EANTC Independent Test Report

Huawei 5G-Ready SDN

Evaluation of the Transport Network Solution

September 2018

Introduction

Huawei commissioned EANTC to conduct an independent test of 5G bearer network features. The vendor selected a range of functional aspects from its product portfolio for the verification. EANTC was invited to review the solution at Huawei global headquarters in Shenzhen, China, in August 2018. There, engineers executed the tests in three different test beds, each focusing on specific aspects of the transport network.

5G is certainly approaching fast; major challenges for the transport network are going to hit operators in 2–3 years when large-scale services will be deployed:

1. Service scalability: Many more cell sites than before will be connected to the network – the numbers are likely to increase by a factor of five to ten, thus increasing the number of paths across the transport network and requiring better protocol scale, automated provisioning and fault management.

2. Bandwidth scalability: For many consumers, the promise of 5G services is mostly regarding throughput. AR/VR, streaming at HD and soon 4K video quality will require massive transport capacity, starting at the cell site connection.

3. Mobile edge computing (MEC): Some cell sites will require service connectivity to one or more MEC sites closer to the edge in addition to the default mobile core connections; this will increase the number of paths in the transport network further.

4. Slicing: There will be differentiated transport network services for a range of services – some requiring low delay, others high availability, again others massive throughput. The bearer network will have to support these services through traffic engineering, a concept that had often been proposed in the past but not widely adopted yet, and monitor the service quality with advanced telemetry solutions.

5. Transport migration: Service providers have deployed vast backhaul infrastructures for the pre-existent LTE/4G networks; these need to be maintained, integrated and migrated to SDN for cost and deployment speed reasons.

Huawei demonstrated a selected set of functional configurations hinting towards these challenges. A 360-degree evaluation of the whole transport network was left for future exercise – specifically regarding the end-to-end transport network integration, real traffic engineering, performance evaluation and service scalability testing. In any case, the functional tests we witnessed highlight a number of key concepts that service providers can explore further.

Executive Summary

The ATN950C and NE40E-M2K passed interface forwarding performance tests of the 25GE and 50GE cards, including 40km dual-fiber 50GigE SFPs and single-fiber BiDi 50GigE SFPs; the 50GE module showed slicing isolation and 1 Gbit/s bandwidth granularity of channelized sub-interfaces.

The two routers also demonstrated a range of Segment Routing functionality, including live migration from LDP or RSVP-TE to Segment Routing, interconnection to MPLS-TP, and interconnection between MP-BGP L3VPN and EVPN-MPLS. The hardware showed support of up to ten stacked labels in segment routing. The out-of-service time at hot-standby protection remained below 50 ms.

The Huawei Network Cloud Engine (NCE) demonstrated its ability to compute an optimal path based on multiple constraints, to provision L3VPN and L2 EVPN services by automatically deploying Segment Routing tunnels and to simplify the maintenance with automation and what-if simulation. NCE showed also streaming Telemetry capability and path optimization using REST-API.

Test Highlights

→ Forwarding Performance of 50GigE and 25GigE module achieved 99.4% link utilization
→ Latency below 15 µs per hop
→ 50GigE module supports 40km dual-fiber SFPs with up to 20 dBm optical budget, as well as single-strand bi-directional SFPs

→ In-service slicing using 50GigE module
→ Hitless bandwidth resizing
→ 1 Gbit/s slice granularity of the channelized interface
→ Proven congestion isolation between slices

→ Less than 41.5ms failover time with SR-TE static path protection for L3VPN services and zero frame loss during path restoration

→ Proven NCE ability to perform path computation of EVPN/Segment Routing using 9 constraints
→ Proven NCE ability to provision L3VPN and L2 EVPN services, to create bandwidth on demand per service and to simplify maintenance using what-if simulation
→ Streaming Telemetry using gRPC and path optimization with RESTful API

The Huawei ATN950C and NE40E-M2K routers, and the Network Cloud Engine (NCE) can be used in many contexts for fixed and mobile services. When Huawei commissioned EANTC to test the latest hardware and software aspects of these solutions, the vendor put the focus on 5G readiness. That said, the evaluated functions can be applied for other network service use cases as well. Test results are described in detail in the following sections.

Hardware and Software

Device Under Test	Chassis	Line Cards	Software Version
NE1	ATN950C	Ethernet 50GE	V300 R005C00
		Ethernet 2x25GE	
NE2	NE40E-M2K	Ethernet 2x50GE	V800 R011C00
		Ethernet 4x25GE	
NE3	NE40E-X2-M8A	Ethernet 10GE	V800 R011C00
NE4	NE40E-X8A	Ethernet 10GE	V800 R011C00
NE5	PTN960	Ethernet 10GE	V100 R007C00
NCE	Network Cloud Engine		V100 R018C00

Table 1: Hardware and Software

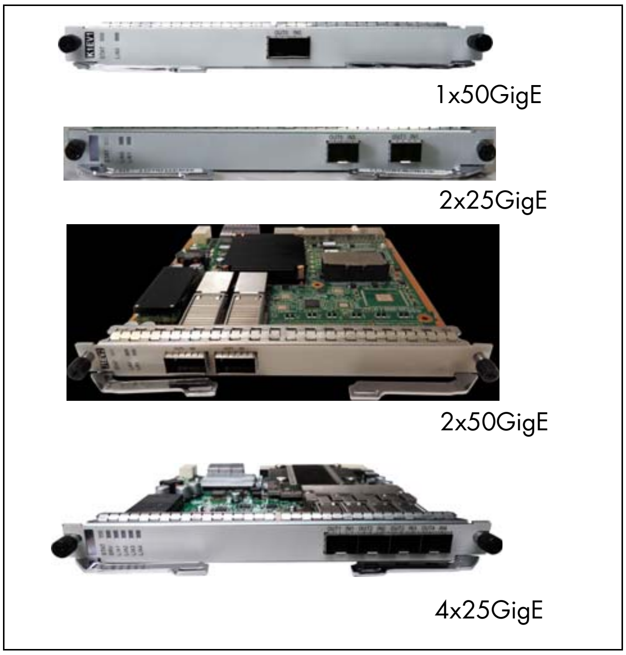

Figure 1: Line Card Under Test

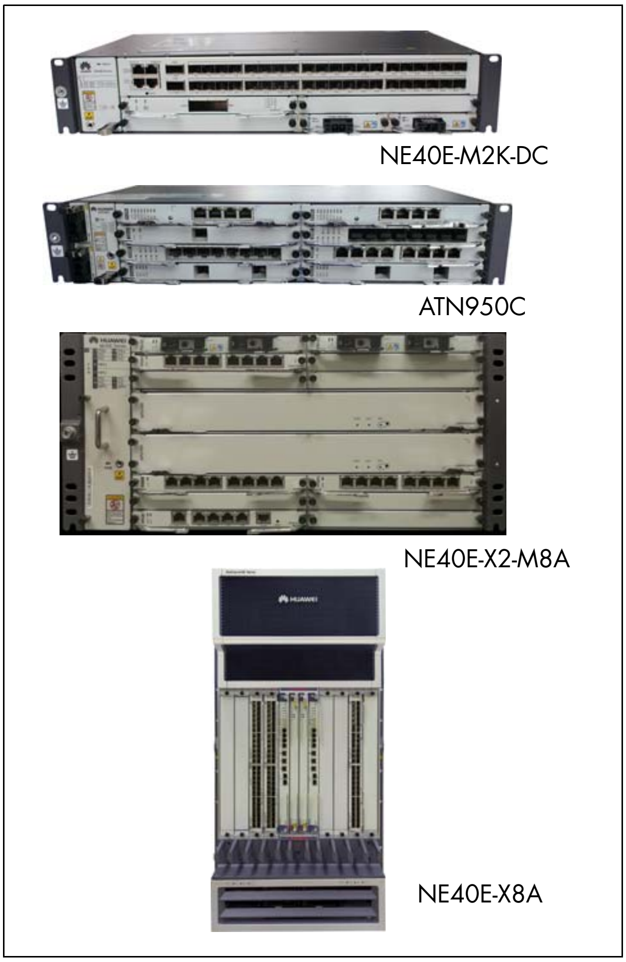

Figure 2: Device Under Test (DUT)

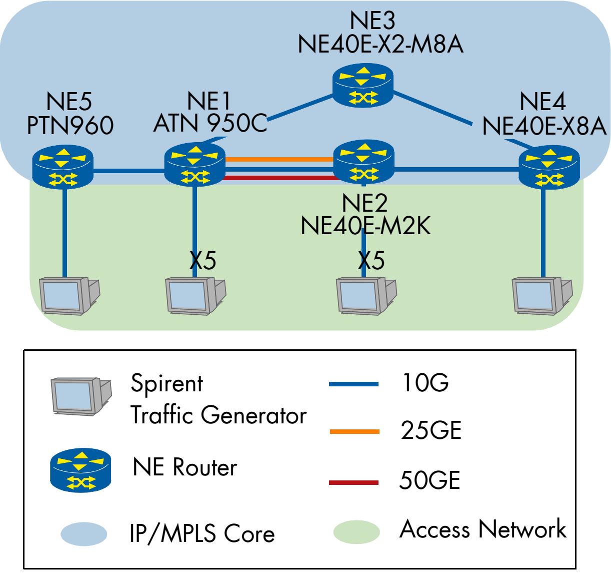

Figure 3: Physical Test Setup

Segment Routing and MPLS

Maximum Label Stack Depth

> The 10GbE line cards on ATN950C and NE40E routers support up to 10 stacked MPLS labels.

Routers are restricted with respect to the label depth supported for a PUSH operation. In this test, we determined the maximum depth of a segment label stack supported on the device under test (DUT). Segment Routing does not require any changes to the operations of the data plane compared with MPLS networks. However, deploying Segment Routing may affect the maximum depth of the MPLS label stack required. As every segment in the list is represented by an additional MPLS label, the length of the segment list directly correlates to the depth of the label stack.

There are several ways to reduce the length of the label stack as discussed in the SPRING draft (spring-segment-routing-mpls-07). Implementing a long path with many explicit hops as a segment list may yield to a deep label stack. Thus, the operator needs to be aware of the routers' limits and take them into account in the design.

We emulated up to 10 segments using three DUTs: ATN950C, NE40E-X2-M8A and NE40E-X8A. To reach the maximum stack depth, Huawei configured an explicit path which passed the ingress node, then traversed back and forth between the latter two nodes until all segments had been added. We captured the packets as shown in Figure 4. We sent traffic via the generated service and did not observe any packet loss, as was expected.

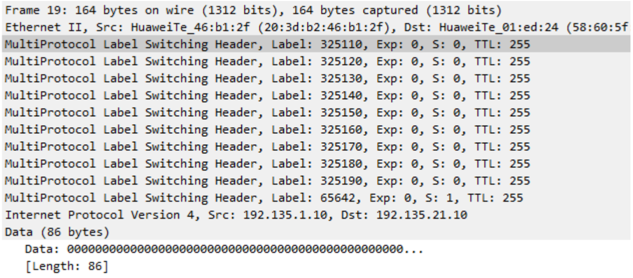

```
Frame 19: 164 bytes on wire (1312 bits), 164 bytes captured (1312 bits)
Ethernet II, Src: HuaweiTe_46:b1:2f (20:3d:b2:46:b1:2f), Dst: HuaweiTe_01:ed:24 (58:60:5f
MultiProtocol Label Switching Header, Label: 325110, Exp: 0, S: 0, TTL: 255
MultiProtocol Label Switching Header, Label: 325120, Exp: 0, S: 0, TTL: 255
MultiProtocol Label Switching Header, Label: 325130, Exp: 0, S: 0, TTL: 255
MultiProtocol Label Switching Header, Label: 325140, Exp: 0, S: 0, TTL: 255
MultiProtocol Label Switching Header, Label: 325150, Exp: 0, S: 0, TTL: 255
MultiProtocol Label Switching Header, Label: 325160, Exp: 0, S: 0, TTL: 255
MultiProtocol Label Switching Header, Label: 325170, Exp: 0, S: 0, TTL: 255
MultiProtocol Label Switching Header, Label: 325180, Exp: 0, S: 0, TTL: 255
MultiProtocol Label Switching Header, Label: 325190, Exp: 0, S: 0, TTL: 255
MultiProtocol Label Switching Header, Label: 65642, Exp: 0, S: 1, TTL: 255
Internet Protocol Version 4, Src: 192.135.1.10, Dst: 192.135.21.10
Data (86 bytes)
   Data: 000000000000000000000000000000000000000000000000...
   [Length: 86]
```

Figure 4: Packet Capture: Ten MPLS Labels

Segment Routing Path Protection

> Segment routing statically configured end-to-end path protection with BFD showed less than 50 ms out-of-service time on primary link failure and zero frame loss during path restoration.

In this test we measured the maximum convergence time in Segment Routing after a link failure event.

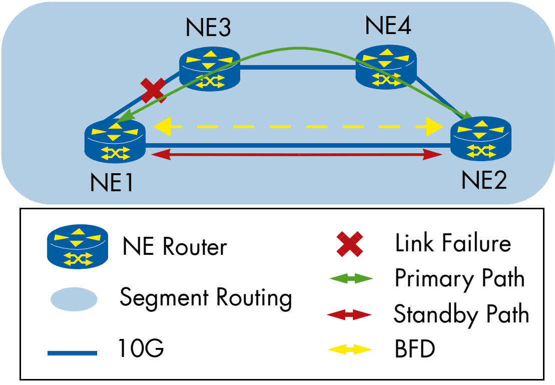

Figure 5: Segment Routing Path Protection

Huawei configured hot-standby protection for Segment Routing and established the primary path between two PE nodes (NE1 and NE2). The endpoints of the path initiated BFD (Bidirectional Forwarding Detection) messages with 10 ms interval to detect network failures over this path. For protection, Huawei configured an explicit backup path manually which worked as RSVP-TE hot-standby-like protection. Huawei explained that the egress node shalldetect remote network failures via BFD.

Huawei asked us to send IPv4-only test traffic at a total rate of 6 Gbit/s for the six L3VPN services. The EANTC team observed via the router's CLI traffic transported over the primary LSP. While traffic was running, we measured the out-of-service time by introducing a link failure (pulling out the cable as shown in the figure). After the link failure we observed that the traffic was switched over to the backup path as expected.

The out-of-service time was measured between 34.7 ms to 41.5 ms which was as expected and below the 50 ms claimed by Huawei. No packet loss was observed during link recovery. We performed this test three times to validate consistent values.

Interconnection of MPLS L3VPN and EVPN-based IP VPN

The NE40E router supports IP-routed splicing (interconnection) between legacy MPLS IP VPNs and MPLS-based IP EVPNs for up to six services simultaneously with 6 Gbit/s traffic in total.

Huawei demonstrated a solution to interconnect legacy MP-BGP based L3VPN (RFC4364) and EVPN-MPLS (RFC 7432). A vendor-specific function on the NE40E-X8A allows interconnection of the two VPN implementations. Huawei explained that an end-to-end IP VPN service is created by a virtual router function through route import. We checked that traffic was forwarded across the different VPN types. Additionally we verified that each pair of VPN services were isolated from each other by using six of such end-to-end L3VPN services (see Figure 6).

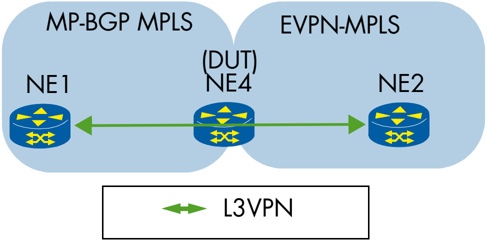

Figure 6: Interconnection MP-BGP Based L3VPN and EVPN-MPLS

As part of the evaluation, we inspected the VRF (Virtual Routing and Forwarding) table on the NE40E-X8A using the CLI. As expected, this table learned the routes of both service types, consisting of the L3VPN routes as well as the EVPN routes, which were imported into the same routing table. For the latter case, we captured the EVPN Segment Routes carried in the MP-BGP updates to verify that the ESI (Ethernet Segment Identifier) was bound per EVPN route via capture. As expected, we observed the requested ESI field in the BGP Network Layer Reachability Information (NLRI) extension.

We also checked the VRF tables on the other two PEs and observed the same set of L3VPN routes. We sent and received 6 Gbit/s test traffic in total for six end-to-end L3VPN services; no packet loss was observed.

MPLS-TP Mode

In a special software configuration for MPLS-TP migration purposes, the ATN950C showed support of bidirectional MPLS-TP tunnels; the MPLS-TP OAM protocol monitored pseudowire status by running continuity check messages.

Huawei asked EANTC to verify the support and connectivity of legacy MPLS-TP pseudowires on the ATN 950C. Supporting MPLS-TP, Huawei explained, provides a migration solution for MPLS-TP using modern, fully supported hardware and software.

The test scenario is depicted in Figure 7. A legacy MPLS-TP router, the Huawei PTN960, was used to terminate the MPLS-TP tunnels. Before starting the verification, we initially checked via CLI that the ATN 950C was correctly configured to enter the MPLS-TP mode. We inspected the correct exchange of MPLS-TP-specific CC (Continuity Check) messages. Additionally we sent 6 Gbit/s bidirectional test traffic which was received without any frame loss.

To verify that the CC function would detect a change of PW status, we disabled the AC site of the PW to emulate the PW "down" status. It was set as anticipated.

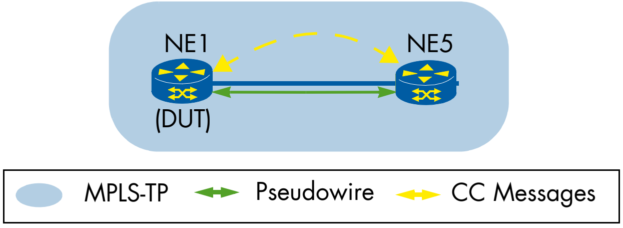

Figure 7: MPLS-TP Mode

EANTC was able to confirm that MPLS-TP functionality can be provided on the ATN950C in a special configuration mode, interoperating with legacy MPLS-TP implementation on the PTN960.

SR-BE (Best Effort) Migration from LDP

ATN950C and NE40E routers support live migration of intra-AS MPLS IP VPN flows from LDP to Segment Routing (Best Effort) with zero packet loss

Live, gradual migration of legacy MPLS services to next-generation Segment Routing services enables a smooth transition. We verified that LDP tunnels could be migrated to Segment Routing while traffic over these tunnels was not affected by the migration.

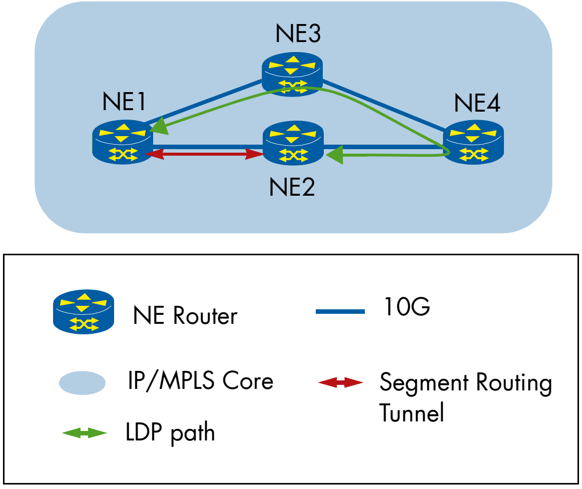

Figure 8: Tunnel Selection for L3VPN

Huawei configured full-mesh LDP paths between all devices of the topology (Figure 8). Simultaneously, Segment Routing was configured with IS-IS TE. The intention, as Huawei explained, was to show that legacy and Segment Routing transport are supported on the same hardware.

We observed fully meshed SR-BE (best effort) paths between all devices. We verified that the SID (Segment Identifier) was exchanged through IS-IS-TE control packets between the segment nodes. This happened according to the IETF draft (ISIS-segment-routing-extensions-18).

To obtain a baseline without any loss we sent test traffic at 6 Gbit/s to the L3VPN services. As expected, no packet loss was experienced. We captured the service packets from the MPLS side and compared the labels carried in the packets to the labels shown in the LDP database. The carried labels matched the database indicating that the test traffic was transported over LDP.

While the traffic was running, Huawei applied the pre-configured policy, allowing the L3VPN to choose the Segment Routing tunnel instead. As expected, the L3VPN service switched over while the test traffic did not show any packet loss. We analyzed the capture of MPLS packets and observed the expected labels as shown in the Segment Routing database, indicating that L3VPN services were transported over Segment Routing.

SR-TE Migration from RSVP-TE

ATN950C and NE40E routers support live migration from RSVP-TE to SR-TE with zero packet loss.

This test case added NCE (Network Cloud Engine), Huawei's network management and control solution extensively tested in the next section below. We verified whether NCE would be able to migrate live L3VPN services running on RSVP-TE tunnels to SR-TE tunnels without affecting services running over these RSVP-TE tunnels.

We started the test by configuring RSVP-TE tunnels using NCE. Then we deployed two L3VPN services over RSVP-TE tunnels via NCE's GUI (see Figure 9). Afterwards, we configured SR-TE tunnels using NCE and verified that the SR labels were different from the RSVP-TE labels by comparing their label database with the DUTs CLI. Then we sent 1 Gbit/s bidirectional traffic through the PEs. While the traffic was running, Huawei triggered the L3VPN service migration action using NCE (see Figure 10). We verified that the traffic switched over to the SR-TE tunnels by checking the SR-TE tunnel counter statistics. The test traffic did not show any packet loss.

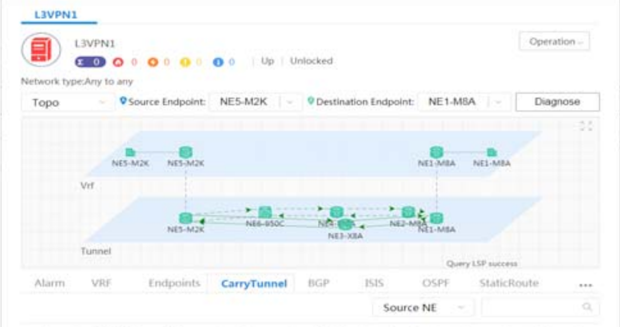

Network type:Any to any

Source Endpoint: NE5-M2K

Destination Endpoint: NE1-M8A

Figure 9: NCE View: L3VPN over RSVP-TE

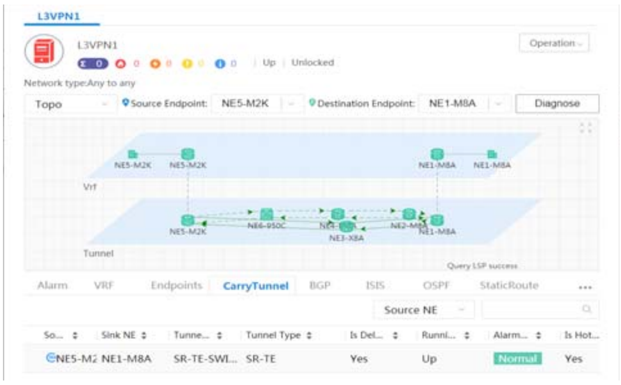

Network type:Any to any

Source Endpoint: NE5-M2K

Destination Endpoint: NE1-M8A

So...	Sink NE	Tunne...	Tunnel Type	Is Del...	Runni...	Alarm...	Is Hot...
NE5-M2	NE1-M8A	SR-TE-SWI...	SR-TE	Yes	Up	Normal	Yes

Figure 10: NCE View: L3VPN over SR-TE

50GE and 25GE Line Cards

As part of the test, Huawei provided new 50GE and 25GE line cards for both router types under test, the ATN950C and the NE40E-M2K. EANTC was asked to determine basic IPv4 forwarding performance of these line cards for VPWS services. By allocating these VPWS services to multiple slices with 2 Gbit/s bandwidth granularity, we also determined the throughput and latency performance of slices using the 50GbE line card.

50Gigabit Ethernet Interface Cards

We measured the interface throughput of the two Ethernet 50GE line cards when they interworked with one of the following optics:

- 40-Kilometer dual-fiber 50 GbE fiber optics
- 50 GbE BIDI (bi-directional) fiber optic

To conduct this test, we connected an ATN950C (NE1) and an NE40E-M2K (NE2) via one 50GE or 25GE connection (port to port), respectively. A Spirent TestCenter traffic generator was used to generate traffic; it was connected with five 10GE ports to each of the two routers. See Figure 3 for details.

Capacity	# Fibers	Maximum Distance	Hardware
50GE	2	40 km	HISILICON OM9380E
50GE	1	10 km	HISILICON OM9382LX101
25GE	2	10 km	Hisense LTF1325-BH+

Table 2: 50GE Parameters

The tested interfaces showed 49.8 Gbit/s data forwarding without packet loss.

Optics Tested	Layer 1 Throughput	
	UNI[a]	NNI[b]
50GE (40 km)	44.4 Gbit/s	49.8 Gbit/s
50GE (BiDi)	44.4 Gbit/s	49.8 Gbit/s

a. 18 Bytes overhead were added on NNI link for VPWS service (14 byte Ethernet plus 4 Byte MPLS label)
b. Including control protocol messages exchanged on NNI

50GE Throughput Performance

The 40 km optics were evaluated with actual 40 km of fiber cables attached (regular Corning fiber, yielding 20 dBm attenuation in each direction). Due to the speed of light, approximately 200 μs of latency were attributed to the transport itself. The average latency was measured at 225.3 μs. The maximum latency remained below 230.3 μs.

The 50GE BiDi (single-fiber) optics were tested in the lab with a cable length of 3 m. In this case, the maximum latency was below 50.4 μs; average latency was measured at up to 29.1 μs.

Sender	Optics Tested	Latency [μs]		
		Min	Avg	Max
ATN950C	50GE (40 km)	216.6	225.3	252.2
NE40E-M2K		218.0	219.0	230.3
ATN950C	50GE (BiDi)	21.6	29.1	50.4
NE40E-M2K		21.0	25.1	38.2

Table 3: 50GE Latency

25Gigabit Ethernet Interface Cards

We performed an interface throughput and latency test for the 25GE line cards on the ATN950C and NE40E-M2K. The test topology was configured similarly to the 50GE test; only three (3) 10GE links were required on each side to connect to the Spirent traffic generator, and only three (3) VPWS services were configured.

The 25GE interface showed 24.9 Gbit/s data throughput without packet loss.

Optics Tested	Layer 1 Throughput	
	UNI[a]	NNI[b]
25GE	22.2 Gbit/s	24.9 Gbit/s
	22.2 Gbit/s	24.9 Gbit/s

a. 18 Bytes overhead were added on NNI link for VPWS service (14 Byte Ethernet plus 4 Byte MPLS label)
b. Including control protocol messages exchanged on NNI

Table 4: 25GE Throughput Performance

The maximum forwarding latency remained below 37.3 μs; we measured an average latency of 22.1 μs.

Sender	Optics Tested	Latency [μs]		
		Min	Avg	Max
ATN950C	25GE	21.3	21.7	37.3
NE40E-M2K		21.7	22.1	30.0

Table 5: 25GE Latency

Channelized Sub-Interfaces for Slicing

Huawei introduced us to a new solution for sub-interfacing implemented on the ATN950C and NE40E-M2K routers. This solution serves an alternative to the hardware-based FlexE slicing. It uses VLAN tagging underneath the MPLS label stack to determine the slices.

For our test, Huawei configured channelized slice interfaces on the 50GE link between the two router types with 2 Gbit/s bandwidth granularity for each slice (2,4,6,8 and 10 Gbit/s). VPWS services were used to transport L2 VPN traffic over these slices. First, we sent test traffic for all slices to obtain baseline throughput without any frame loss at rates shown in Figure 11. We also measured the latency during the test. As expected, the latency values were below 45 μs for all slices (see Figure 12).

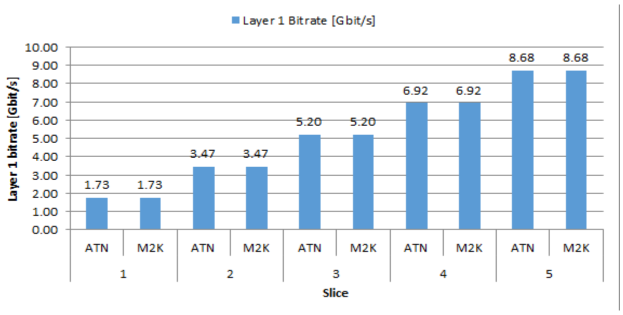

Figure 11: Slice Baseline Throughput (Layer 1)

We verified the bandwidth adjustment by a granularity of 1 Gbit/s and ensured that bandwidth changes for a particular slice did not affect other slices during this process.

At the beginning, Huawei selected a particular slice with the baseline traffic running to increase the bandwidth of this slice from 4 Gbit/s to 5 Gbit/s via CLI. To verify that the bandwidth change had taken place, we increased the traffic rate of the second slice from 3.47 Gbit/s to 4.30 Gbit/s. As expected, all traffic was received without packet loss. The latency value did not increase and stayed consistent to the baseline value (consistent with Figure 12).

Then, we repeated this step by reducing the bandwidth of the slice's bandwidth back to 4 Gbit/s. The baseline traffic passed as expected and the latency value remained consistent.

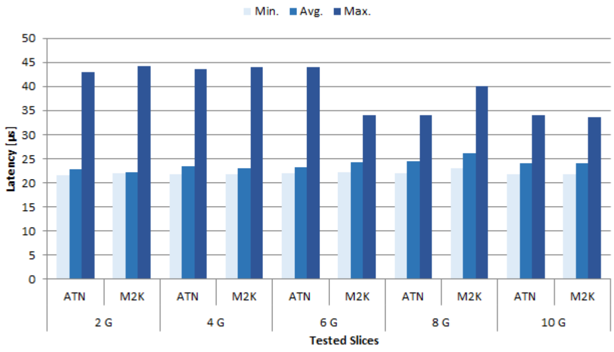

Figure 12: Slice Latency

We also verified that when the traffic exceeded the rate of a particular slice, the excess traffic was discarded and did not affect the traffic of other slices. By increasing the traffic rate of the second slice to 4.30 Gbit/s, as expected, there was no loss for 3.47 Gbit/s traffic passing the slice while no impact was shown on throughput and latency on other slices.

NCE (Network Cloud Engine)

Once we had configured 5G readiness test cases presented by Huawei for network services on the data plane and control plane, we investigated manageability aspects. In large-scale SDN-based transport networks, it will be crucial to provision and maintain services efficiently and to utilize resources optimally.

Huawei introduced NCE, the Network Cloud Engine, to our test and demonstrated an impressive range of features. Specifically, the features focused on SDN orchestration with the Path Computation Element Protocol (PCEP), on EVPN service provisioning, What-If simulation and telemetry.

Constraints-based Path Calculation

> NCE showed path computation for Segment Routing tunnels based on nine constraints.

We analyzed if the NCE is able to perform path computation based on several constraints: cost, bandwidth, latency, explicit-path, hop-limit, Hot-standby, co-route, SRLG and Affinity.

Figure 13 shows the test setup. NCE was connected to NE3 via BGP-LS and by capturing packets we verified that the NCE performed path computation through PCEP protocol and send the configuration commands using SSH -> Netconf.

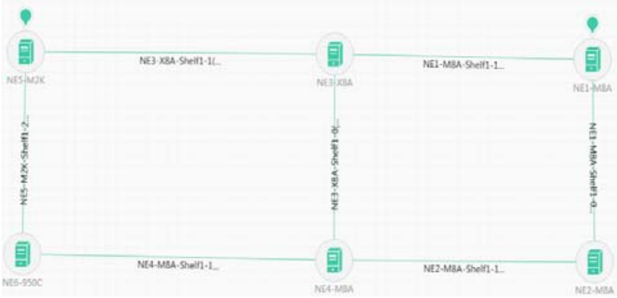

Figure 13: NCE Path Calculation Setup

Before the test, we ensured that the DUT starts with a clean configuration without any tunnels or VPN services. Then Huawei ran the SR-TE tunnel and EVPN-MPLS service in the NCE and applied the configuration from the NCE to the network elements. We tested each constraint separately to get the baseline path. Afterwards, we validated the new path by modifying the configuration constraints in the NCE.

We tested the created path after the NCE computation was done by checking the change in the NCE GUI, as well as performing a trace path command on the DUT's CLI and by sending bidirectional traffic between PEs(NE1 and NE5) and monitoring ports statistics.

Table 6 on page 10 shows the initial path for the initial constraint values and the new path after changing the constraints value.

For cost constraint, the NCE computed the best path according to the shortest path. We verified successfully the path change after increasing the cost value on a particular segment.

Bandwidth constraint was also tested successfully. We confirmed that if the demand bandwidth of a particular tunnel exceeds the total bandwidth reservation for a particular LSP segment, than the NCE is able to change the LSP path.

We also observed that the NCE was able to change the LSP initial path once the latency value of the LSP exceeded the configured threshold.

We tested the Hop-limit constraint and verified that the shortest new path is within the Hop-limit threshold value after changing the initial path cost.

We looked into the NCE's ability to specify and compute the hot-standby paths based on the shortest path. We sent traffic once when Path1 was primary and once again when Path2 was primary and verified no traffic loss.

The NCE was able also to configure the SR tunnel in co-route mode. We verified that the NCE configured two tunnels on the same path in both directions. Each path direction used the shortest path based on the cost attribute. We sent traffic in both directions and measured that it was used in both tunnels.

We saw that the NCE can compute the best path once we explicitly defined a desired node in the path. Figure 16 shows the selected path we defined once in the NCE. Additionally, we sent bidirectional traffic and verified the computed path.

The NCE was able to compute the LSP path based on the Affinity attribute. We verified that the Affinity is able to make up the LSP as primary by configuring affinity on all Path2 segments (see Figure 18).

Finally, we tested and confirmed that the NCE is able to use the SRLG (Shared Risk Link Group) constraint to build the best path. We started the test by configuring NE4-NE3 as primary path and NE4-NE2-NE1-NE3 as standby path. Then, NE4-NE3 and NE4-NE2 links were configured in one SRLG group by the NCE. We verified that the NE3-NE4 remained primary and the standby path changed to NE4-NE6-NE5-NE3 (see Figure 17). We sent bidirectional traffic between NE4 and NE3 and observed that the traffic took the primary path.

All the constraint parameters were configured via the GUI. Figure 14 shows an example of the cost configuration using the GUI and Figure 15 shows an example of CLI verification after changing the cost value attribute:

Figure 14: Configuration per GUI

Setup	Constraint Value		
	Required by SR tunnel	Path 1 (NE1-NE3-NE5)	Path 2 (NE1-NE2-NE4-NE6-NE5)
Legend: 🟩 selected path (NE1-NE5) matching required constraint value			
		Cost	
Initial	Shortest	🟩 120	230
Change		1520	🟩 230
		Bandwidth	
Initial	4 Gbit/s	🟩 10 Gbit/s	10 Gbit/s
Change		0.8 Gbit/s	🟩 10 Gbit/s
		Latency	
Initial	<=800 µs	🟩 400 µs	800 µs
Change		1150 µs	🟩 800 µs
		Hop-Limit	
Initial	<=4 hops	🟩 2 hops (equal to 120 cost)	4 hops (equal to 230 cost)
Change		2 hops (equal to 1520 cost)	🟩 4 hops (equal to 230 cost)
		Affinity	
Initial	Primary include Affinity, Standby exclude Affinity	🟩 Standby	🟩 Primary
		Hot-Standby	
Initial	Primary + Standby	🟩 Primary	🟩 Standby
		Co-Route	
Initial	same path in both directions	🟩 cost of the path higher in one direction	

Table 6: Constraints-based Path Calculation

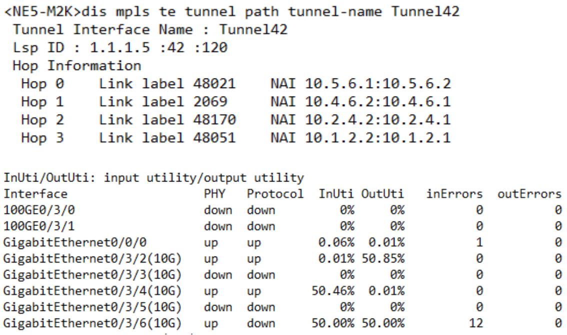

```
<NE5-M2K>dis mpls te tunnel path tunnel-name Tunnel42
 Tunnel Interface Name : Tunnel42
 Lsp ID : 1.1.1.5 :42 :120
 Hop Information
  Hop 0    Link label 48021    NAI 10.5.6.1:10.5.6.2
  Hop 1    Link label 2069     NAI 10.4.6.2:10.4.6.1
  Hop 2    Link label 48170    NAI 10.2.4.2:10.2.4.1
  Hop 3    Link label 48051    NAI 10.1.2.2:10.1.2.1

InUti/OutUti: input utility/output utility
Interface                  PHY    Protocol  InUti OutUti   inErrors  outErrors
100GE0/3/0                 down   down         0%     0%          0          0
100GE0/3/1                 down   down         0%     0%          0          0
GigabitEthernet0/0/0       up     up        0.06%  0.01%          1          0
GigabitEthernet0/3/2(10G)  up     up        0.01% 50.85%          0          0
GigabitEthernet0/3/3(10G)  down   down         0%     0%          0          0
GigabitEthernet0/3/4(10G)  up     up       50.46%  0.01%          0          0
GigabitEthernet0/3/5(10G)  down   down         0%     0%          0          0
GigabitEthernet0/3/6(10G)  up     down     50.00% 50.00%         12          0
```

Figure 15: Selected Path (value displayed via CLI)

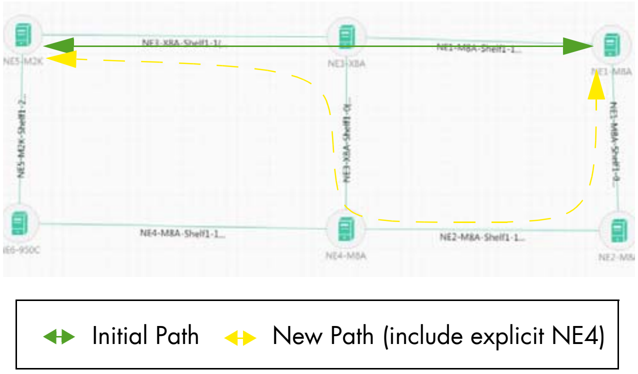

Figure 16: Explicit Path

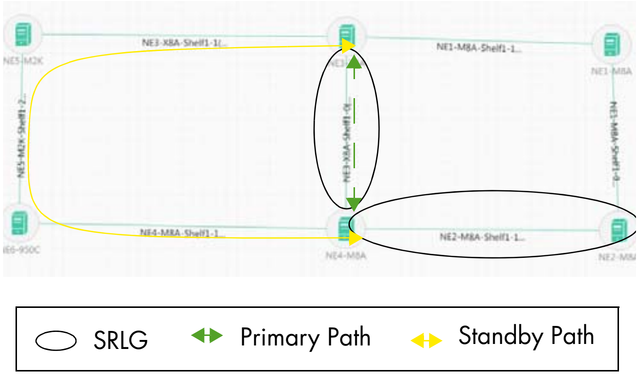

Figure 17: SRLG (Shared Risk Link Group) Constraint

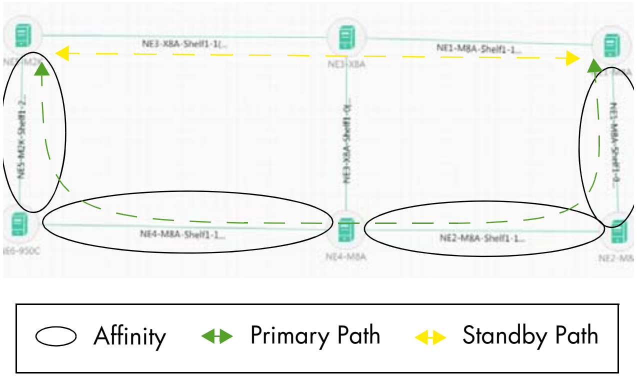

Figure 18: Affinity

L3VPN/EVPN Service Provisioning

NCE provisioned L3VPN-MPLS and EVPN-MPLS services, automatically deployed Segment Routing tunnels by establishing the services and configuring Bandwidth on Demand per service.

We witnessed that the NCE (Network Cloud Engine) was able to provision and remove L3VPN services and L2EVPN services using the NCE GUI and its ability to configure CIR (Committed Information Rate) per service.

Figure 13 shows the test setup. Before the start, we ensured that the DUT is set with a clean configuration without any VPN services. SR only was enabled on all DUTs. We verified the service creation by checking the change in the NCE's GUI as well as by checking the routing table and configuration on the DUT's CLI. We also sent bidirectional traffic between the PEs(NE1 and NE5) and monitored ports statistics.

Huawei created one L3VPN between NE5 and NE1 by using NCE. Few steps were required for the configuration: specifying the L3VPN template, selecting the NE node, choosing a name for a VRF and specifying the access interface with an IP address. The Route Distinguisher value and importing and exporting the Route Targets were dynamically allocated by the NCE. A CIR value matching 1Gbit/s was also configured on the UNI interface. After applying the configuration, we verified that the configuration on-boarded on the DUT by checking the DUT's CLI configurations and the PEs' routing tables (Figure 19 shows the on-boarded service from the NCE's GUI). First, we sent 1Gbit/s bidirectional traffic using a 1518Byte frame size between the PEs and observed no packet loss as expected. Then, we increased the traffic above the CIR value and noticed packet drops. Again as expected.

Then Huawei performed the same steps to deploy L2EVPN services on both NE1 and NE5. We verified the on-boarded configurations on the DUTs' CLI and EVPN MAC and routing table on the PEs. We sent1Gbit/s bidirectional traffic using a 1518Byte frame size between the PEs and verified no packet loss as anticipated. Then we increased the traffic above the CIR value and noticed packet drops as expected again.

Finally, we looked into the service deletion action. The NCE was able to remove the selected L3VPN and EVPN services. We checked the DUTs configuration files and all VPN instances were removed.

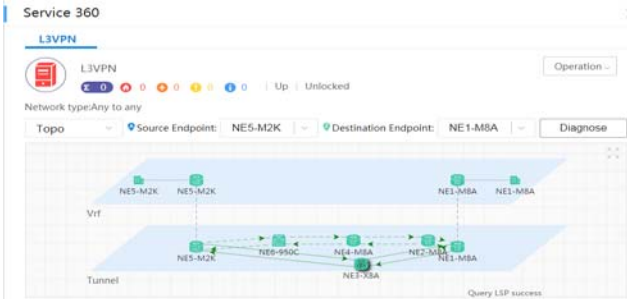
Service 360

L3VPN

L3VPN

Σ 0 0 0 0 0 | Up | Unlocked

Network type:Any to any

Topo | Source Endpoint: NE5-M2K | Destination Endpoint: NE1-M8A

Query LSP success

Figure 19: L3VPN Service (Displayed by NCE)

What-if Simulation

What-if analysis simulated a link failure, computed the optimal path based on failure, predicted the idle time for maintenance using 24 hours records of traffic load and live traffic statistics and performed maintenance operation automatically.

The what-if simulation feature supported by the NCE is a way to automate maintenance implementation plans. This can be achieved by keeping records of the historical traffic load of the network to accurately estimate the idle time of maintenance and by running what-if analysis functions to pre-compute the best path after simulating a failure. It is based on the precise result of the what-if algorithm as error-free maintenance is desired. We witnessed that the NCE is able to accomplish a what-if simulation and analysis function within a reasonable time.

The same setup as depicted in Figure 13 was used. Before starting the test, we ensured that the SR-TE and L3VPN services between the PEs were already configured. We kept the 5Gbit/s bidirectional traffic running between the PEs. At first, Huawei selected from the NCE simulated traffic types, historical tunnel flows that were previously recorded for 24 hours as well as live traffic flows, which were used as simulated traffic to predict the maintenance time. The NCE used the average rate of the flows within the historical records to obtain the suitable maintenance time. Secondly, we simulated a failure on the NE5-NE3 link and run the what-if analysis. Path computation, service transition and displaying the result on GUI took around one minute.

The left hand side of Figure 20 shows the path before simulation and the right hand side shows the path computation result after simulating the link failure. Figure 21 illustrates the best time to apply the maintenance policy. Finally, Huawei applied the maintenance policy. After a few seconds we noticed that the maintenance operation was successfully implemented. That was as expected in regards to the simulation result (see Figure 22).

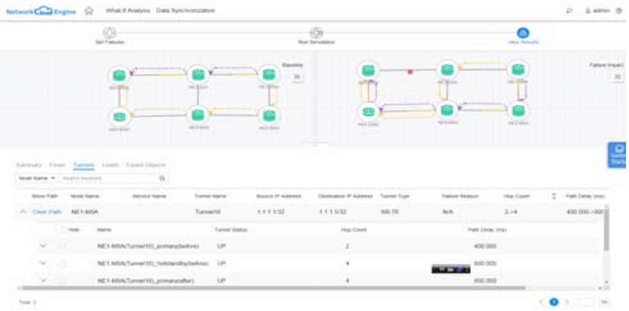

Figure 20: What-if Simulation Path Computation

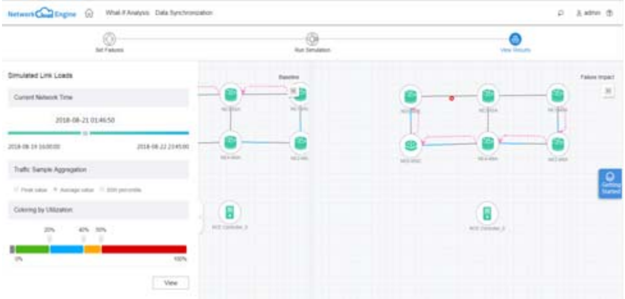

Figure 21: What-if Simulation Time Prediction

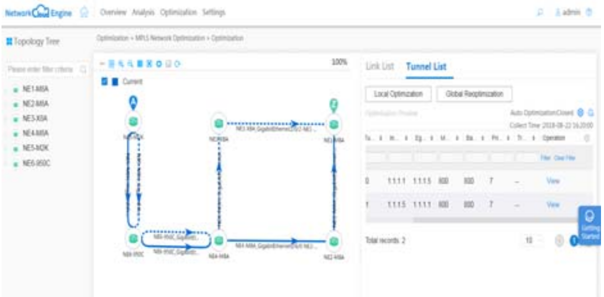

Figure 22: Maintenance Policy Result

Telemetry

Telemetry collected interface statistic via gRPC (Google Remote Procedure Call) based on 1s collection interval and showed live graph based on one minute sample interval.

The SNMP protocol is often used for network performance monitoring. However, the SNMP pull model cannot scale for today's large number of devices. Streaming-Telemetry with its push model mechanism is an alternative to the SNMP protocol and overcomes the weakness of the pulling model. According to Huawei, the NCE uses OpenConfig streaming telemetry mechanisms via gRPC protocol (Google Remote Procedure Call).

In this test we verified that the NCE is able to subscribe to specific data items in the DUTs and that the NCE is able to collect live traffic statistics from network elements.

The same setup as in Figure 13 was used. Before starting the test, we ensured that the SR-TE and L3VPN services between the PEs were already configured and no telemetry subscription existed on the DUTs.

Huawei started the test by configuring health-check monitor instances for all NEs nodes. gRPC was configured as the collect protocol and 10 seconds as collection interval were set. The NCE was able to show a live health report included live CPU and Memory statistics represented in a live graph based on a one minute sample interval.

Then, interfaces monitor instances were configured for all NE nodes and tunnel monitor instances were configured for all PE nodes (NE1 and NE5). Again gPRC was configured as the collect protocol with one second as collection interval. We verified the telemetry subscription on the DUTs using its CLI. All DUTs showed two gRPC subscription sessions on the NCE, indicating health-check and interface monitor instances. In other hand, PE nodes had three subscription sessions to the NCE due to tunnel monitor instance. Then we sent 5Gbit/s bidirectional traffic between the PEs and monitored the interfaces and tunnels report on the NCE. The NCE was able to show live statistics information for both UNI and NNI interfaces such as: average and peak of transmitted, received bandwidth utilization and traffic rate in addition to CRC error packets. Furthermore, we increased the traffic by 25% and we noticed that the NCE was able to show the change within the one minute interval as expected. Figure 23 shows the Telemetry live traffic tunnel report once the traffic increased.

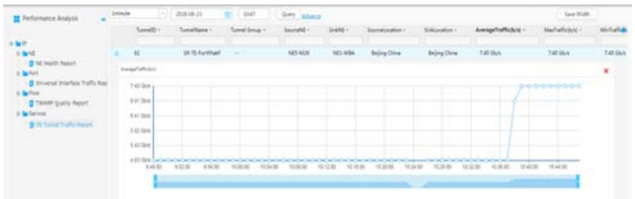

Figure 23: Telemetry Tunnel Report

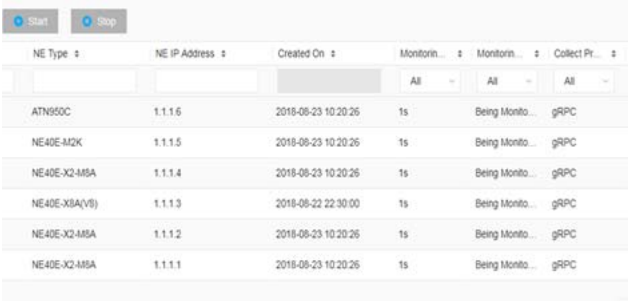

NE Type	NE IP Address	Created On	Monitorin...	Monitorin...	Collect Pr...
			All	All	All
ATN950C	1.1.1.6	2018-08-23 10:20:26	1s	Being Monito...	gRPC
NE40E-M2K	1.1.1.5	2018-08-23 10:20:26	1s	Being Monito...	gRPC
NE40E-X2-M8A	1.1.1.4	2018-08-23 10:20:26	1s	Being Monito...	gRPC
NE40E-X8A(V8)	1.1.1.3	2018-08-22 22:30:00	1s	Being Monito...	gRPC
NE40E-X2-M8A	1.1.1.2	2018-08-23 10:20:26	1s	Being Monito...	gRPC
NE40E-X2-M8A	1.1.1.1	2018-08-23 10:20:26	1s	Being Monito...	gRPC

Figure 24: 1 s Collection Interval

REST-API

NCE supports REST-API for adjustment of link cost for path optimization.

In this verification, we saw that the Huawei NCE's REST-API calls support a modifying link cost for path optimization.

The same setup as shown by Figure 13 was used. Before starting the test, we ensured that the SR-TE was already configured with default cost values. We analyzed the path optimization by checking the change in the NCE's GUI, as well as performing trace path commands on the DUT's CLI.

Huawei started the test by checking the initial LSP path which was NE5-NE3-NE1. The engineer prepared four API calls and used the Postman tool to send API to receive and post commands to the NCE's API northbound interface.

The purpose of the first API call was to query information about the pre-configured tunnels and to obtain the tunnel-id to change its cost. The API received successfully and information returned back in XML format. The second API call was a "put" command to change the cost link value from 100 to 10 of the NE5-NE6 and NE6-NE4 links. We verified the new cost value from the DUT's CLI . A third API call was added as a "post" command. It was sent to trigger the tunnel computation process on the NCE. The fourth call was to query the computation task status. The response indicated that the computation task finished with successful results. All in all, we verified that the new LSP path changed to NE5-NE6-NE4-NE2-NE1. Figure 25 depicts this result.

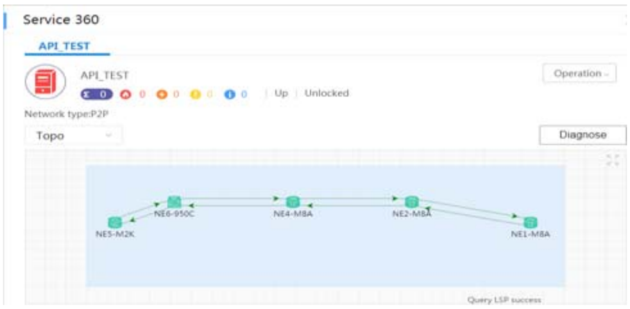
Service 360

API_TEST

API_TEST

Σ 0 0 0 0 0 | Up | Unlocked

Network type:P2P

Query LSP success

Figure 25: API Tunnel Computation Result

Conclusion

Huawei demonstrated its 5G-ready transport network solutions through an agreed list of tests on products, protocols, and the SDN controller. In our campaign, we tested the ATN950C and NE40E-M2K being added to Huawei's 5G-ready transport network product portfolio. Both routers support 25GE/50GE interfaces with slicing capabilities and Segment Routing/EVPN, which are key features for 5G readiness of the transport network. Additionally Huawei showed its Network Cloud Engine (NCE) as the *brain* of its transport network solution. EANTC undertook an extensive test of NCE covering a wide range of features including service provisioning and path calculation. In short, Huawei's products showed very good results delivering its 5G ready transport network solution.

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.

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

EANTC AG
Salzufer 14, 10587 Berlin, Germany
info@eantc.com, http://www.eantc.com/
[v1.3 20180910]