EANTC Independent Test Report

Huawei 4T Line Processing Unit

Performance, Power Efficiency, Scalability and Functionality

January 2018

HUAWEI

EANTC

About EANTC

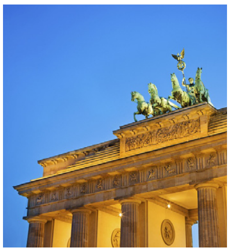

EANTC (European Advanced Networking Test Center) is internationally recognized as one of the world's leading independent test centers for telecommunication technologies.

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.

Table of Contents

Introduction

Huawei Technologies commissioned EANTC to verify the functionality and performance of its new 4T Line Processing Unit (LPU) for Huawei's NE9000 Series Routers. The 4T LPU is powered by Huawei's in-house designed fourth-generation Solar 5.0 Network Processing (NP) chipset which Huawei designed to support 1 Terabit/second full-duplex forwarding capacity. The card houses 40 (forty) 100Gigabit-Ethernet (100GbE) interfaces.

Test Highlights	
Throughput	➔ Line rate (4 Tbit/s) IP forwarding performance with IMIX packet distribution and uniform packet sizes of 256 bytes or larger
Power Efficiency	➔ 0.35 Watt per Gbit/s weighted energy consumption rate per 4T LPU
Scalability	➔ BGP scalability up to 80 Million IPv4 routes, or simultaneously 40 Million IPv4 ➔ FIB scalability up to 5.12 Million IPv4 and 2.048 Million IPv6 routes
Functionality	➔ Full system virtualization support with Huawei's Virtual Systems feature ➔ Ethernet VPN (EVPN) functionality with both Segment Routing and MPLS, and VXLAN encapsulation ➔ Next Generation Multicast VPN (NG MVPN) support

As the name "4T" indicates, the line card aims at providing four (4) Terabit/second throughput in a single Huawei NE9000 slot. The goal is indeed impressive and is likely to be a welcome upgrade to service providers especially those that already have the NE9000 installed in their network, and are running out of capacity.

EANTC ran a range of tests on the network card, with particular focus on cloud data center and service provider infrastructure use cases. Our test initially corroborated the functional aspects of the network card, followed by measuring the data plane forwarding performance, control plane scalability and finally the power consumption.

NE9000-8

NE9000-20

4T LPU

Executive Summary

Huawei set the bar high with its new 4T Line Processing Unit (LPU), by providing a 4 Tbit/s per slot full-duplex forwarding performance and a theoretical maximum 80 Tbit/s full-duplex forwarding capacity in a single NE9000-20 chassis.

The core router features we tested functioned as expected and as Huawei advised.

We found that the new LPU is not only suitable for IP core application; it is also a good fit for provider edge and Internet peering scenarios, metro aggregation as well as data center gateways due to its large forwarding and routing (FIB and RIB) scale and its support for EVPN, segment routing, MPLS, VXLAN and NG-MVPN functionality.

Tested Devices and Equipment

EANTC tested the 4T LPU line card installed on an Huawei NE9000 Converged Backbone Router. Depending on the test case, one, two or three LPU were installed in a single chassis.

The Huawei NE9000-20 tested configuration for the first device under test (DUT1) included 8 Switched Fabric Units (SFU) and 2 Main Processing Unit (MPU). The implementation could support up to 20 LPUs. The chosen SFU redundancy model was 7+1 (seven active units plus one standby unit), and the MPUs worked in active/standby mode.

For the performance, scalability and power consumption test cases, we used the Huawei NE9000-20 router directly connected with test equipment while in the functionality tests, we deployed a network topology using Huawei NE9000-20(DUT1), two Huawei NE9000-8 (DUT2 and DUT3), all three routers are with 4T LPU Line cards installed, one additional Huawei NE40E customer edge router and emulated customer edge routers as shown in Figure 1.

Huawei claims that the NE9000-20 and NE9000-8 routers are programmed with the same scope of functionalities, but designed with different chassis sizes, forwarding capacities and power consumption etc.

All core routers were tested with Huawei Versatile Routing Platform VRP (R) software version 8.160, while the NE40E customer edge router was running Huawei Versatile Routing Platform VRP (R) software version 8.150. Single-mode QSFP28 transceivers were used in all tests.

An Ixia XGS12 traffic generator was used with four Lava AP40/100GbE-2P modules.

As shown in the table below, EANTC defined a particular IPv4 and IPv6 traffic mix ("IMIX") with a range of packet sizes to ensure a realistic IP traffic load in the core network.

Frame Size (Bytes)	Weight	Percent of Packets (%)
64 (IPv4) or 78(IPv6)	3	5
100	26	41
373	6	10
570	5	8
1300	6	10
1518	16	25
9000	1	2

Table 1: IMIX Traffic Definition

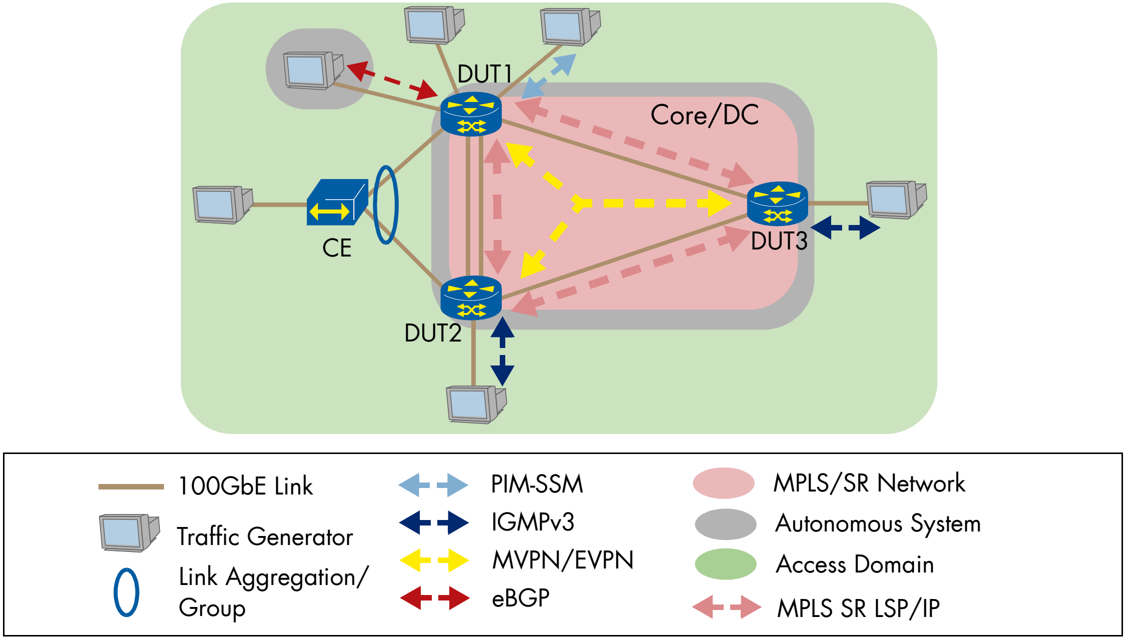

Figure 1: Overall Test Topology

Test Results: Performance

The goal of the test was to measure the forwarding performance of the new line card. We tested IP/MPLS forwarding performance with two 4T LPUs as shown in Figure 2.

Since a total of 80 tester ports were required, we used a "snake" configuration between LPU 1 and 2 using a single-mode fiber patch cord. Huawei's engineers additionally configured 40 VRFs between LPUs interfaces. This setup allowed us to test the 4 Tbit/s forwarding capacity of the line card, using a 100 Gbit/s data rate between two test generator ports TG1 and TG2.

We performed two series of test runs. In the first series we sent IPv4-only test traffic. Per Huawei's request, we sent the traffic with 100 IPv4 prefixes per VRF, resulting in a total of 4,000 prefixes.

In the second run we sent a mix of IPv4 and IPv6 traffic with a 80:20 ratio and 100 IPv4 plus 100 IPv6 prefixes per VRF, resulting in a total of 8,000 prefixes. The service type was L3VPN as defined in RFC 4364.

Huawei requested to select 256 byte-sized packets as the smallest size for both IPv4 only and IPv4/IPv6 mix traffic test runs. Additionally, we performed test runs using uniform frame sizes of 512, 1024, 1518 and 9000 bytes. In addition, we conducted one more test run with IMIX frame sizes as specified in Table 1.

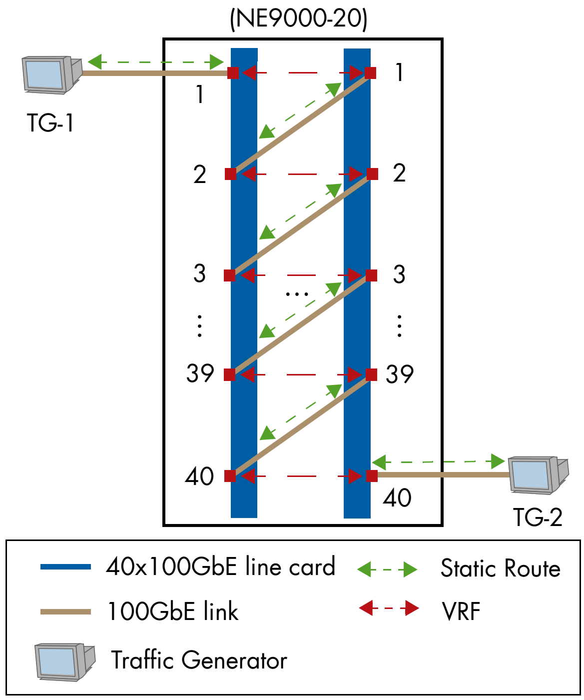

Figure 2: Test Setup — Forwarding Performance

For each frame size, we sent full line rate of bidirectional traffic for the duration of 600 seconds. The 4T LPU forwarded the traffic with zero packet loss in each case.

➔ 4T LPU achieved a maximum packet forwarding rate of 4 Tbit/s bidirectional traffic with 256 bytes frames

➔ The minimum/average/maximum latency for any IPv4/IPv6 packets forwarded across the L3VPN services were below 12 µs/20 µs/21 µs (microseconds)

During the test, we also measured forwarding latency for each frame size; however, the use of "snake" configuration meant that each packet traversed the router 40 times. We divided the measured latency by 40 to obtain forwarding latency per port pair (input and output port). Figure 3 and Figure 4 show the forwarding latency for each frame size. In both test runs (IPv4-only and IPv4&IPv6 mix), the measured latencies was approximately identical.

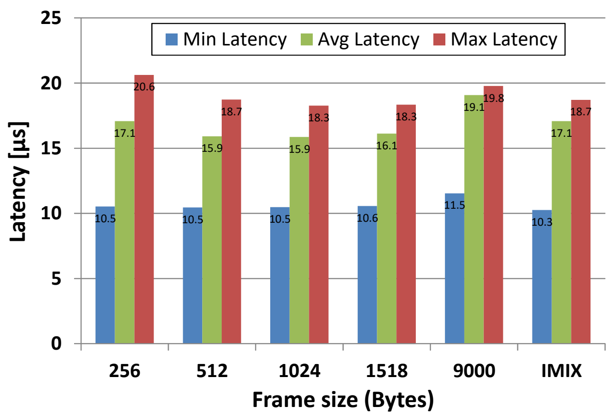

Figure 3: IPv4 Traffic Latency

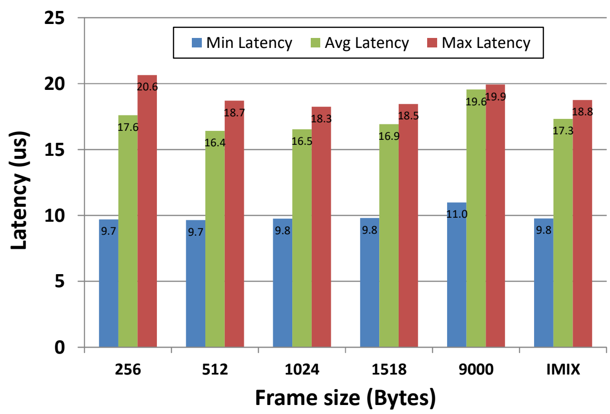

Figure 4: IPv4& IPv6 Mix Traffic Latency

EANTC concluded that Huawei's 4T LPU can support 4 Tbit/s full duplex forwarding capacity with maximum of 21 µs latency. The smallest packet size tested was 256 bytes.

Test Results: Power Efficiency

Service providers are sensitive to the electrical power consumed by network devices. The more energy-efficient a device is, the more a service provider can save energy costs, which leads to reduction of operational expenditures (OpEx).

Huawei was keen to demonstrate that the new 4T LPU is even more power-efficient than the LPUI-1T line card we tested previously[1].

Initially, we measured the power consumption of the NE9000-20 chassis with two 4T-LPU installed as shown in Figure 5. Each 4T-LPU used the same "snake" configuration mode with two ports connected to a Traffic Generator. The Huawei NE9000-20 device was configured in basic (default) power mode.

➔ Energy Consumption Rating Weighted (ECRW) of the 4T LPU reached to 0.35 Watt per Gbit/s

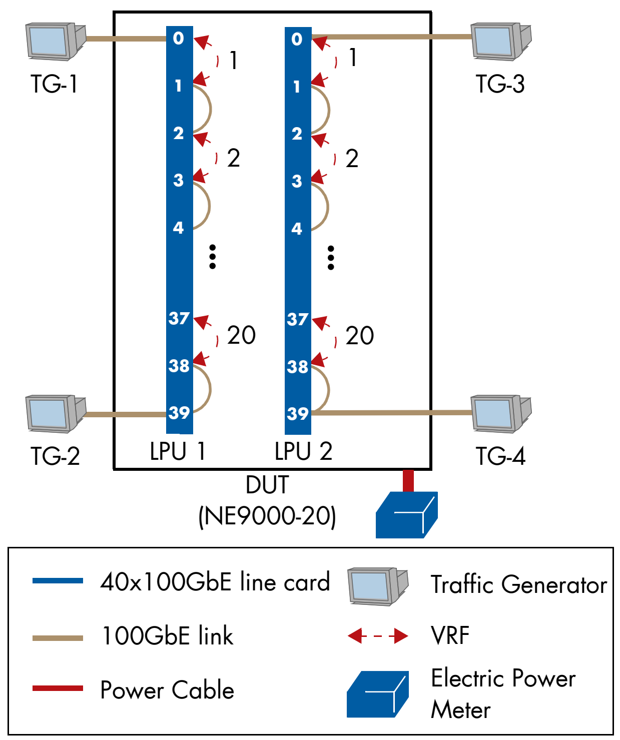

Figure 5: Test Setup — 4T LPU Power Consumption

Based on test methodology defined by the Alliance for Telecommunication Industry Solutions (ATIS) ATIS-0600015.03.2016 standard, we measured the weighted power consumption of the base NE9000-20 chassis including two 4T LPU in three instances: no traffic, 30% of the traffic rate, and finally 100% of the traffic rate.

As a next step, we removed one 4T LPU and performed the same three measurements again.

The weighted power consumption of the 4T LPU line card reached to 1,402 Watt in total.

Finally we calculated ECRW (Energy Consumption Rating Weighted) This metric describes the amount of weighted energy consumption required by the device to transmit one Gigabit of line-level data per second.

The calculated ECRW was 0.35 Watt per Gbit/s.

Test Results: Scalability

Once we confirmed the line rate forwarding performance of the 4T-LPU line card, we asserted the control plane scalability of the line card.

➔ 4T LPU required max. 62 seconds to install 5.12 Million IPv4 routes

➔ 4T LPU required max. 94 seconds to install 5.12 Million IPv4 plus 2.048 Million IPv6 routes

➔ NE9000 series router with 4T LPU remained stable while routes exceeded its maximum capacity

FIB Scalability

The Forwarding Information Base (FIB) size is one of the key metric of control plane scalability for the operator when designing and operating the network.

For the FIB scalability test we used BGP to advertise the maximum number of routes that the 4T LPU line card could install on the forwarding engine.

The setup is depicted in Figure 6.

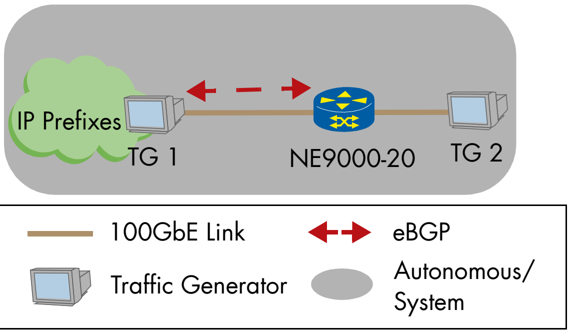

Figure 6: Test Setup — FIB Scalability

We performed the following test runs:

- Run 1: IPv4 only 5,120,000 unique routes
- Run 2: 5,120,000 unique IPv4 and 2,048,000 unique IPv6 unique routes

In run 1, we emulated one BGP router with a single eBGP session between traffic generator and the Huawei NE9000-20. To reflect the reality, we decided to advertise the routes with different prefix lengths ranging from /8 to /32. We monitored CPU utilization of the 4T LPU during and after the route population. The maximum CPU utilization was 31% while learning the routes and it was reduced to 11% after the routes were installed in to the FIB. The 4T LPU utilized the memory usage from 12% to 30%.

In the second test run, we configured a IPv6 eBGP peer to advertise IPv6 routes. We configured traffic generator to advertise IPv4 and IPv6 prefixes simultaneously from both BGP peers. This time, the maximum CPU utilization was 30% while learning the routes and it was reduced to 11% after the routes were installed in to the FIB. The line card utilized the memory usage from 11% to 23%.

In both test runs, we measured the time to install the routes in to RIB and FIB.

The following table shows a summary of the results.

FIB entries	RIB population time [s]	FIB population time [s]
5,120,000 IPv4	45	62
5,120,000 IPv4 & 2,048,000 IPv6	IPv4: 83 IPv6: 91	IPv4: 88 IPv6: 94

Table 2: RIB/FIB Population Time

To verify that all routes are installed into the FIB, we sent IMIX test traffic towards all the advertised prefixes and observed no packet loss.

After verifying the great scaling numbers of the 4T LPU, we proceed to verify how the FIB would react when receiving more prefixes than the actual supported.

For this purpose, we advertised 10% of prefixes more than the supported maximum number of routes. As expected, the additional routes were not installed in the FIB. NE9000-20 generated alarm and syslog messages indicating that the number of received routes exceeded the system's capacity. We observed that the router was stable and it was still manageable by SSH, and also the traffic designated to the original received routes were unaffected.

BGP Route Scalability

BGP is a routing protocol that was designed to carry a massive amount of routes. Since BGP routes are growing year by year, route table capacity is an important scaling factor for a network operator when designing and operating networks.

➔ Huawei NE9000-20 router sustained 80 Million IP4 BGP routes with 100 BGP sessions

➔ NE9000-20 router supported 40 Million IPv4 routes plus 20 Million IPv6 routes with 100 IPv4 plus 100 IPv6 BGP sessions

The goal of this test was to verify that the Huawei NE9000-20 is able to sustain 80 million routes. To verify this, we used four interfaces of 4T LPU directly connected with the traffic generator as shown in Figure 7. We configured 25 IPv4 eBGP sessions and 25 IPv6 eBGP sessions on each interface, resulting a total of 100 IPv4 eBGP and 100 IPv6 eBGP sessions.

We performed two different test runs:

- Run 1: 100 IPv4 eBGP sessions, with 800,000 same set of prefixes per session. A total of 80 Million prefixes.
- Run 2: 100 IPv4 eBGP sessions and 100 IPv6 eBGP sessions, with 400,000 IPv4 prefixes per IPv4 session and 200,000 IPv6 prefixes per IPv6 session. A total of 40 Million IPv4 prefixes and 20 Million IPv6 prefixes.

The setup is depicted in Figure 7.

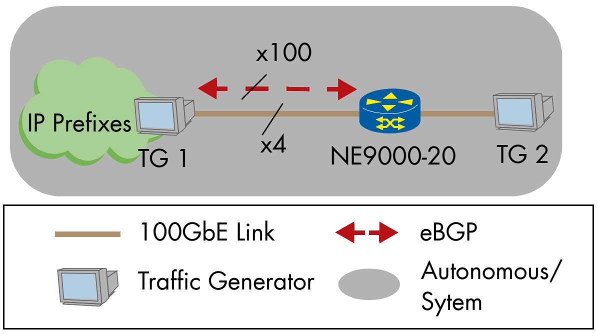

Figure 7: Test Setup — BGP Route Scalability

We verified that 80 million IPv4 routes were installed into the BGP routing table and as expected 800,000 routes were installed into the FIB table. Finally, we sent traffic towards all the advertise prefixes and as expected the 4T LPU line card forwarded the traffic to 800,000 unique prefixes with the best path.

We can conclude that the platform can easily support 100 eBGP full-route Internet feeds and it can perfectly work as an Internet Peering Router/IGW.

In the first test run, the memory utilization of the MPU was increased from 6 % to 43 % after the routers were installed. In the second test run, we observed a memory utilization of 34 %.

Test Results: Functionality

As previously mentioned, our tests started with the functionality test area. In this area we tested the following features:

- Virtual System (VS) separation
- Equal-cost Multi-path (ECMP) load balancing
- Ethernet VPN (EVPN) with MPLS data-plane and Segment Routing (SR)
- EVPN with Virtual eXtensible Local Area Network (VXLAN) data-plane
- Next Generation Multicast VPN (NG-MVPN)

Virtual System Separation

Virtual System (VS) is a technology that provides virtual separation of physical resources installed within the same chassis. By means of VS, a router can be shared among different organizational units or tenants, while keeping all traffic separated and secured. Additionally, Virtual System technology optimizes physical resource allocation and reduces capital and operational costs.

➔ Huawei's Virtual System (VS) can be managed and operated independently

➔ Virtual System isolated traffic from other VS

The goal of this test was to verify that virtual systems are completely separated and independent.

As shown in Figure 8, we configured two VS on Huawei's NE9000-20 router. We assigned one Autonomous System (AS) to each VS and while using a single 4T LPU, we assigned 5 ports to each VS. Furthermore, we advertised 10,000 IPv4 prefixes and 10,000 IPv6 prefixes by eBGP from Traffic Generators and transmitted a dual stack IPv4/IPv6 traffic.

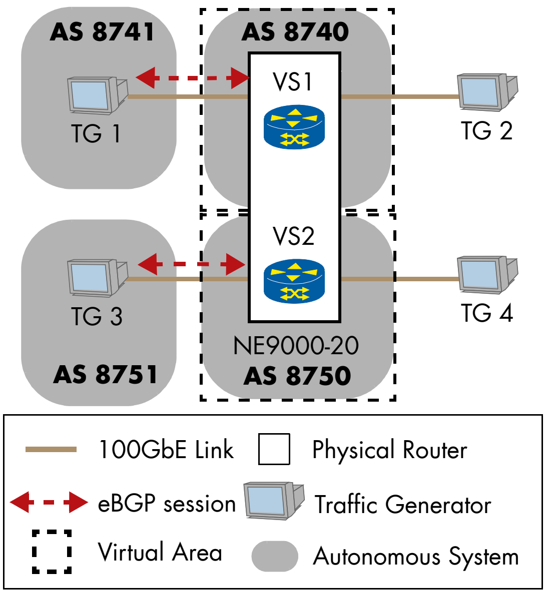

Figure 8: Test Setup — VS Separation

During the test, we concluded that:

- IPv4 and IPv6 traffic forwarded within the VS with no packet loss and no traffic leaking across the VSs
- VSs can be managed independently
- Each VS maintained separate configuration file
- Interfaces from same or different line card can be allocated separately for each VS
- Reboot of one VS did not impact the traffic serving other VS

Each VS can maintain separate syslog.

Equal Cost Multi-Path (ECMP) Load Balancing

ECMP Load balancing allows the network devices to use multiple equal cost paths to reach the destination, thus optimizing resource utilization.

According to Huawei, Huawei's ECMP implementation can use several hashing algorithms such as MAC, IP, TCP/UDP ports, MPLS labels etc for multiple scenarios like IP network, MPLS and VXLAN. In our test, Huawei used destination IP field based hashing algorithm for L3VPN traffic and destination MAC field based hashing algorithm for L2VPN traffic load distribution.

Figure 9 shows the test setup, where Traffic Generator 1 (TG1) generated unidirectional traffic to TG2, and traffic was load balanced on DUT1. Initially we performed the test separately for two different services. In the first run, we used L3VPN service with two VPN instances to generate 5,000 IPv4 flows and 5,000 IPv6 flows. In the second run, we used L2VPN (3 VPWS services) to generate 10,000 layer 2 traffic flows.

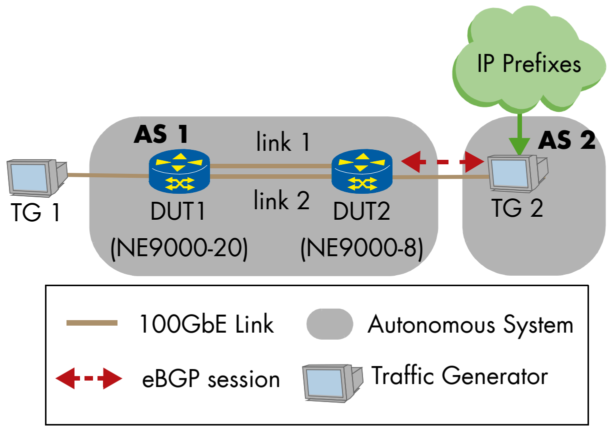

Figure 9: Test Setup — ECMP Load Balancing

In both test runs, we sent unidirectional traffic at the rate of 50 Gbit/s with fixed 256 Bytes Packet size. As expected, that traffic was equally distributed across Link 1 and 2. We did not observe packet loss nor added latency.

In real network environment, the core network devices perform ECMP load balancing for multiple services simultaneously. Therefore, we asked Huawei engineers to perform one more test run with multi-service traffic. In this test run, we generated 5,000 IPv4 flows, 5,000 IPv6 using L3VPN services and simultaneously generated 10,000 Layer 2 flows using L2PN services. We sent unidirectional traffic at line rate (100 Gbit/s).

The traffic was equally distributed for all services on both interfaces.

➔ 4T LPU line card supported ECMP load balancing with multi-service (L2VPN and L3VPN) traffic at 100 Gbit/s data transmission rate

Ethernet VPN (EVPN)

As defined in RFC 7432, Ethernet VPN (EVPN) is a BGP-based next generation solution to provide L2VPN multipoint services by addressing the requirements of Carrier Ethernet and Data Center Interconnect (DCI) market segments. EVPN addresses some of the VPLS limitations when it comes to redundancy, multicast optimization, and provisioning simplicity.

EVPN provides separation between the data plane and control plane, which allows the use of different encapsulation mechanisms in the data plane such as MPLS services and Virtual eXtensible Local Area Network (VXLAN).

EVPN with VXLAN scenario is mostly used in Data center scenarios. It interconnects a layer 2 Ethernet segment over a layer 3 network.

Over the last couple of years, we have recognized growing demand for EVPN-based solutions at EANTC; as a consequence, we decided to include EVPNs as a major test area in this test suite.

We divided the EVPN test section into two parts:

First, we tested EVPN using an MPLS data plane and segment routing as the label allocation/distribution protocol. Afterwards, we tested EVPN with Virtual eXtensible Local Area Network (VXLAN) encapsulation on top of a pure IP core.

- ➔ Huawei NE9000 series routers with 4T LPU supported EVPN deployment over MPLS data plane based on Segment Routing
- ➔ Huawei NE9000 series routers with 4T LPU supported EVPN deployment with VXLAN as data plane encapsulation

EVPN with MPLS/SR Transport

In the first part of our EVPN test, Huawei engineers configured the network core with Intermediate System to Intermediate System (IS-IS) as the Interior Gateway Protocol (IGP). Segment Routing (SR) capabilities were enabled to allocate and propagate Segment IDs (SID) with its respective MPLS label. Following two types of SID defined in the standard; Prefix/node SID and Adjacency SID are advertised based on IGP. Prefix/node SID is used for shortest path. Adjacency SID is used for explicit path. In out test scenario, Huawei used Prefix/node SID type SR as transport signalling protocol.

The scenario consisted of two customer sites, Site A and Site B, where the former was a multi-homed site in all active configuration and the latter a single-homed site.

A Link Aggregation Group (LAG) using LACP protocol was configured between DUT1 (NE9000-20) and DUT2 (NE9000-8) and load balancing was enabled based on destination MAC addresses.

During the test we generated Broadcast, Unknown Unicast and Multicast (BUM) traffic, and unicast traffic from TG3 and TG1. We used a custom IMIX pattern as specified in Table 1.

We verified that DUT 3 (NE9000-8) load balanced the unicast traffic between DUT1 and DUT2 as expected.

We also corroborated that DUT 1 (Huawei NE9000-20) was elected as the Designated Forwarder (DF) and it was responsible for forwarding BUM traffic to Site A. The non-DF router (DUT 2) simply dropped the BUM traffic coming from the core as defined in the standard.

During the test EANTC verified EVPN functionality worked as defined in the standard.

The test setup is depicted in Figure 10.

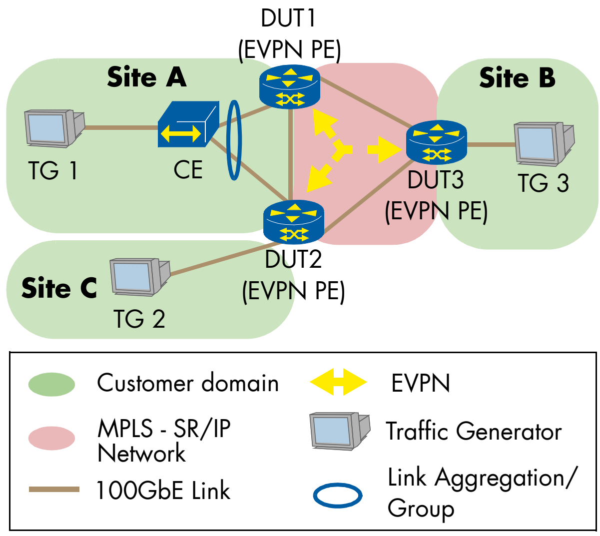

Figure 10: Test Setup — EVPN

EVPN with VXLAN Encapsulation

In the second part of the EVPN test, we executed a similar test with VXLAN encapsulation over an IP-only core. The IGP and IP setup remained the same as in previous test, we only added a new Traffic Generator (TG2) to emulate a third site in the EVPN as shown in Figure 10.

We sent any-to-any bidirectional traffic with the same IMIX pattern as defined in the first part and we observed no packet loss.

We verified EVPN signaling by inspecting how MAC address were propagated over the core using BGP.

During the test, we analyzed attributes of BGP update messages; especially EVPN route type 2 and 3. Based on this analyzes we can conclude that the EVPN functionality with VXLAN worked as defined in the standard.

Next Generation Multicast VPN (NG-MVPN)

NG-VPN leverages the use of BGP to signal customer multicast sources and receivers for defined multicast groups and to signal the multicast distribution tree mechanism in the core of the network.

NG-MVPN uses RSVP-TE or mLDP signaling to create multicast distribution tree or P-tunnels.

In this test, Huawei chose RSVP-TE Point to Multi-Point (P2MP) signaling protocol to demonstrate Traffic Engineering. We configured two link administrative groups as shown in Figure 11.

By doing so, we forced the multicast traffic to flow through DUT3, where multicast traffic destined to TG3 and TG2 was replicated.

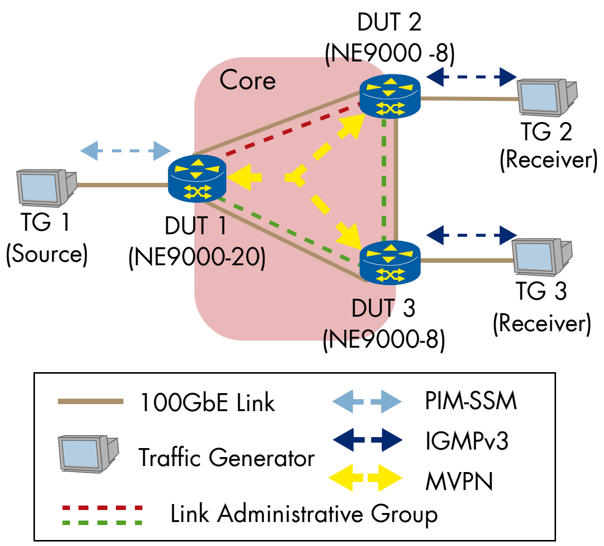

Figure 11: Test Setup — NG-MVPN Test Setup

PIM Source Specific Multicast (PIM-SSM) was used between TG1 and DUT1 and Internet Group Management Protocol version 3 (IGMPv3) between TG2 and DUT2, and between TG3 and DUT3.

➔ Demonstrated NG-MVPN functionality using RSVP-TE as P2MP signaling protocol

➔ Demonstrated Traffic Engineering functionality by configuring link administrative groups

During this test, the source sent a constant rate of 20 Mbit/s per multicast group with a fixed packet size of 1500 bytes. EANTC tested 100 Multicast groups, totalizing 2 Gbit/s of multicast traffic.

We observed no packet loss during the test and we verified the correct signaling of P-Tunnels and VPN multicast signaling.

Conclusion

During two weeks of the test campaign, EANTC verified the forwarding performance, power efficiency, scalability and functionality of Huawei's new 4T Line Processing Unit (LPU) installed in Huawei NE9000 series routers.

Initially we concluded the full line rate packet forwarding performance with various packet sizes and we verified the energy efficiency of the 4T line card. Once we concluded the performance of the line card, we turned our attention to the control plane scalability. We verified the maximum capacity of FIB and BGP routing entries as Huawei claimed.

Finally, we verified the core router features such as Virtual System Separation, Equal Cost Multi-Path, EVPN, Segment Routing, VXLAN, and Next Generation Multicast VPN. All tested features functioned as expected.

EANTC AG
Salzufer 14, 10587 Berlin, Germany
info@eantc.com, http://www.eantc.com/
[v1.7 20180219]