looking forward to testing with Huawei more advanced protection scenarios like SRv6 TI-LFA Shared Risk Link Group (SRLG).

Transport Network Evolution and Migration

SRv6 data plane has a proven simplicity and flexibility of service provisioning compared to the classical MPLS. SRv6 attracts the attention of the network operators to start preparing the migration plans from the classical MPLS to the SRv6 paradigm. In the classical MPLS networks, the primary two label distribution protocols are LDP and RSVP. LDP is usually deployed to establish MPLS-BE tunnels with less administrative effort. RSVP is more known for MPLS-TE tunnels signaling based on one or set of predefined constraints.

Huawei SRv6 solution facilitates the migration from LDP and RSVP MPLS tunnels in very smooth and straightforward migration steps. Moreover, the migration to SRv6 does not require all the midpoint routers in the network to be SRv6-aware. The minimum requirement is to enable IPv6 forwarding across all the transit routers, plus the ingress and egress routers must be SRv6-aware routers. In the following test cases, Huawei demonstrated the needed procedures for each MPLS migration scenario.

Scenario 1: Migration of MPLS LDP Tunnel to SRv6-BE Tunnel

The objective of this test case is to show the required steps to migrate the L3VPN service between two VPN sites, which is transported by MPLS-LDP tunnel. Also, to verify the traffic switchover to the SRv6 tunnel without service interruption or any packet loss.

We run this test case in three stages:

1. The L3VPN service was established between DUT11 and DUT3 through DUT5. Huawei enabled the LDP on DUT11, DUT5, and DUT3 to allocate the transport labels. And Huawei configured MP-BGP VPNv4 to exchange the VPN labels between the PE's (DUT11 and DUT3). Under the VPN-instance section, Huawei set the LDP policy to enforce the outbound traffic toward remote PE to flow through the MPLS LDP tunnel. We generate bidirectional traffic between CE1 and CE2 to ensure the service establishment.
2. In the second stage, Huawei configured IS-IS IPv6 on all the routers, enabling SRv6 on the ingress and egress routers (DUT11 and DUT3), assign the node and VPN SIDs, establish IBGP session and enable VPNv4 neighbor by using the IPv6 address of the PE's (DUT11 and DUT3), creating new SRv6-BE tunnels over DUT11 and DUT3. After that, Huawei replaced the configuration of the LDP tunnel policy under the VPN instance with the SRv6-BE tunnel, to switch the traffic flowing from the LDP labeled tunnel to the new SRv6-BE. We generate bidirectional traffic between CE1 and CE2 to ensure the service establishment, and we captured a sample of the packets on the line to verify the new SRv6 packet encapsulation and the zero presence of MPLS encapsulated packets.
3. Finally, we removed all the MPLS-related configurations from DUT3, DUT5, and DUT11 and keep only SRv6 configurations to make sure there are no configurations dependency still exists.

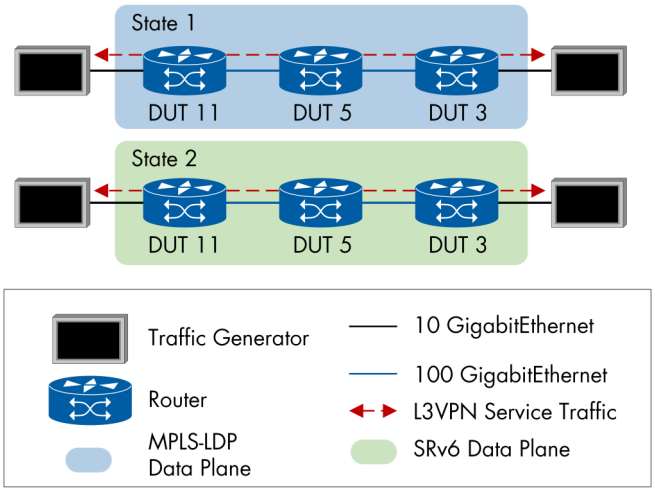

Figure 7: Migration from MPLS-LDP to SRV6 Data Plane

Scenario 2: Migration of MPLS RSVP Tunnel to SRv6-TE Tunnel

In this scenario, we used an RSVP-TE tunnel for an L3VPN service. To check the seamless switchover from the RSVP-TE transport tunnel to the SRv6-TE tunnel, we went through the following steps:

1. The L3VPN service was established between DUT9 and DUT1 through DUT5 and DUT3. Huawei enabled the RSVP-TE tunnel on DUT9, DUT5, DUT3, and DUT1 to allocate the transport labels. And Huawei configured MP-BGP VPNv4 to exchange the VPN labels

between the PE's (DUT9 and DUT1). Under the VPN-instance section, Huawei applied the RSVP-TE policy to enforce the outbound traffic toward remote PE to flow through the MPLS RSVP-TE tunnel. We generate bidirectional traffic between CE1 and CE2 to ensure the service establishment.

2. In the second state, Huawei configured IS-IS IPv6 on all the routers, enabling SRv6 on the ingress and egress routers (DUT9 and DUT1), assign the node and VPN SIDs, establish IBGP session and enable VPNv4 neighbor by using the IPv6 address of the PE's (DUT9 and DUT1), creating new SRv6-TE tunnels over an explicit path (DUT9-DUT5-DUT3-DUT1). After that, Huawei replaced the configuration of the RSVP-TE tunnel policy under the VPN instance with the SRv6-TE tunnel, to switch the traffic flowing from the RSVP-TE labeled tunnel to the new SRv6-TE. We generate bidirectional traffic between CE1 and CE2 to ensure the service establishment, and we captured a sample of the packets on the line to verify the new SRv6 packet encapsulation and the zero presence of MPLS encapsulated packets.

3. Finally, we removed all the MPLS-related configurations from DUT1, DUT3, DUT5, and DUT9 and keep only SRv6 configurations to make sure there are no configurations dependency still exists.

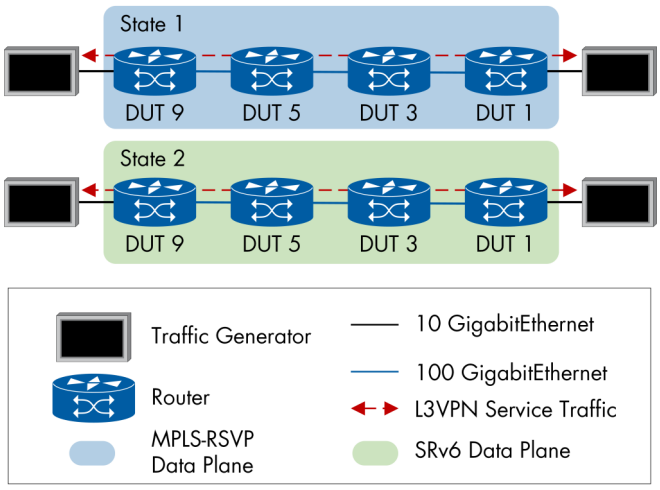

Figure 8: Migration from MPLS-RSVP to SRV6 Data Plane

We verified the non-disruptive behavior of the protocol migration to SRv6 and without any compromises to establish a new SRv6-TE tunnel based on the old RSVP tunnel's constrains. During the test, we didn't observe any packet loss or service interruption.

Hardware Architecture and Capacity

The second chapter of this report sheds light on the newly introduced router series (Huawei NetEngine 8000 X8). More precisely, we tested throughput and power consumption for the 4T (40x100GE) line card. Moreover, Huawei has added a new line card to its portfolio, this card supports (8x400GE and 8x100GE) ports. We checked the throughput rate of the 400GE along with latency measurements.

To test the forwarding throughput, EANTC uses different traffic flows of IPv4, IPv6, and a combination of both. This methodology reflects the actual throughput performance of the routers in the typical deployments of service providers. Based on Huawei's request, we performed the throughput test cases for the NetEngine 8000 X8 4T line card and 400GE ports using IPv4 only traffic.

Throughput Capacity and Power Consumption of NetEngine 8000 X8 4T (40x100GE) Line Card

In this test, we verified the throughput performance of the 4T (40x100GE) line card. We connected the traffic generator (Ixia IxNetwork) to the DUT (40x100GE line card) through 40 links. We configured a full mesh traffic flow between all the 40 ports to validate the maximum throughput performance on the line card. According to Huawei's request, we generated only IPv4 traffic based on different frame sizes, as defined in the RFC 2544. Table 2 summarizes the achieved results.

The other part of this test is to verify the power consumption of the line card on it operates in full load. So, we requested the Huawei team to measure the input power toward the power supplies of the chassis by a power meter. We started by measuring the input power without plugging the line card. The measured power was 1970.4 watt. To measure the power consumption, we plugged the line card to the chassis and generated IPv4 traffic with 550 Byte frame size. The measured output power was 3100.3 wat which means the line card consumed 1129.9W in a full load.

Throughput Capacity and Forwarding Latency of 400GE port on NetEngine 8000 X8

With the 400GE port technology, the packet network capacity is standing on a new edge that is a key enabler for 5G-ready networks. In this test, we verified the performance of a pair of 400GE ports in terms of throughput and forwarding latency. We connected the traffic generator (Spirent Test Center) with the DUT (Huawei NetEngine 8000 X8 8X400GE+8X100GE Line Card).

Frame Size (Byte)	Throughput		Latency (ns)		
	FPS	Gbit/s*	Min	Max	Average
256	1,811,594,094.1	4000	8537	20365	16278.699
512	939,849,570.6	4000	8727	18400	16014.438
1024	478,927,175.7	4000	8852	19475	15972.851
1280	384,615,363.2	4000	8870	18160	15937.381
1518	325,097,511.3	4000	8970	19260	16018.330
9200	54,229,932.7	4000	10150	18797	16762.815

Table 2: Throughput and Latency Results of NetEngine 8000 X8 4T Line Card
*The raw Ethernet throughput Gbit/s includes 8 bytes of the preamble and 12 bytes of Inter Frame Gap on wire

Frame Size (Byte)	Throughput		Latency (ns)		
	FPS	Gbit/s*	Min	Max	Average
256	362,318,840	800	9,180	27,500	14,214
512	187,969,924	800	9,180	30,920	14,122
1024	95,785,440	800	9,150	29,200	14,206
1280	76,923,076	800	9,220	28,190	14,125
1518	65,019,505	800	9,130	27,340	14,146
9200	10,845,987	800	10,120	28,390	15,098

Table 3: Throughput and Latency Results of 400GE Port
*The raw Ethernet throughput Gbit/s includes 8 bytes of the preamble and 12 bytes of Inter Frame Gap on wire

FIB Scalability

This test aims to verify the datasheet's number of entries that can be stored in the Forwarding Information Base (FIB) without any packet loss or increased forwarding latency. According to the Huawei team, the NetEngine 8000 X8 platform processes the routing information base (RIB) by the Main Processing Unit (MPU). Each active module receives a copy of the active and most recent FIB from MPU and caches it in the local Line Processing Unit (LPU). Each IP address version has an independent FIB with a distinct capacity.

To verify the published FIB capacities in the NetEngine 8000 X8 40x100GE line card datasheet, we connected the DUT to the Spirent test center by two 100GE ports, as shown in Figure 13. After that, we established a BGP session to advertise 4.1 million IPv4 prefixes and 2.1 million IPv6 addresses.

We selected the advertised prefixes to be diverse and incontiguous prefixes with a variety of prefix lengths.

During the prefixes exchanging between the Spirent TC and the Huawei NetEngine 8000 router, we were checking the count of the learned prefixes. After a couple of minutes, the number of received BGP routes on the DUT was 4.1 million IPv4 and 2.1 million IPv6. Then, we checked the count of installed routes in the FIB through the router's CLI. We observed 2 million IPv4, and 1 million IPv6 prefixes were only installed in the FIB. This matches with the published values in the datasheet of Huawei NetEngine 8000 X8 router. The following table summarizes the capacity of the achieved result of the FIB table.

Prefix Type	Advertised Prefixes Count	FIB Installed Prefixes Count
IPv4 only	4,100,000	4,000,000
IPv6 only	2,100,000	2,000,000
IPv4+IPv6	IPv4: 4,100,000 IPv6: 2,100,000	IPv4: 2,000,000 IPv6: 1,000,000

Table 4: FIB Capacity Testing Results

BGP Routes Learning Rates

In this test, we measured the speed of learning BGP IPv4, IPv6, or combination of both prefixes type and installing the prefix in the forwarding information table (FIB). We established an MP-BGP session between Spirent Test Center (STC) and Huawei NetEngine 8000 X8. The BGP peering was configured to support IPv4 and IPv6 prefixes exchange. On the STC, we defined 4 million IPv4 prefixes and 2 million IPv6 prefixes. We started our test by the IPv4 prefixes. Spirent advertised 4 million IPv4 prefixes to the DUT.

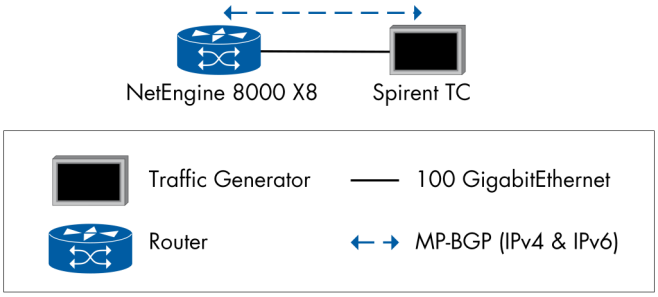

Figure 9: FIB Scalability and BGP Routes Learning Rate Test Topology

Simultaneously, Spirent was forwarding traffic toward the DUT. The DUT can only be capable of forwarding the received traffic once the received BGP learned routes are installed in the FIB. The delta time that is between 100% and 0% frame loss; this is the required time to thoroughly learn the 4 million IPv4 prefixes and install them actively in the FIB. Table 5 summarizes the learning rate for each prefix type.

Prefix Type	Count	Max Delta Time	Learning Rate
IPv4 only	4,000,000	37s	108,108 routes/s
IPv6 only	2,000,000	35s	57,142 routes/s
IPv4 + IPv6	IPv4: 2,000,000 IPv6: 1,000,000	30s	100,000 routes/s

Table 5: BGP Routes Learning Rate

Long-Distance Laser Support Capability

The direct long-distance connectivity between the routers in different geographical locations requires a special type of optical module to transmit enough power till the other end of the link. We evaluated three types of optical modules with different distances and different interface speeds. We used an actual fiber cable (manufactured by Corning) with a length of 40KM and 80KM to serve the purpose of this test. After we plugged the optical modules in the routers and brought up the ports, we generated IPv4 and IPv6 traffic to detect any failure or abnormal performance. Table 6 summarizes the achieved results.

	Model	Layer 1 Throughput	Frame Loss Percentage	Max Latency (µs)	Average Latency (µs)
50GbE BIDI 40KM (Single Fiber)	QSFP28-50G-1309TX/1295RX-40km-SM-PAM4	49.95 Gbit/s	0%	217	214
50GbE 80KM	QSFP28-50G-1295~1310nm-80km-SM-01	49.95 Gbit/s	0%	414	412
100GbE 80KM	QSFP28-100G-1295~1310nm-80km-SM-01	99.9 Gbit/s	0%	823	587

Table 6: Optical Module Throughput and Latency Measurements

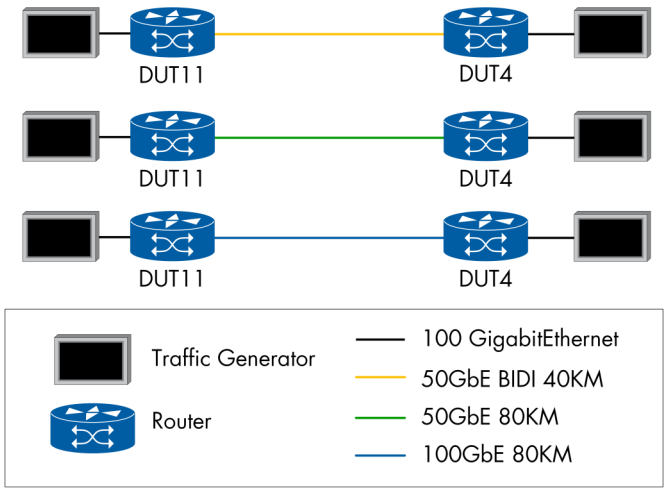

Figure 10: Different Optical Module Types Test Connectivity

Network Cloud Engine

The third chapter of this test focuses on centralized control, automated service provisioning, and next-generation network services using Huawei iMaster Network Cloud Engine for IP domain (NCE IP Domain).

Referring to the Huawei website, "iMaster NCE (IP Domain) centrally manages, controls, and analyzes IP devices such as NetEngine, ATN, CX, and PTN series NEs in a unified manner. Ideal for IP private line, IP core, 5G transport, and metro network scenarios, it provides functions such as device plug-and-play and service automation to enable automated full-lifecycle network management and maintenance. With real-time monitoring of network traffic and quality, iMaster NCE (IP Domain) leverages big data analytics to identify network trends in real-time and implement proactive maintenance and closed-loop optimization through service control and optimization".

Huawei introduced NCE (IP Domain) to us with a lot of modules and features. In our intended test, we demonstrated and used only the following modules; Network Management, Network Slice, Network Performance Analysis, and Network Path Management modules. Huawei deployed redundant instances of NCE installed on a physical server (Huawei Tai Shan 200). The used NCE version was (V100R019C00SPC600). Huawei testing team informed us this version is a pre-commercial or beta software version. The following table lists the details of the testbed setup of the NCE.

The test topology of NCE (IP Domain) chapter consists of 12 routers from different models of NetEngine 8000 and ATN platforms. Table 8 list down the DUTs hardware platform and software version.

We focused on three areas to demonstrate and verify the capabilities of NCE to:

1. Calculate SRv6 policy paths based on different constraints
2. Provision L3VPN services over SRv6 data plane
3. Applying next-generation OAM techniques using iFIT Monitoring

Component	Version
NCE (IP Domain)	V100R019C00SPC600
Data Base	GaussDB V100R003 Gauss100 OLTP V300R001
OS	EulerOS 2.8
Hypervisor	FusionCompute 8.0.0
Physical Server	TaiShan 200

Table 7: NCE Test Setup

DUT Code	Hardware Platform	Software Version
DUT1	NetEngine 8000 M8	V800R012C00
DUT2	NetEngine 8000 M8	V800R012C00
DUT3	NetEngine 8000 X4	V800R012C00
DUT4	NetEngine 8000 X8	V800R012C00
DUT5	NetEngine 8000 M14	V800R012C00
DUT6	NetEngine 8000 M14	V800R012C00
DUT7	NetEngine 8000 F1A	V800R012C00
DUT8	NetEngine 8000 M8	V800R012C00
DUT9	ATN980C	V300R006C00
DUT10	ATN910C-G	V300R006C00
DUT11	NetEngine 8000 M1A	V800R012C00
DUT12	NetEngine 8000 M6	V800R012C00

Table 8: DUT Hardware and Software Details

SRv6 Policy Path Calculation based on various Constraints

Huawei positioned NCE (IP Domain) as centralized software-defined networking (SDN) controller for the IP-based networks. As a centralized LSP path computation controller, NCE collects the underlying-network topology and IP reachability information using the BGP-LS protocol. Moreover, the Huawei team explained to us the theoretical mechanism to report the latency of the underlying network's links.

In this test case, we verified the creation of SRv6 paths based on different network constraints, including IGP cost, Link latency, Explicit-path, and bandwidth balancing. We started the test by enabling the NCE to automatically discover the underlying network and visualize the topology using the network management module. Figure 11 shows the discovered physical topology of 12 routers and the physical links between them.

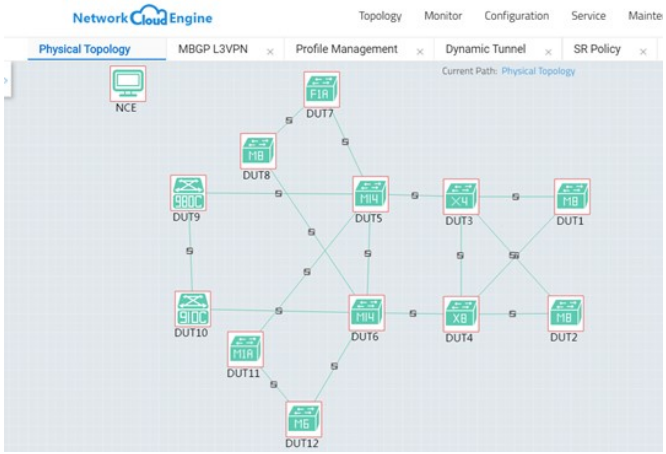

Figure 11: Physical Topology

SRv6 Path Calculation based on IGP Cost

NCE was capable of collecting the configured values of the IS-IS cost on the physical links. Based on the selected head and tail of the SRv6 path, NCE recommended the lowest-cost path. We selected DUT7 and DUT3 as the endpoints of the SRv6 tunnel, and NCE chose the lowest-cost path, as shown in Figure 12. Through DUT7-DUT5-DUT6-DUT4-DUT3. If the network operator confirms the recommended path by NCE and clicks on the "Configure" button, NCE will propagate the SRv6 policy to the ingress router (DUT7) using PCEP. We verified the installation of the newly configured SRv6 policy on DUT7 using the direct CLI access to DUT7 and generating traffic (IPv4 and IPv6) from DUT7 to DUT3 without any frame loss or routing loops.

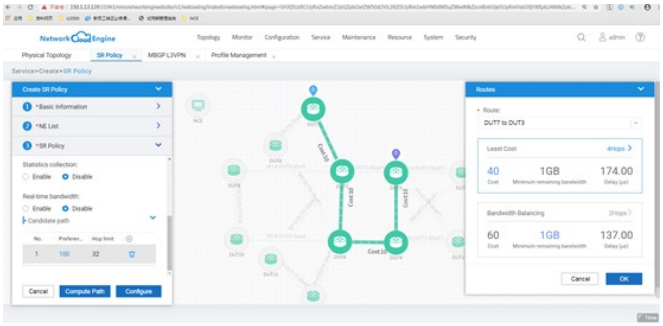

Figure 12: SRv6 path calculation based on the IGP cost

SRv6 Path Calculation based on Latency

Link latency is another attribute that NCE can consider to build an SRv6 path. Huawei team explained to us the used mechanism to report the latency on the physical links. Initially, Tow-Way Active Measurement Protocol (WAMP) was enabled on the physical links and signal the measured values to the IS-IS instance on the local router. IS-IS advertised the link attributes, including the latency to the locally enabled BGP-LS instance. Then, the local BGP-LS NLRI was propagated to NCE via the route reflector (DUT6). Due to the limited time during the test execution, we couldn't verify in detail the TWAMP session and the interaction with the locally configured IS-IS instance.

To apply an actual link latency, Huawei used a 20 km long fiber cable between DUT7 and DUT5. NCE reported the latency between the DUT7 and DUT5 120 microseconds. After that, we selected DUT7 and DUT3 again to be the ends of the new SRv6 path. Figure 13 shows the calculated SRv6 path based on the lowest-delay constrain. After we applied the selected minimum delay path, we verified the installation of the new SRv6-TE policy again and confirmed the traffic flow through the path (DUT7-DUT8-DUT6-DUT5-DUT3).

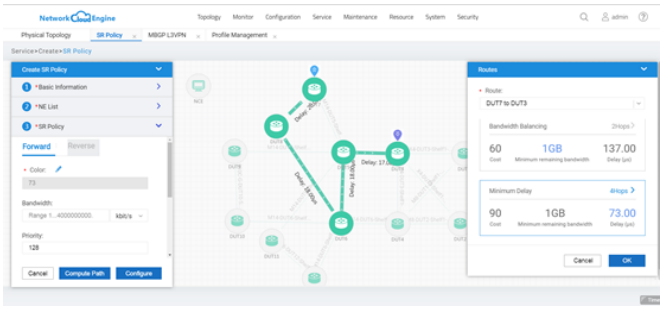

Figure 13: SRv6 path calculation based on the minimum delay

The second goal of this test is to check the responsiveness time of NCE to adjust the optimal path once the concerned attribute is changed. So, we asked Huawei to replace the fiber cable between DUT7 and DUT5 with a shorter fiber cable (length of 5M). In less than 1 minute, NCE was capable of proposing a new lowest-day path between DUT7 and DUT3 based on the new link latency value, as shown in Figure 14.

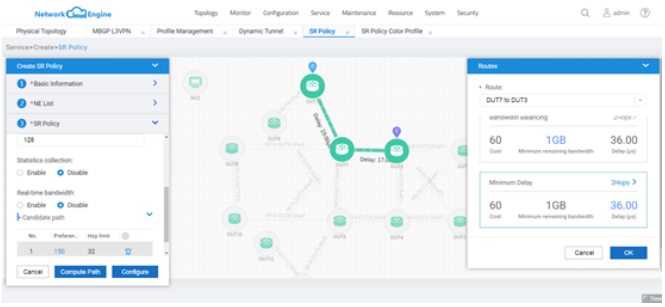

Figure 14: SRv6 path re-optimization based on the minimum delay

SRv6 Path Calculation based on Bandwidth Utilization

It is common in the service provider or enterprise networking to get a request for setting up a traffic engineering tunnel based on a specific bandwidth. The MPLS-RSVP tunnel allocates the requested bandwidth on the level of the logical tunnel. A shortcoming of RSVP-TE tunnel is the bandwidth of the logical tunnel is allocated without real-time monitoring of the actual available bandwidth on the physical links. Huawei demonstrated the capabilities of the NCE controller to establish an SRv6-TE tunnel based on the actual remaining links bandwidth of the End-to-End path. We started the test by creating an SRv6-TE policy that allocated 1Gbps bandwidth between DUT7 and DUT3 through DUT5. The name of this policy was "Policy 1". Then, we generated traffic in the rate of 1Gbps to completely utilize all the available bandwidth.

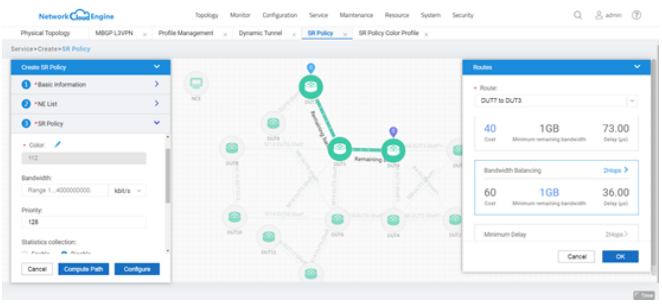

Figure 15: SRv6 path calculation based on the available bandwidth

After that, we tried to create another policy between the same tunnel's ends (DUT7 and DUT3) with the 1Gbps bandwidth requirement. NCE proposed a new path that matches the requested service bandwidth with the remaining bandwidth on each physical link through the path.

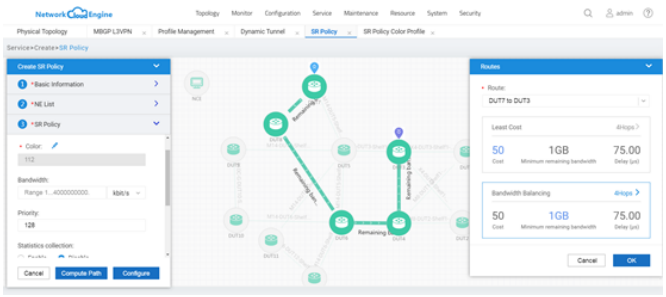

Figure 16: SRv6 path re-optimization based on the available bandwidth

SRv6 Path Calculation based on Explicit Path

The set up of the SRv6-TE path is not limited to the actual link attributes or the IS-IS link cost. The network operators have the freedom to establish TE tunnels based on the explicit paths. We verified the possibilities and options to design an explicit path. We asked Huawei team to design SRv6-TE policies based on include or exclude specific node or link, include node strictly or loosely. As an example, we set up a path between DUT7 and DUT3 based on an explicit path rule which excludes DUT5. We verified the configured SRv6-TE policy on DUT7 and the SIDs that are listed in the SRH, as shown in Figure 17 and Figure 18.

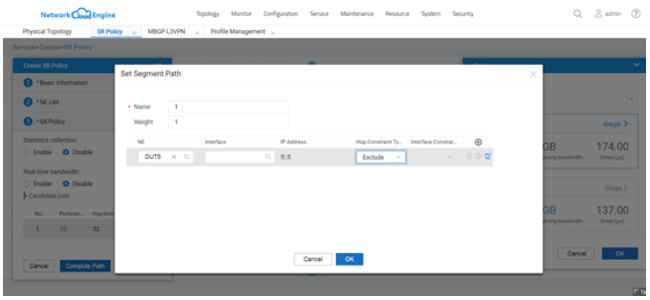

Figure 17: Defining SRv6-TE Explicit Path Policy

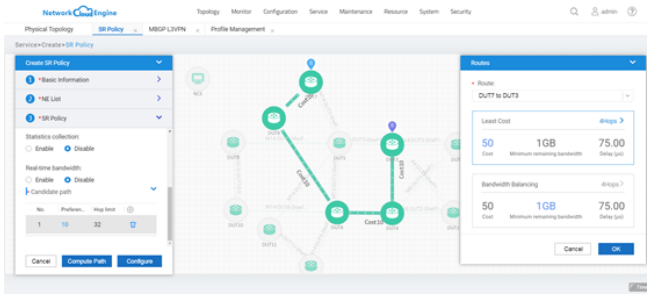

Figure 18: SRv6 path setup based on the explicit path

NCE (IP Domain) demonstrated a flexible and interactive way to program SRv6-TE policies based on various constraints. We verified the policy propagation and implementation through the Path Computation Element Protocol (PCEP) between the NCE (PCE) and the DUTs (PCC).

Service over SRv6 Policy provisioning

Continuing to demonstrate the possible options of provisioning networking services by Huawei NCE (IP Domain) controller, the Huawei team asked us to verify the MBPG L3VPN over SRv6 service provisioning through the NCE (IP Domain). From the "Network Management" module, a Huawei test engineer started to define the service template which includes the service type (MBGP L3VPN), the service nodes (DUT3 and DUT9), the VRF address family (IPv4, IPv6 or Both), routing policy (PE-CE protocol), and finally the data plane encapsulation (MPLS or SRv6) and the selected tunnel associated between the PE's. After completing the configuring of all the required parameters, a Huawei test engineer applied the service template and we verified the new configurations on DUT3 and DUT9.

To confirm the traffic flow between the CE sites, we generated IPv4 and IPv6 traffic between the emulated CE sites, as shown in Figure 19. The traffic was flowing between the sites as expected without any frame loss.

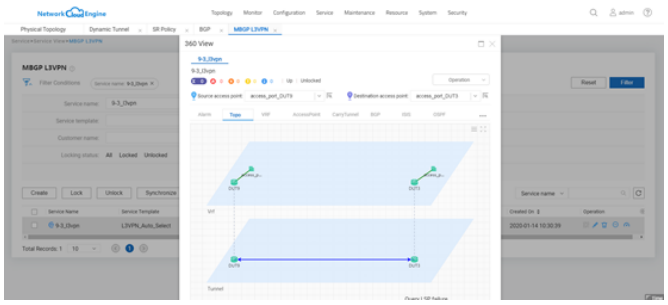

Figure 19: Logical Topology of L3VPN Service

iFIT Monitoring

According to the Internet draft "draft-song-opsawg-ifit-framework-00", In-situ Flow Information Telemetry (iFIT) is a framework for applying techniques such as In-situ OAM (iOAM) and Postcard-Based Telemetry (PBT) in networks.

In this test, the Huawei team introduced iFIT framework with NCE (IP Domain) to detect report packet loss or packet latency exceeding a predefined threshold based on the Hop-by-Hop and End-to-End paths. Huawei configured an SR-MPLS L3VPN tunnel. The same tunnel has two traffic flows, one is GTP-emulated traffic and the other is SCTP-emulated traffic. Huawei started the configurations of iFIT by defining the thresholds of packet loss and delay for each traffic flow. The Huawei configured the triggering policy whenever a predefined consecutive threshold-crossing occurrence. Clearing policy also configured to set the number of consecutive restoration time is required to clear the alarm.

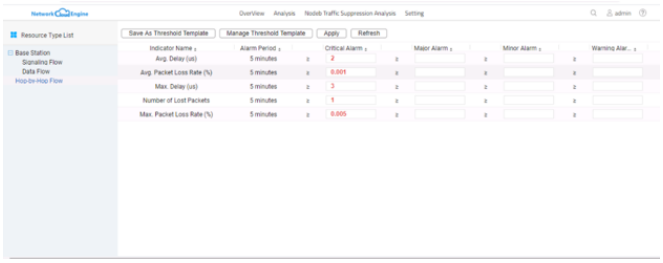

Figure 20: Configuring Hop-by-Hop Flow Analysis

We started generating GTP and SCTP-emulated traffic because there was no packet loss, and the delay was under the predefined threshold, no alarm was triggered, and the delivered services met the SLA. To emulate the packet loss incident, Huawei enabled a traffic shaping policy on the link between DUT5 and DUT3. NCE detected the packet loss within 1 minute. After that, the hop-by-hop test was triggered, as shown in Figure 21. The hop-by-hop test can be enabled on-demand and it will stop automatically after the trigger released (the packet loss is less than the threshold).

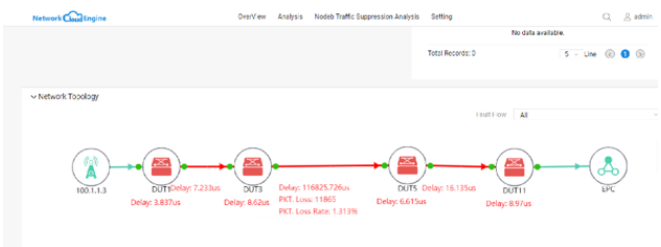

Figure 21: Result of Hop-by-Hop Flow Analysis

Conclusion

Huawei has provided the family of NetEngine 8000 routers including new line cards and a new NCE software version for our test. The EANTC team is able to confirm all of Huawei's functional, performance, high availability and manageability claims made by Huawei for this project. Specifically, Huawei's advances in the SRv6 implementation were impressive: This network architecture, combined with the Network Cloud Engine (NCE), allows more advanced end-to-end configurations with better network utilization and more intelligent traffic engineering constraints – while at the same time simplifying network provisioning and operations.

This is specifically important for large enterprises, who are looking to deploy a unified end-to-end architecture ready for the scale and diversity of future network requirements: Whether in the data center or in the WAN, SRv6 together with the NCE management solution showed that it can cover any network connectivity requirement. The new Huawei NetEngine 8000 line cards for 100GigabitEthernet and 400GigabitEthernet exhibited impressive port density, an aspect important for data center networks.

As seen in our tests, Huawei's unified carrier and enterprise solutions enable Huawei to offer advanced, cutting-edge technologies to both target audiences at the same time.

About EANTC

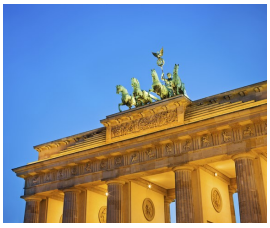

EANTC (European Advanced Networking Test Center) is internationally recognized as one of the world's leading independent test centers for telecommunication technologies. Based in Berlin, the company offers vendor-neutral consultancy and realistic, reproducible high-quality testing services since 1991. 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.com, http://www.eantc.com/
[v1.1 20200605]