EANTC Independent Test Report

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

April 2020

EANTC

HUAWEI

Introduction

Huawei commissioned EANTC to independently validate advanced aspects of the NetEngine 8000 family of routers and the ATN series, including new line card types, together with the Network Cloud Engine (NCE). At Huawei's Shenzhen headquarters, we conducted extensive tests with Huawei's team in January 2020.

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. Segment Routing over IPv6, in short SRv6, constituted the heart of the architecture. The transport services were provisioned and monitored by NCE. Additionally, we evaluated the performance of new hardware including a 400GigabitEthernet line card, and the performance of precision time protocol (PTP) clocking implementations.

Related to all three goals of 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 (L2EVPN) 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 protection was tested independent of the network architecture (TI-LFA), showing less than 20 milliseconds detection and rerouting time for a single node or link failure.

In fact, SRv6 is a relatively new technology not widely deployed in service provider networks yet. It is important for such a technology to efficiently support live migration of existing MPLS services to SRv6 with minimal impact. As part of our test, we verified Huawei's implementation of zero-loss migration to SRv6 from MPLS LDP and MPLS RSVP-TE, respectively.

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. Related to 5G use case scenarios, slicing support was verified with NCE using channelized interfaces on the routers – we created multiple slices, confirmed they are not affecting each other, and tore them down using NCE. Finally, Huawei's implementation of the draft standard iFIT was verified to monitor loss and latency of selected types of real service traffic on the application layer.

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. Additionally, we benchmarked a NetEngine 8000 M14 chassis fully loaded with 100GE cards. 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 BiDi single fiber 50GigabitEthernet.

Finally, we verified the precision of boundary clocks (T-BC) implemented in the NetEngine 8000 family and the ATN series. We confirmed boundary clock precision complying with G.8273.2 Class C requirements which is more than suitable for the latest 5G requirements.

Overall, the NetEngine 8000 family is designed for carrier, cloud provider and enterprise markets. The EANTC team focused on network simplification, intelligent connectivity, and high availability functions suitable for all customer groups. 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 carrier-grade routers and management solution.

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 Tunnel Creation and 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, and check some of the OAM features like SRv6 tunnel ping and traceroute.

SRv6 Traffic-Engineering (SRv6-TE) Tunnel Creation

The source routing is the main idea of the Segment Routing (SR). Meaning that the ingress node of the service is in charge of encoding the path which the packet will follow to reach the destination (egress node). The path is encoded in the Segment Routing Header (SRH) as a list of IPv6 addresses that represents the locators and the functions. The ingress node computes the SRv6 Best-Effort (SRv6-BE) path based on the shortest path calculations provided by an Interior Gateway Protocol (IGP) (i.e., IS-IS).

Moreover, the ingress node can compute an SRv6 Traffic-Engineering (SRv6-TE) path based on explicit network constraints, which are requested by the network administrator. The first objective of this test section was to verify the functional capability of the Huawei NetEngine 8000 series routers to establish SRv6-TE tunnels. The second objective was to verify the operation of SRv6 OAM using the tunnel Ping and traceroute SRv6 encoded function (OAM Endpoint with Punt, for short "END.OP").

Figure 1 shows the network topology of these test cases. Huawei enabled the SRv6 related configuration on DUT1, DUT3, DUT7, and DUT8. This includes:

- Configure IPv6 addresses in all the interfaces between the DUTs; all the DUTs must be IPv6-capable
- Configure ISIS IPv6 instance in all the DUTs
- Enable SRv6 globally on the DUT1, DUT3, DUT7, and DUT8 and define a local locator for each and the required function based on the test case (i.e., END, END.OP)

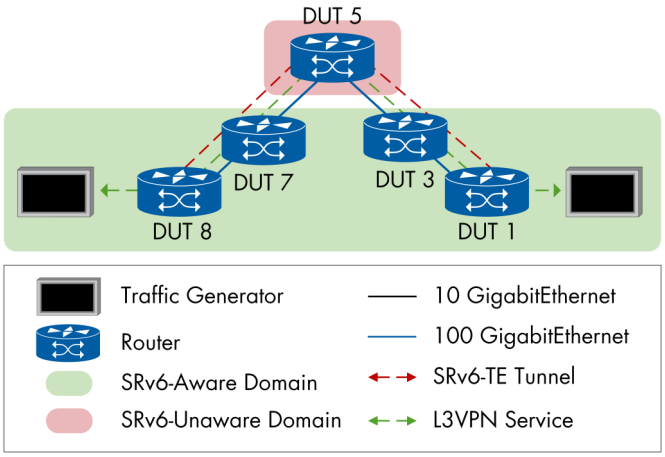

Figure 1: SRv6-TE Tunnel Cross SRv6-Unaware Domain

The Huawei team disabled all the SRv6 related configurations on DUT5 intentionally, to prove the capabilities of forwarding the SRv6 traffic within the SRv6-unaware domain. Huawei created an explicit-path SRv6-TE tunnel between DUT8 and DUT1. The path was defined to include DUT7, DUT4, and DUT1. The Huawei software demonstrated the capability to program an explicit-path based on including strict or loose nodes. To verify the SRv6 functional support and the end-to-end traffic flow, Huawei created an L3VPN service between DUT8 and DUT1 and generated the traffic between the two locations of the VPN instance. EANTC verified the SRv6-TE tunnel establishment, and we assured the traffic was flowing between the VPN sites as defined in the explicit-path list. Because the DUT5 was an SRv6-unaware router, it can not process the SRH. However, DUT5 was still able to forward the traffic based on the outer destination IPv6 address successfully.

In the second test case of this section, EANTC verified the operation of SRv6 Operation, Administration, and Maintenance (OAM) using tunnel ping and traceroute. The Internet-Draft draft-ali-spring-srv6-oam-02 defines the OAM building blocks and mechanism the can be implemented using these building blocks in the SRv6 data plane. Huawei testing team allocates the OAM SID END.OP on the ingress and egress routers (DUT8 and DUT1, respectively) for bidirectional traffic scenarios. Using the CLI of DUT1 and DUT8, we executed the ping and tracert commands. Figure 2 and 3 shows the output of the executed commands.

Huawei testing team allocates the OAM SID END.OP on the ingress and egress routers (DUT8 and DUT1, respectively) for bidirectional traffic scenarios. Using the CLI of DUT1 and DUT8, we executed the ping and tracert commands. Figure 2 and 3 shows the output of the executed commands.

```
[~M8-DUT1]ping srv6-te policy endpoint-ip 8::8 color 21 end-op 108::1080 -a 1::1
  PING srv6-te policy : 100  data bytes, press CTRL_C to break
  srv6-te policy's segment list:
  Preference: 21; Path Type: primary; Protocol-Origin: local; Originator: 0, 0.0.0.0;
Discriminator: 21; Segment-List ID: 1; Xcindex: 1; end-op: 108::1080
    Reply from 108::1080
    bytes=100 Sequence=1 time=1 ms
    Reply from 108::1080
    bytes=100 Sequence=2 time=1 ms
    Reply from 108::1080
    bytes=100 Sequence=3 time=1 ms
    Reply from 108::1080
    bytes=100 Sequence=4 time=1 ms
    Reply from 108::1080
    bytes=100 Sequence=5 time=1 ms

  --- srv6-te policy ping statistics ---
    5 packet(s) transmitted
    5 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 1/1/1 ms
```

Figure 2: Ping Command Output from DUT1 to DUT8

```
[~M8-DUT8]tracert  srv6-te policy endpoint-ip 1::1 color 21 end-op 101::1010 -a 8::8
  Trace Route srv6-te policy : 100  data bytes, press CTRL_C to break
  srv6-te policy's segment list:
  Preference: 21; Path Type: primary; Protocol-Origin: local; Originator: 0, 0.0.0.0; Discriminator: 21; Segment-List ID: 65;
Xcindex: 65; end-op: 101::1010
  TTL    Replier                                  Time    Type      SRH
  0                                                       Ingress   [SRH: 107::107, 103::103, 101::101, 101::1010, SL=3]
  1      2000:7:8::1                              3 ms    Transit   [SRH: 107::107, 103::103, 101::101, 101::1010, SL=3]
  2      2000:5:7::1                              3 ms    Transit   [SRH: 107::107, 103::103, 101::101, 101::1010, SL=2]
  3      2000:3:5::1                              2 ms    Transit   [SRH: 107::107, 103::103, 101::101, 101::1010, SL=2]
  4      101::1010                                2 ms    Egress    [SRH: 107::107, 103::103, 101::101, 101::1010, SL=1]
```

Figure 3: Traceroute Command Output from DUT8 to DUT1

In this section, Huawei demonstrated a real functional capability to create an SRv6-TE tunnel that crosses SRv6 and IPv6-Only-capable domains. Also, We verified the readiness of the Huawei NetEngine 8000 platform to process SRv6 OAM END.OP SID. For future test campaigns, we recommend combining functional testing with performance benchmarking to assess the scalability of SRv6-TE tunnels creation.

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 DUT3, 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-BE tunnels, as shown in Figure 4. 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 DUT10-DUT4 as an SRv6 data plane for VPWS services.

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 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 5 shows the output of the Local-SID End.DT4 forwarding table for DUT3. Figure 6 shows the destination IPv6 address on the outer IPv6 header, which matches the VPN SID of the VPN ID 23.

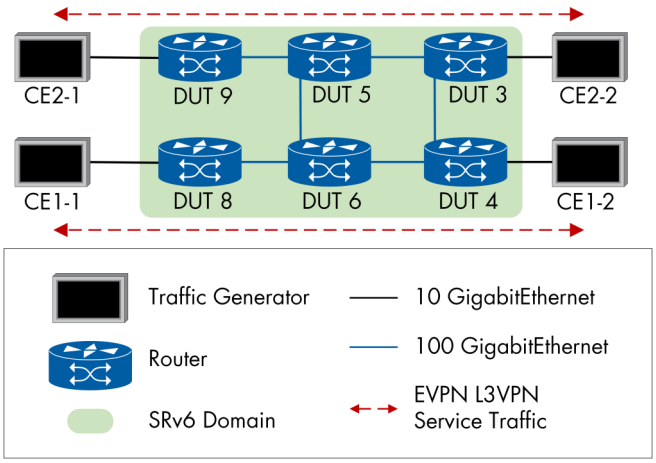

Figure 4: EVPN L3VPN over SRv6-BE

```
[~X4-DUT3]display segment-routing ipv6 local-sid end-dt4 forwarding

                  My Local-SID End.DT4 Forwarding Table
                  -------------------------------------

SID         : 103::1:0:3D/128                       FuncType : End.DT4
VPN Name    : 2100                                  VPN ID   : 22
LocatorName: as                                     LocatorID: 1

SID         : 103::1:0:3E/128                       FuncType : End.DT4
VPN Name    : 290                                   VPN ID   : 34
LocatorName: as                                     LocatorID: 1

SID         : 103::1:0:3F/128                       FuncType : End.DT4
VPN Name    : ifit                                  VPN ID   : 36
LocatorName: as                                     LocatorID: 1

SID         : 103::1:0:40/128                       FuncType : End.DT4
VPN Name    : 230                                   VPN ID   : 23
LocatorName: as                                     LocatorID: 1

SID         : 103::1:0:42/128                       FuncType : End.DT4
VPN Name    : 231                                   VPN ID   : 24
LocatorName: as                                     LocatorID: 1
```

Figure 5: L3VPN SID Allocation

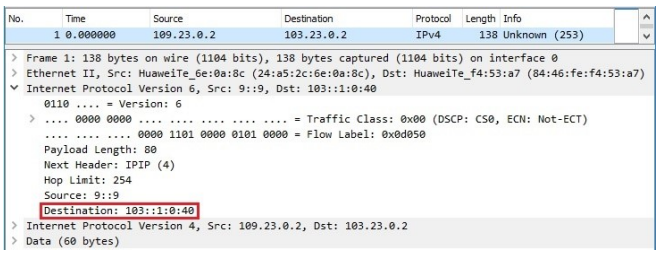

Figure 6: L3VPN Packet Capture

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 and ATN 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 performance of Huawei NetEngine 8000 and ATN 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 (TI-LFA) and SRv6 Traffic Engineering Fast Re-Route (SRv6-TE FRR) handle the link and transit node failures with restoration time less than 50 ms. Moreover, SRv6 Path Egress Protection affords an extra 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 Micro-loop-avoidance

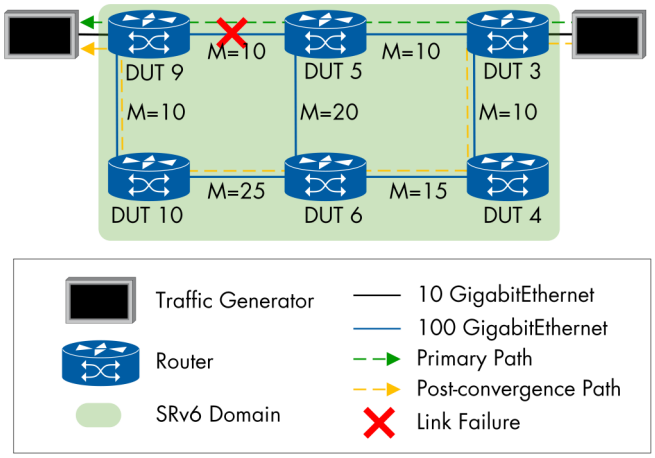

Figure 7: SRv6 TI-LFA Topology

Topology-Independent Loop-Free Alternate provides full convergence in less than 50 ms in case of link or node failure. TI-LFA prevents the micro-loops, which usually occur because of the suboptimal routing during the topology convergence. Segment routing can encode the post-convergence path in the segment-list.

This enforces the packet to follow a loop-free path and avoid the micro-loops. We started the test by setting the IS-IS link metrics, as shown in Figure 7. TI-LFA was enabled under the IS-IS IPv6 process. DUT3 was selected to be the ingress node, and DUT9 was the egress node for ten L3VPN instances. We send IPv4 bidirectional traffic by the rate of 100,000 frames per second between DUT3 and DUT9.

Then, we requested the Huawei test engineer to break the link physically between DUT5 and DUT9. After that action immediately, we observed frame loss in both directions of the traffic flow. To proof the results consistency, we repeated the test three times.

The maximum number of lost frames was 246 frames. This indicates that the out-of-service time was 2.46 ms. To check the restoration behavior, we asked the Huawei team to bring up the physical link again. Subsequent to the network convergence period, the optimal SRv6 tunnel was established and included the link between DUT5 and DUT9. The traffic was switched back without frame loss.

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 fast reroute. 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 8. 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. After repeating the test three times, the maximum number of lost frames was 1085 frame. 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 cross through DUT5. The traffic was restored with 0 frame loss.

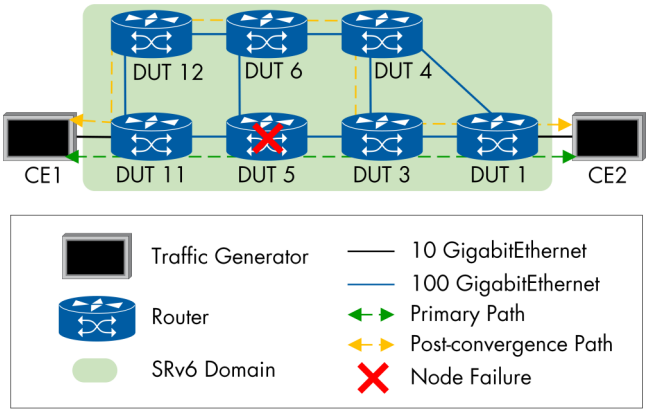

Figure 8: SRv6 TI-LFA FRR Topology

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".

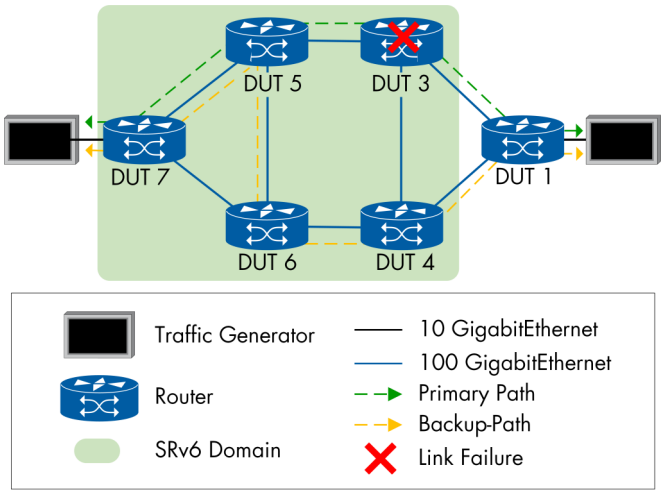

Figure 9: SRv6 Path Egress Protection Topology

Figure 9 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 DUT4. 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. After the reboot cycle completed, DUT3 brought up and the SRv6 tunnel was switched back through it. During the switching back or restoration period, we didn't observe any frame loss in the flowing traffic. 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 (30~50 ms). 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. On the other hand, 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 doesn't 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 DUT9 and DUT3 through DUT5. Huawei enabled the LDP on DUT9, DUT5, and DUT3 to allocate the transport labels. And Huawei configured MP-BGP VPNv4 to exchange the VPN labels between the PE's (DUT9 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 (DUT9 and DUT3), assign the node and VPN SIDs, establish IBGP session and enable VPNv4 neighbor by using the IPv6 address of the PE's (DUT9 and DUT3), creating new SRv6-BE tunnels over DUT9 and DUT3. After that, Huawei replaced the configura-

tion 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 DUT9 and keep only SRv6 configurations to make sure there are no configurations dependency still exists.

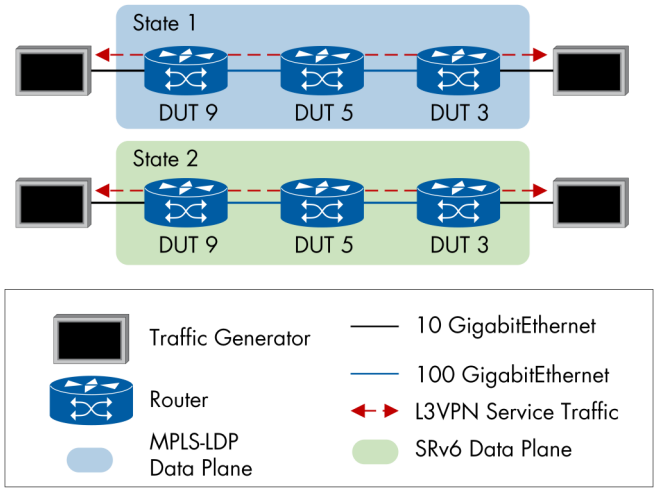

Figure 10: Migration from MPLS-LDP to SRV6 Data Plane

During the whole period of the test, we didn't observe any packet loss or service interruption. We verified the non-disruptive behavior of the protocol migration from MPLS LDP to SRv6.

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

In the second scenario, we used an RSVP-TE tunnel instead of an LDP tunnel for the same L3VPN service that we utilized in scenario 1. 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 DUT3 through DUT5. Huawei enabled the RSVP-TE tunnel on DUT9, DUT5, and DUT3 to allocate the transport labels. And Huawei configured MP-BGP VPNv4 to exchange the VPN labels between the PE's (DUT9 and DUT3). 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 stage, Huawei configured IS-IS IPv6 on all the routers, enabling SRv6 on the ingress and egress routers (DUT9 and DUT3), assign the node and VPN SIDs, establish IBGP session and enable VPNv4 neighbor by using the IPv6 address of the PE's (DUT9 and DUT3), creating new SRv6-TE tunnels over an explicit path (DUT9-DUT5-DUT3). 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 DUT3, DUT5, and DUT9 and keep only SRv6 configurations to make sure there are no configurations dependency still exists.

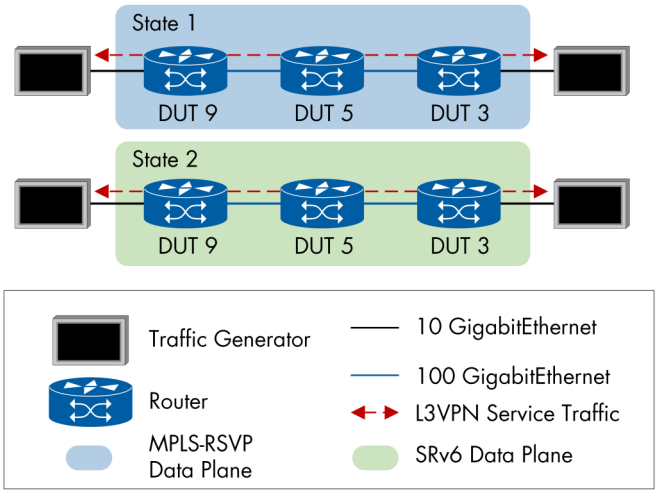

Figure 11: Migration from MPLS-RSVP to SRV6 Data Plane

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.

EANTC verified the support of Huawei NetEngine 8000 and ATN platforms to both MPLS and SRv6 data planes. Huawei proved the migration from the classical MPLS data plane to the SRv6 data plane without any service interruption or packet loss. Natively, the migration from LDP to SRv6 requires the ingress and egress routers to be SRv6-aware. The midpoint routers can be IPv6-capable only.

SRv6 Service Function Chaining

Service Function Chaining (SFC) is commonly used to describe a series of connected network functions that accomplish an end-to-end network service. The service function chain can run on a physical appliance or a virtual instance or a combination of both. Optimally, SFC path instructions should be programmed by the ingress node of the chain.

SRv6 steers the packet in the network based on the source routing paradigm natively. Conceptually, source routing enables the ingress node to encode an explicit path which the packet will follow to reach the destination node. The SRH lists the SIDs of all the SRv6-aware nodes in the path. The node can be a simple transit forwarding node (router) or a node that processes the packet based on a particular network service function (SF). If the SF can't process or handle the SRH, then the service is described as an SRv6-unaware service. In case an SRv6-unaware SF is inserted in an end-to-end SFC, the SFC proxy component is needed. SFC proxy is a logical element that removes and inserts SRH on behalf of an SFC-unaware service function.

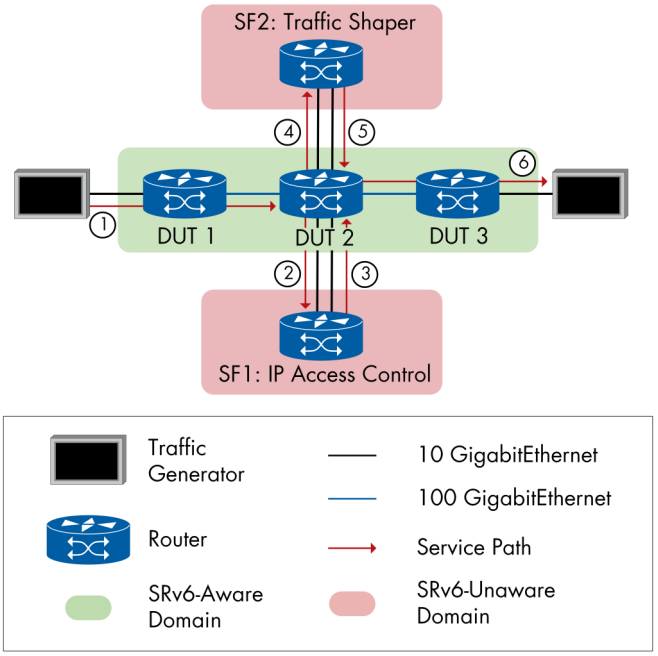

Figure 12: SFC using SRv6 Test Topology

For SFC testing, the Huawei testing team build a new testbed that was configured to demonstrate the capabilities of Huawei software and hardware to create an SFC using SRv6 data plane, which includes two SRv6-unaware SFs.

Figure 12 depicts, Huawei configured the intended SFC to span through DUT1 (ingress node), DUT3 (egress node), and DUT2 (SFC Proxy). Huawei configured on DUT2 (SFC Proxy) with the proxy SID. Once the packet reached DUT2 with the Proxy SID, DUT2 striped the outer IPv6 header and sent out the original IPv4 payload to the SRv6-unaware SF1. Huawei configured SF1 router as a simple IP access control function to block the incoming traffic based on a specific source and destination IPv4 address. DUT2 received back the allowed IPv4 traffic by SF1 router. Then DUT2 forwarded the IPv4 payload to the second function. SF2 was inserted to the service function chain to apply bandwidth limiting/shaping policy on the incoming traffic, and reroute the traffic back to DUT2. After that, DUT2 inserted a new IPv6 outer header along with a new SR list and forwarded it to DUT3.

EANTC verified the traffic flow of this SFC and the operations on the traffic by the SFC proxy (DUT2), SF1 and SF2. Also, we tested the behavior of the static SR proxy as defined in the Internet-Draft draft-xuclad-spring-sr-service-chaining. Because of the SFC proxy function is a compute and memory-intensive operation, we recommend to test the scalability and capacity of Huawei NetEngine 8000 platform than can be handled.

Next Generation Multicast Services

MVPN over BIERv4

Bit Index Explicit Replication for IPv4 (BIERv4) is an efficient protocol to forward IPv4 multicast traffic without engaging the intermediate routers in the tree-building process of the multicast topology. Moreover, the states of the multicast flows are not required to be maintained on the intermediate routers. BIER is designed to build a stateless forwarding for the multicast traffic.

This test aims to verify the functional capabilities of the DUTs to support Multicast VPN (MVPN) service using BIERv4 and MPLS data plane. Huawei prepared the testbed by the following configurations and procedures:

- Configuring MPLS data plane in the DUTs of Figure 13, and an L3VPN instance in the DUT3, DUT7, and DUT9
- Configuring BIER and enable BIER in IS-IS on all the DUTs
- Establishing BGP MVPN peer relationships between Root BFIR (DUT3) and the Leaves BFER (DUT7 and DUT9), and configuring BGP to advertise A-D and C-multicast routes
- Configure DUT3 (Root), DUT7 and DUT9 (Leaves) to transmit multicast traffic over BIER tunnels

- Enabling PIM on the interfaces between the DUTs (root and leaves) and the traffic generator (customer edge router) to allow a VPN multicast routing table to be established to guide multicast traffic forwarding

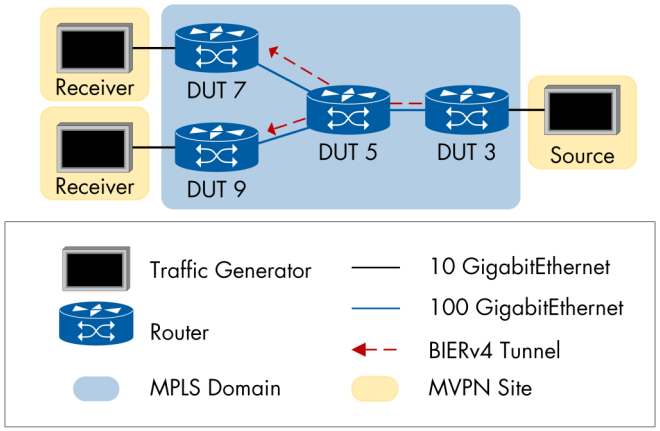

Figure 13: MVPN over BIERv4 Test Topology

After that, we asked Huawei to configure one IPv4 multicast group in the testbed. Spirent TG emulated the source and receivers of the multicast traffic for this group. We verified the BFR-ID advertisement through IS-IS, BIER IPMSI status, and the auto-discovery and C-multicast routes by from each leaf by MVPN BGP session. After the control plane verification, we generated multicast traffic from DUT3 (BFIR) toward the leaves DUT7 & DUT9 (BFER).

We verified the multicast stream was delivered to each leaf identically and without any frame loss. Also, we didn't observe any routing loops in the core network by checking the links utilization continuously.

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 the 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 a full load.

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

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 1970.4W. 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 3100.3W. Which means 1129.9W was consumed by the line card 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 TestCenter) with the DUT (Huawei NetEngine 8000 X8 8x400GE+8x100GE Line Card). Table 3 enumerates the achieved throughput and latency results per frame size.

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 TestCenter (STC) 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 STC 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

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

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 FIB. We established an MP-BGP session between the 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. The Spirent advertised 4 million IPv4 prefixes to the DUT.

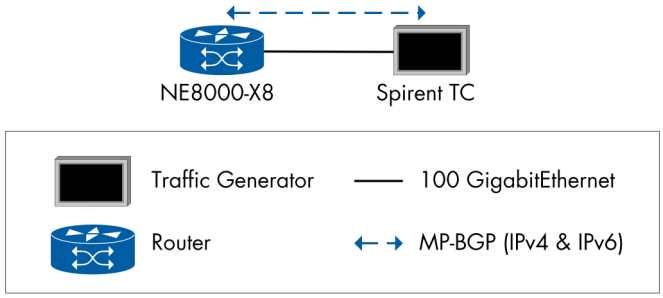

Figure 14: FIB Scalability & BGP Routes Learning Rate Test Topology

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

Table 5: BGP Routes Learning Rate

Simultaneously, the STC 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.

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 & IPv6 traffic to detect any failure or abnormal performance. Table 6 summarizes the achieved results.

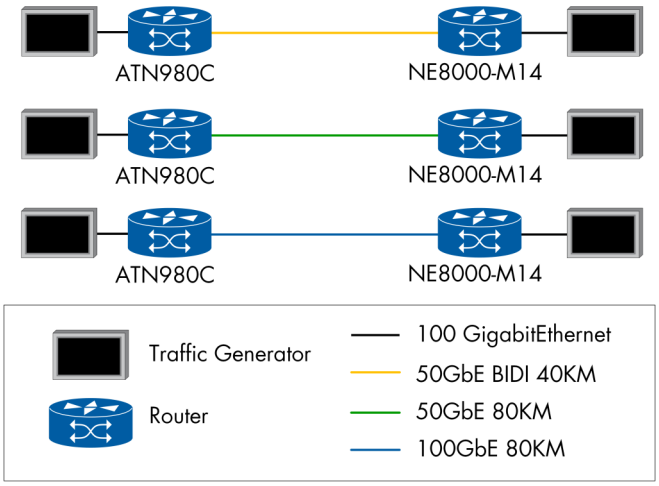

Figure 15: Different Optical Module Types Test Connectivity

	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

Network Cloud Engine

The second 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). NCE version was (V100R019C00SPC600). The Huawei testing team informed us this version is a pre-commercial or beta software version. Table 7 lists the details of the testbed setup of the NCE.

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 Testbed Components and Software Versions

For the NCE (IP Domain) testing, we focused on four areas to demonstrate and verify the capabilities of NCE to:

1. Calculate SRv6 policy paths based on different constraints
2. Provision L3VPN service over SRv6 data plane
3. Manage 5G network slices
4. Applying next-generation OAM techniques using iFIT Monitoring

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 16 shows the discovered physical topology of 12 routers and the physical links between them.

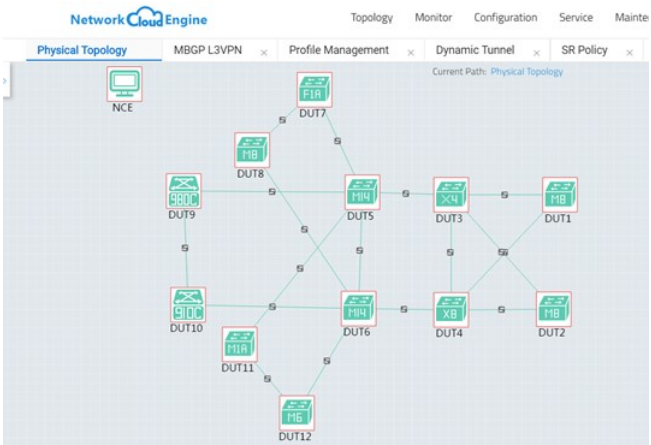

Figure 16: Testbed Physical Topology

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 17. 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.

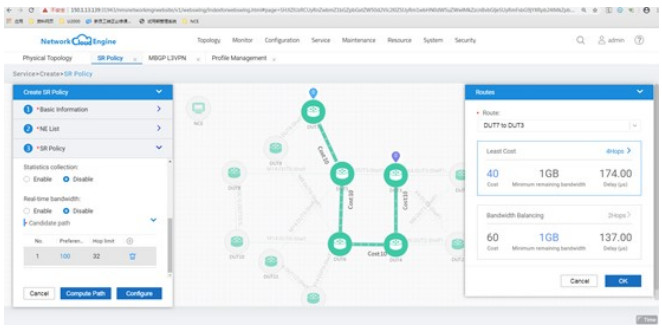

Figure 17: SRv6 Path Calculation based on the IGP Cost

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, Two-Way Active Measurement Protocol (TWAMP) 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).

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 18 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).

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 19.

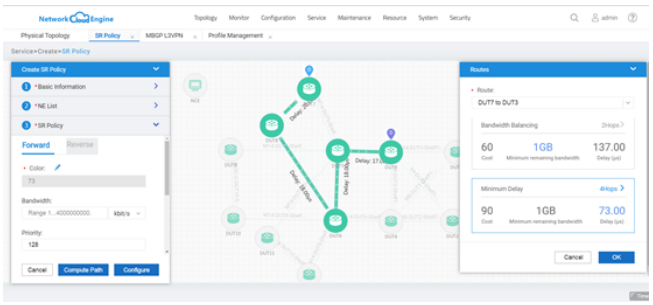

Figure 18: SRv6 Path Calculation based on the Minimum Delay

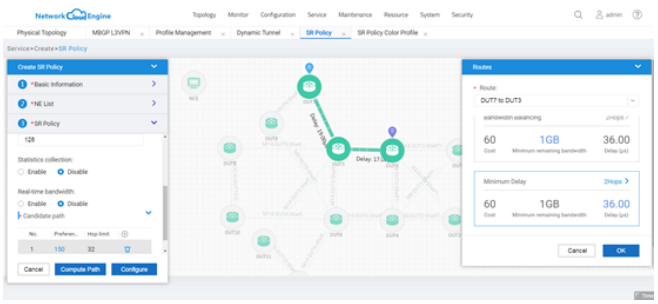

Figure 19: SRv6 Path Re-optimization based on the Minimum Delay

SRv6 Path Calculation based on Bandwidth Utilization

It's 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. 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 as depicted in Figure 21.

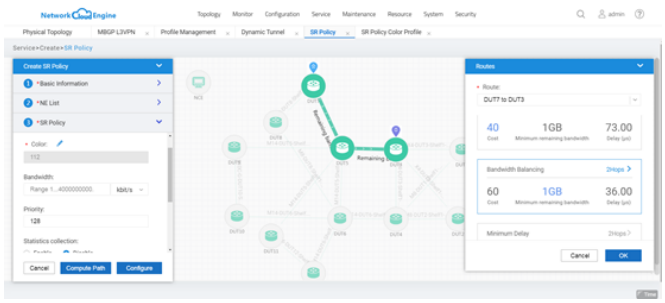

Figure 20: SRv6 Path Calculation based on the Available Bandwidth

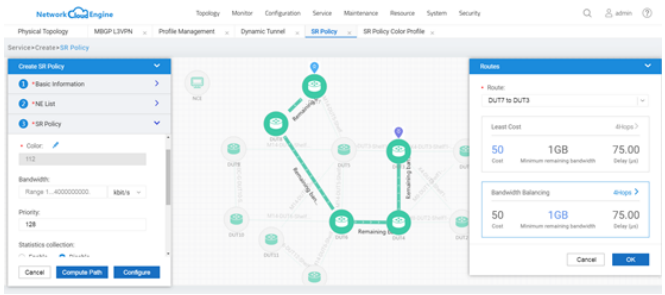

Figure 21: SRv6 path re-optimization based on the available bandwidth

SRv6 Path Calculation based on Explicit Path

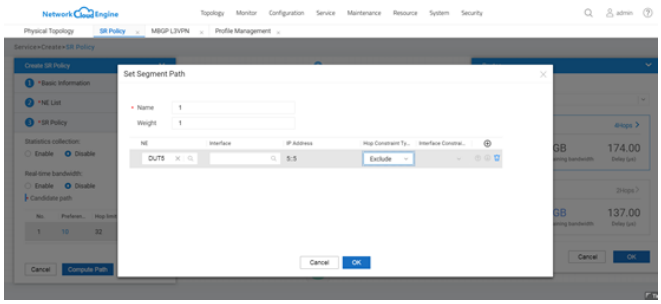

Figure 22: Defining SRv6-TE Explicit Path Policy

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 22 and Figure 23.

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).

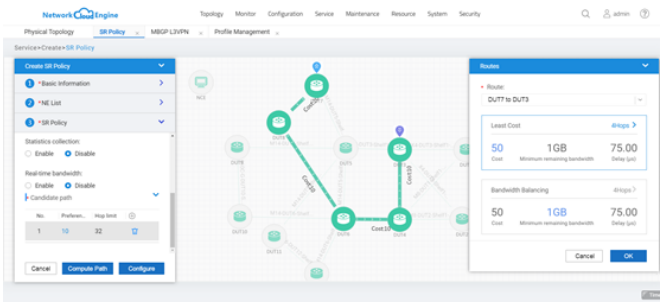

Figure 23: SRv6 Path Setup based on the Explicit Path

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, 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 select tunnel associated between the PE's. After completing the configuring of all the required parameters, the Huawei test engineer applied the service template, and we verified the new configurations on DUT3 and DUT9.

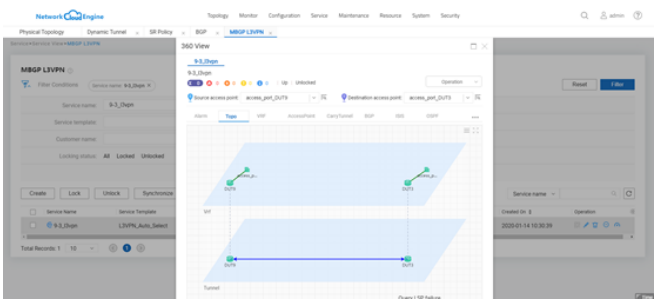

Figure 24: Logical Topology of L3VPN Service

To confirm the traffic flow between the CE sites, we generated IPv4 and IPv6 traffic between the emulated CE sites, as shown in Figure 24. The traffic was flowing between the sites as expected without any frame loss.

5G Network Slicing Management

The Capabilities of NCE (IP Domain) for Transport Network Slicing

5G Network slicing is a key network architecture that enables multiple 5G services with distinct requirements over the same network infrastructure (Radio Access Network, Transport Network, and the Mobile Core Network). The deployment of the network slicing concept in the data plane of the transport network has two flavors; the first is "Soft Slicing", which accomplishes the logical isolation between the services by the logical networking concepts like VPN or H-QoS. The second flavor is "Hard Slicing" which dedicates a specific physical bandwidth or networking resource (i.e., lambda or TDM channel) for each network slice (5G service).

However, provisioning the network's physical resources and managing the life cycle of a network slice is not a piece of cake operational task for the network operators. This requires a kind of SDN controller that can match the service SLA with the available network resources. To fill this gap, Huawei developed a new module in NCE (IP Domain) called "Network Slice".

Network Condition	Value
Physical Bandwidth Isolation Technology	VLAN Channelized Sub-Interfaces
Huawei platforms that NCE (IP Domain) will manage	NetEngine 8000, ATN910C, ATN980C
Data Plane Encapsulation for Network Slicing	MPLS-SR
Service Type	L3VPN

Table 8: Network Slicing Deployment Conditions

According to Huawei documents, NCE (IP Domain) can support multiple physical bandwidth isolation technologies, managing multiple routing platforms and provisioning network slicing over different data plane encapsulations. To generate reproducible test results, EANTC and Huawei testing team fix the network conditions to check the capabilities of NCE (IP Domain) to support network slicing as listed in Table 8.

Network Slice Template and Network Traffic Isolation

We started the test by defining the slice templates of three network slices with the following requirements:

Network Slice Name	Type	SLA	DSCP
National Power Slice	URLLC	BW: 1Gbps, Latency: 1 ms	10
CloudVR Game Slice	eMMB	BW: 3Gbps, Latency: 30ms	20

Table 9: Definitions of The Network Slices

After applying the three network slices, we checked the changes in the network topology. Initially, we observed enabling the "VLAN Channelized Sub-Interfaces" mode in all the supported physical links. Each physical link has two new sub-interfaces with specific bandwidth for each, as described in Table 9. Based on the current capabilities of the NCE (IP Domain) controller, Huawei created MBPG L3VPN services between two sites for each slice. The L3VPN services were transported via the MPLS-SR data plane. To confirm the physical isolation between the slices, we started generating IPv4 traffic in the rate of (100 Mbps) for the national power slice with DSCP value (10), and (500 Mbps) for the CloudVR game slice with DSCP value (20). Each traffic flow steered dynamically into a related tunnel of the network slice based on the DSCP value. We didn't observe any frame loss in any slice. Also, we checked the maximum latency in each slice.

Then we increased the generated traffic throughput for CloudVR game slice up to (5 Gbps). Because the generated traffic (5 Gbps) is higher than the allocated bandwidth of the CloudVR game slice, we observed around 40% frame loss. Also, the max latency was extremely increased because of the expected traffic queuing. However, all these changes in the CloudVR game slice don't cause any impact on the traffic of national power slice.

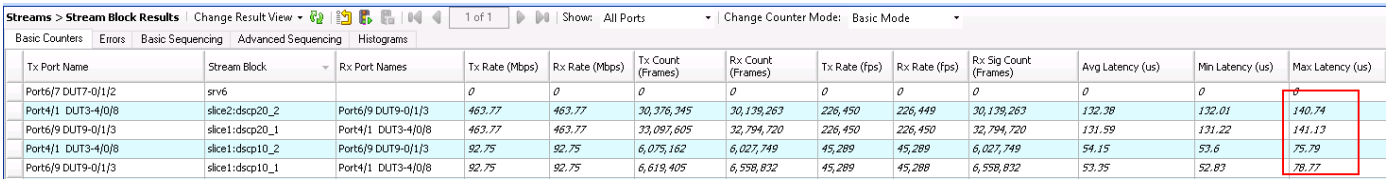

Tx Port Name	Stream Block	Rx Port Names	Tx Rate (Mbps)	Rx Rate (Mbps)	Tx Count (Frames)	Rx Count (Frames)	Tx Rate (fps)	Rx Rate (fps)	Rx Sig Count (Frames)	Avg Latency (us)	Min Latency (us)	Max Latency (us)
Port6/7 DUT7-0/1/2	srv6		0	0	0	0	0	0	0	0	0	0
Port4/1 DUT3-4/0/8	slice2:dscp20_2	Port6/9 DUT9-0/1/3	463.77	463.77	30,376,345	30,139,263	226,450	226,449	30,139,263	132.38	132.01	140.74
Port6/9 DUT9-0/1/3	slice1:dscp20_1	Port4/1 DUT3-4/0/8	463.77	463.77	33,097,605	32,794,720	226,450	226,450	32,794,720	131.59	131.22	141.13
Port4/1 DUT3-4/0/8	slice1:dscp10_2	Port6/9 DUT9-0/1/3	92.75	92.75	6,075,162	6,027,749	45,289	45,289	6,027,749	54.15	53.6	75.79
Port6/9 DUT9-0/1/3	slice1:dscp10_1	Port4/1 DUT3-4/0/8	92.75	92.75	6,619,405	6,558,832	45,289	45,288	6,558,832	53.35	52.83	78.77

Figure 25: Generated Traffic Streams Per Network Slice - Full Load Throughput

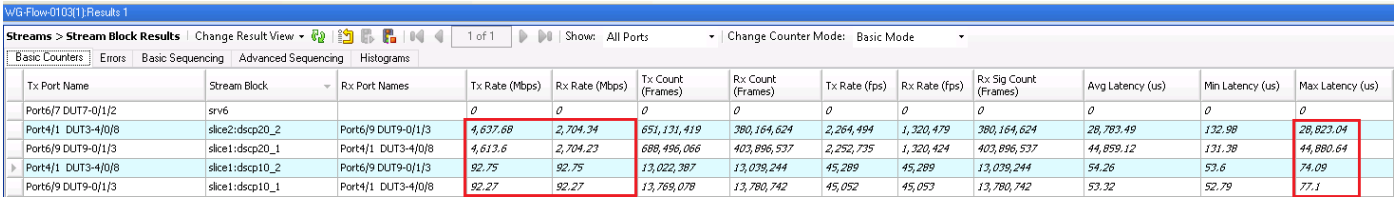

Tx Port Name	Stream Block	Rx Port Names	Tx Rate (Mbps)	Rx Rate (Mbps)	Tx Count (Frames)	Rx Count (Frames)	Tx Rate (fps)	Rx Rate (fps)	Rx Sig Count (Frames)	Avg Latency (us)	Min Latency (us)	Max Latency (us)
Port6/7 DUT7-0/1/2	srv6		0	0	0	0	0	0	0	0	0	0
Port4/1 DUT3-4/0/8	slice2:dscp20_2	Port6/9 DUT9-0/1/3	4,637.68	2,704.34	651,131,419	380,164,624	2,264,494	1,320,479	380,164,624	28,783.49	132.98	28,823.04
Port6/9 DUT9-0/1/3	slice1:dscp20_1	Port4/1 DUT3-4/0/8	4,613.6	2,704.23	688,496,066	403,896,537	2,252,735	1,320,424	403,896,537	44,859.12	131.38	44,880.64
Port4/1 DUT3-4/0/8	slice1:dscp10_2	Port6/9 DUT9-0/1/3	92.75	92.75	13,022,387	13,039,244	45,289	45,289	13,039,244	54.26	53.6	74.09
Port6/9 DUT9-0/1/3	slice1:dscp10_1	Port4/1 DUT3-4/0/8	92.27	92.27	13,769,078	13,780,742	45,052	45,053	13,780,742	53.32	52.79	77.1

Figure 26: Generated Traffic Streams Per Network Slice - Overload Throughput

Non-disruptive Network Slice Life Cycle Management

The goal of this test case is to demonstrate the capability of NCE to create a new network slice or to delete an existing network slice in a live environment. We asked the Huawei team to create a third network slice "Test" based on a different network slice template. And to terminate this network slice after a while. During that time, we were generating traffic in all the existing slices. We didn't observe any frame loss or impact on the traffic latency. We confirmed the capability of the NCE (IP Domain) controller to create and terminate network slice without service impact on the neighbor slices.

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 team 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. 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.

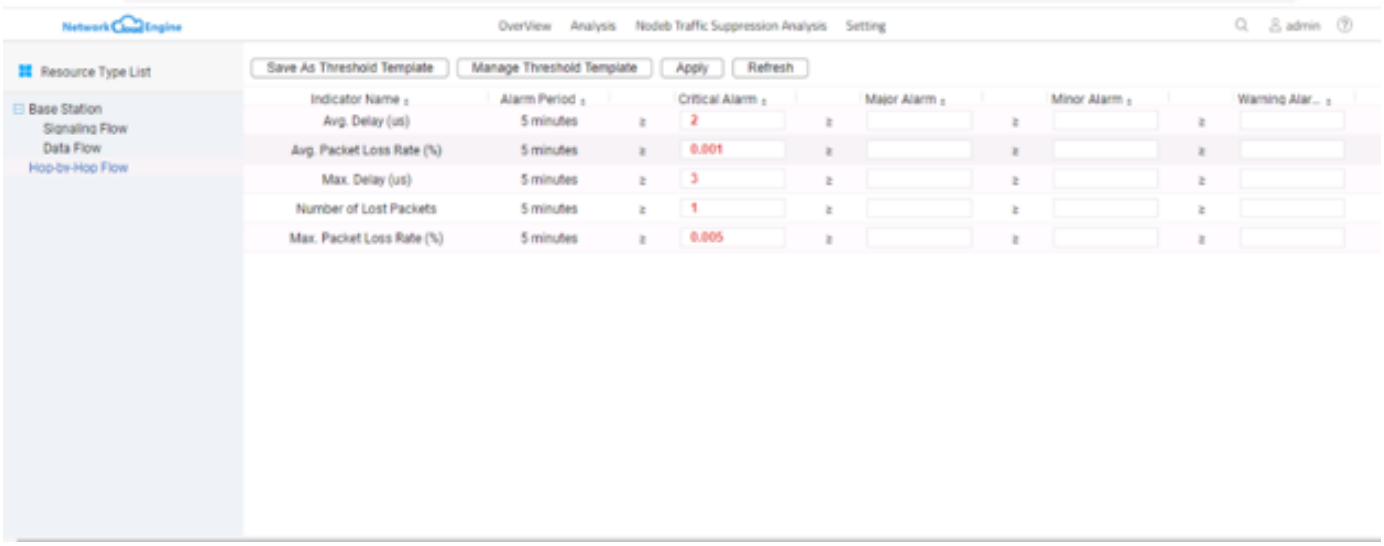

Figure 27: Configuring Hop-by-Hop Flow Analysis

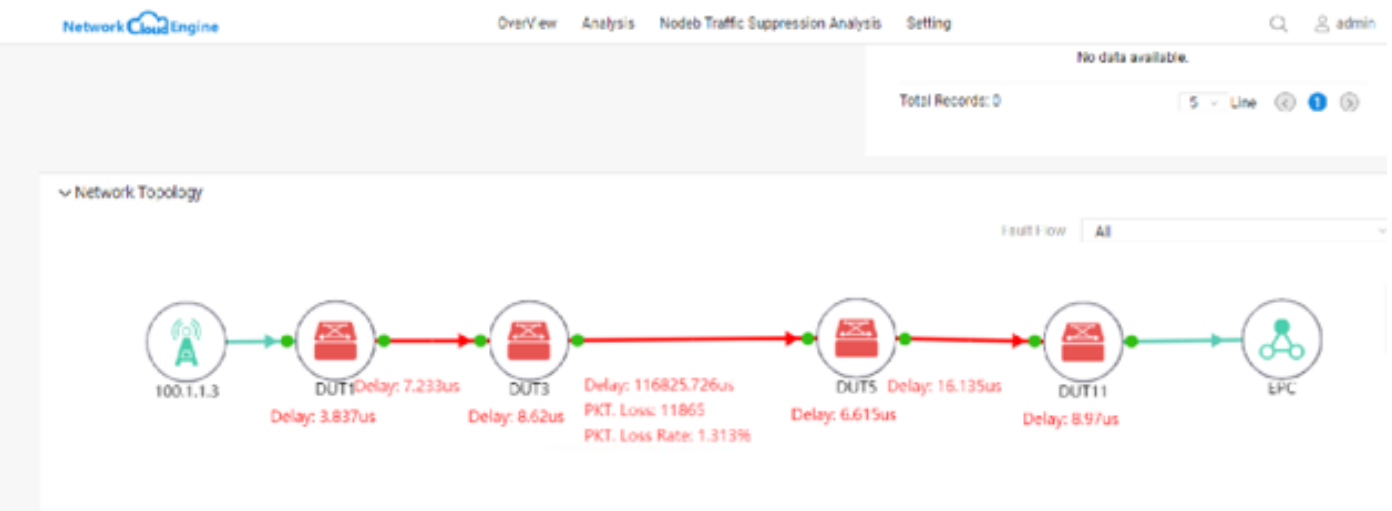

Figure 28: Result of Hop-by-Hop Flow Analysis

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 28. 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).

Packet Network Clock Synchronization

The third chapter of this test focuses on clock synchronization support of the Huawei NetEngine 8000 and ATN910C routers under test.

Introduction

Since the advent of mobile network services, cell sites need to be provided with synchronized time. While base stations were mostly equipped with GNSS (including GPS, GLONASS, GALILEO) satellite receivers, in the beginning, this has shown to be impractical in many cases. On one hand, GNSS receivers are expensive; on the other hand, a view of the sky is not always easily available. Thus, the Precision Timing Protocol (PTP), specified in IEEE-1588:2008, has been the dominant solution to provide synchronized timing to cell sites for several years. A GNSS receiver is connected to a Grandmaster Clock (GM) which initiates PTP connections. On the way through the network, Telecom Boundary Clocks (T-BC) are located at intermediate routers to distribute the timing information downstream. Slave Clocks (T-SC) are located at cell site gateways and terminate the PTP sessions.

Typically, most or in some cases even all intermediate routers need to implement T-BC functions for LTE or 5G scenarios. The very precise timing information required at cell sites can only be achieved if boundary clocks are run at every hop in the packet network.

There are three types of timing information: Frequency, phase, and time of day. Frequency sync is the most basic, required for all applications to ensure that the frequency (the speed of the clock) is maintained correctly. Phase information helps to align the phase offset in addition to the frequency – synchronizing for the correct phase requires much more precision. Finally, time of day (TOD) synchronization is closely linked to phase sync; it aligns the actual time of the day with the reference by sharing a precise indication of the beginning of each second.

Test Scope

Huawei asked us to verify the PTP implementations of the three main router types involved in the test, specifically the NetEngine 8000 X8 and M14 and the ATN910C. At EANTC, we have conducted many tests with previous generation NetEngine routers – specifically the NE40E – as part of our ongoing series of multi-vendor interoperability tests over the years (Huawei 5G-Ready SDN: Evaluation of the Transport Network Solution). However, we had not evaluated the brand new NetEngine 8000 yet. For each of the three routers, we validated the T-BC implementation with three test cases:

- Performance and precision over the long term
- Stability of operations when confronted with noisy input signals from the GM
- Correct detection of degraded primary GM service with subsequent failover to a secondary GM

These are three of the most important basic T-BC clock quality tests. Of course, many additional tests could be executed to evaluate the boundary clock's features and performance. The NetEngine 8000 family participated in the EANTC multi-vendor interop event in March 2020.

Huawei positions the PTP clocks to match "Class C", one of four new classes recently introduced in ITU-T G.8273.2 describing clock performance requirements at the end application. The ITU-T standard describes performance requirements in detail, and we will refer to it in the following subsections. Class C defines the most strict performance requirements known today, as Class D requirements are still under study.

Single-Node Boundary Clock Performance Test

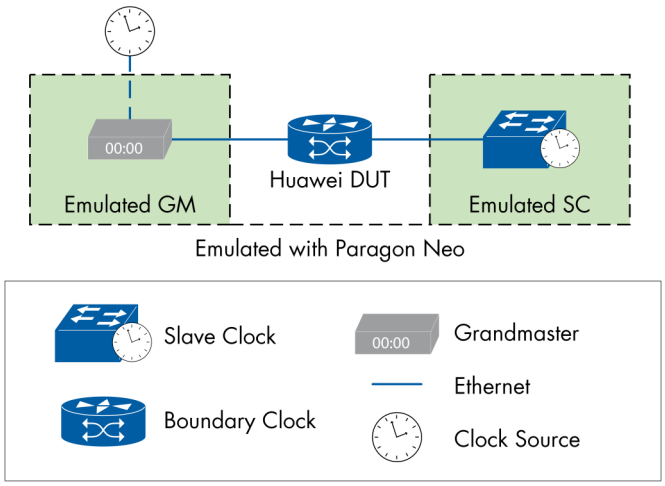

Figure 29: Single-Node Boundary Clock Performance Testbed

In this first clocking test, we baselined the boundary clock quality under normal, locked operating conditions when an ideal input reference packet timing signal was present.

G.8273.2 states that a Class C clock must have a maximum absolute time error, or max|TE| less than plus/minus 30 nanoseconds (max|TE| < 30 ns). Huawei suggested to tighten this requirement further and promised that the NetEngine 8000 T-BC would be in line with the very strict precision of max|TE| < 10 ns. Although today's foreseeable 5G scenarios require an end-to-end clock precision of 30 ns only, Huawei suggested testing for these values to be prepared for future deployment scenarios.

We connected each of the three routers under test to a reference grandmaster clock and a reference slave clock, implemented by a Calnex Paragon-Neo emulator. The test tool, in turn, was connected to a GNSS clock source. This setup is preferred because it eliminates any kind of interoperability issues and allows complete control of the testbed.

Once connected, we ran a long-term clock stability test over 12 hours. The emulated grandmaster input clock provided an optimal PTP signal. The boundary clock's ability to provide a likewise optimal PTP signal and a Synchronous Ethernet (SyncE) output downstream was constantly checked by the Calnex emulator. In accordance with G.8273.2, all results were unfiltered.

Results: We verified that each of the three routers matched Huawei's claims:

- NetEngine 8000 X8: max|TE| = 5.7 ns over 12 hours connected via 100GE
- NetEngine 8000 M14: max|TE| = 6.8 ns over 12 hours connected via 100GE
- ATN910C: max|TE| = 5.2 ns over 12 hours connected via 10GE

All three routers matched and exceeded the G.8273.2 requirements for Class C clocks by far.

Single-Node Boundary Clock Tolerance Test

A reference test with perfect clock inputs, as shown above, is certainly a good baseline to qualify the implementation under perfect lab conditions. In reality, a telecom boundary clock faces imperfect input signals. For this reason, a separate test is usually carried out verifying the tolerance of the boundary clock to legitimate, but imperfect input clock signals.

We conducted this test with all three boundary clocks as before. The Grandmaster and Telecom Slave Clock were emulated by a Calnex Paragon-Neo as before. Over a period of 20 minutes, the Calnex Paragon-Neo manipulated its PTP input signal at the Grandmaster to show various patterns of time error. Figure 30 depicts that.

The only task for the boundary clock is to stay synchronized. There are no specific performance requirements, as the massive input time error cannot be expected to be fully compensated. Thus, the downstream PTP signal of the boundary clock cannot be used to analyze the pass/fail criteria for this test. Instead, it is necessary to capture all logs of the router under test and to check for messages that might show the system is no longer locked.

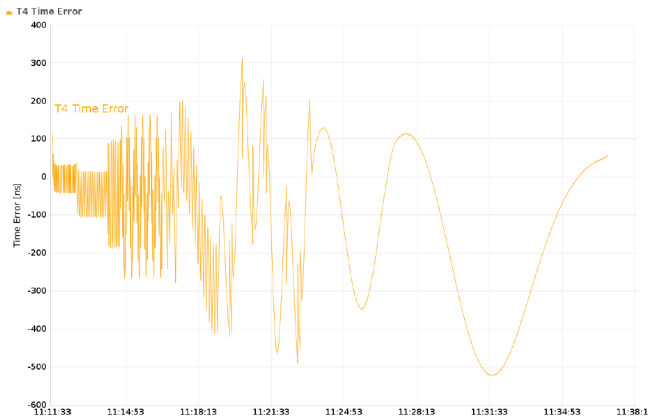

Figure 30: Time Error Pattern

Results: All three routers under test, the NetEngine 8000 X8, the NetEngine 8000 M14, and the ATN910C, remained locked to the Grandmaster clock signal all the time despite the (legitimate but hard to follow) variations of the PTP input signal.

The three systems mastered the maximum input tolerance with different frequency patterns, as required by the ITU standards for Class C devices.

Grandmaster Clock Source Failover

In this last of three baseline clocking tests, we verified the correct failover between two Grandmaster clocks. Usually, a T-BC is connected to multiple upstream GMs for redundancy. It selects the best suitable GM as deemed feasible by the T-BC. In case this, the GMs input signal is degraded – as indicated by PTP messages sent from the GM – the boundary clock must failover to the backup GM. During this time and after the failover, it needs to maintain output clock quality as before. The maximum time error (see first clock test above) must remain within the boundaries of G.8273.2, in this case, max|TE| < 30 ns.

We conducted this test with all three routers under test, using a slightly different setup as shown below. In each case, we stimulated a failover by disconnecting GNSS antenna from the primary GM, causing it to signal clock degradation downstream. The T-BC was expected to failover to the backup GM. After a few minutes, we reconnected the GNSS antenna to the primary GM. Subsequently, its clock quality got restored and the boundary clock was expected to switch back to the primary GM. We repeated each test scenario three times and compared values, reviewing them for consistency.

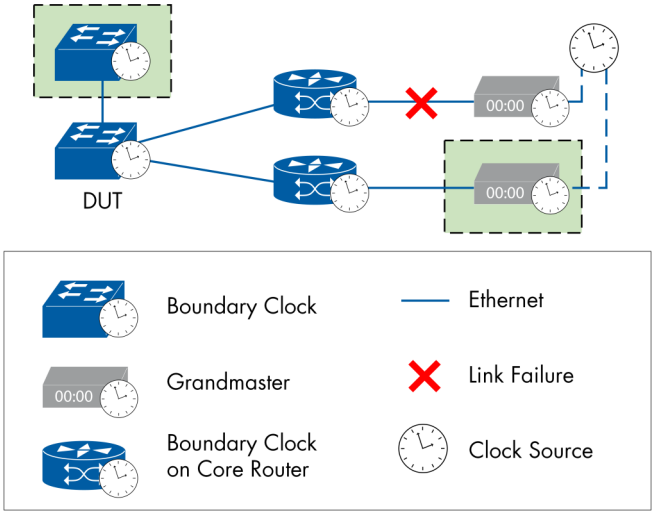

Figure 31: Grandmaster Clock Source Failover Topology

Results: All three routers under test mastered the challenge, maintaining different maximum time error values. Looking at the details, the NetEngine 8000 X8 maintained a max|TE|< 9 ns in the worst of three measurements. The primary and backup GM were not perfectly calibrated with each other in the lab, due to different cable lengths. This results in some offset between primary and backup GM, but does not affect the results otherwise.

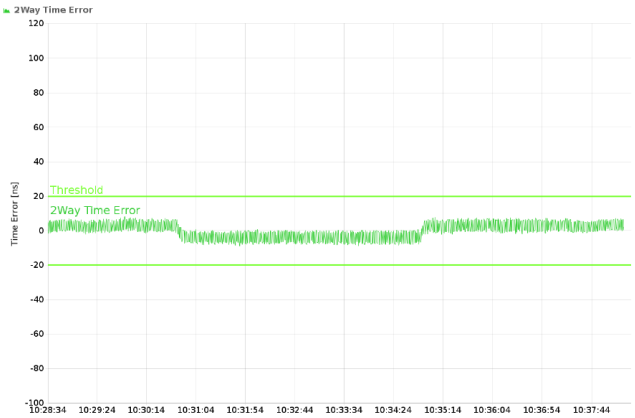

Figure 32: 2Way Time Error of NetEngine 8000 X8

The software version of NetEngine 8000 M14 was still in an early stage during the test. Huawei updated the software during the test process. We ran the test based on the official software release V800R012C00. Across all test iterations, the worst M14 performance was confirmed with max|TE| < 5.96 ns as indicated by the diagram below.

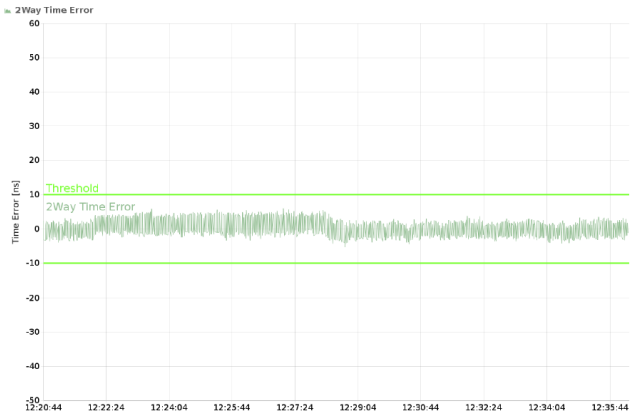

Figure 33: 2Way Time Error of NetEngine 8000 M14

Finally, the ATN910C router exhibited smooth values with max|TE| < of 12.3 ns in the worst observed case during our test.

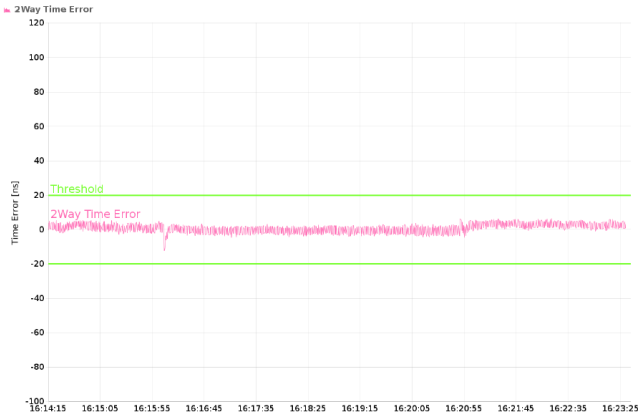

Figure 34: 2Way Time Error of ATN910C

The results of our baseline clocking tests confirm that the Huawei NetEngine 8000 X4, M14, and the ATN910C routers successfully meet and exceed Class C clock requirements in terms of long-term stability, noise tolerance, and GM failover quality.

Conclusion

Huawei has provided the family of NetEngine 8000 and ATN routers including new line cards and 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, 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.

About EANTC

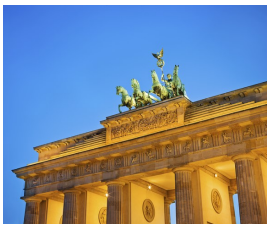

EANTC (European Advanced Networking Test Center) is internationally recognized as one of the world's leading independent test centers for telecommunication technologies. Based in Berlin, the company offers vendor-neutral consultancy and realistic, reproducible high-quality testing services since 1991.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