# EANTC Independent Test Report

Huawei NetEngine 8000 X-Series Router
With 19.2T Line Processing Unit (LPU)
Performance, Advanced Functionality and Security Feature Test

February 2024

EANTC

HUAWEI

## Contents

## Executive Summary

EANTC independently evaluated the performance, service scalability, and energy efficiency of the new Huawei 19.2 Terabit/s line card for the NetEngine 8000 X-series routers. We found that the 19.2T line card lives up to the expectations of its name – it supports a staggering 19.2 Terabit/second line rate throughput across its high-density configuration with 36x 400GigabitEthernet ports and 6x 800GigabitEthernet ports. It takes less than 2,200 Watt electrical energy while running at full load across all ports.

In our test, the NE8000 router and the 19.2T line card exhibited versatile feature support of Segment Routing over IPv6 (SRv6) with compressed segment IDs (SIDs), Ethernet VPN services, interworking with legacy MPLS services, and more. We witnessed large-scale routing and forwarding database (RIB/FIB) support, and robust protection against BGP peering overload.

Finally, EANTC validated extensive security functions of the NetEngine 8000 software, protecting the router's management interface and routing protocols (BGP, ISIS, OSPF) from intrusions and denial of service attacks.

The 19.2T line card marks another milestone in Huawei's series of line processing units (LPUs) for the NetEngine 8000 X-series router family, well combining performance, energy efficiency, service scalability, extensive feature support and security functions.

### Key Findings—19.2T Line Card

☑ 19.2 Tbit/s throughput; 294 Million packets per second (Mpps) forwarding speed

☑ Power use between 1,920 W (idle) and 2,210 W (full throughput) per line card

☑ Ethernet VPN (EVPN) service termination support for L3VPNs, VPWS over SR-TE and VPWS over SRv6-BE

☑ SRv6 routing with up to 10 SIDs

☑ E-Tree service support over SRv6 with 16-bit µSIDs and 32-bit µSIDs

☑ Service-level interworking between MPLS L3VPNs and EVPN L3VPNs over SRv6; between MPLS "Martini" VPLS and EVPN SRv6 VPLS using 16-bit µSIDs

☑ FIB capacity 6.3 Million IPv4+IPv6 routes, RIB capacity at least 120 Million IPv4+IPv6 routes

☑ Protection against BGP processing overload

☑ Versatile BGP, ISIS, OSPF, and LDP protocol protection

### Key Findings—Other line card/ NetEngine 8000 router and iMaster NCE in general:

☑ Power saving of 66 W through SFU warm-backup, 220 W per 4T LPU through slice sleep mode

☑ Risk assessment of management access

## Introduction

EANTC (European Advanced Networking Test Center), an independent test lab based in Berlin (Germany), was commissioned by Huawei to verify the functionality and performance of the new 19.2 Terabit/s line card for the family of NetEngine 8000 X-series routers. In short, this line card is called the "19.2T" model. The tests were conducted in January 2024 at Huawei premises in Beijing, China.

Our vendor-independent router tests focused on performance and advanced functionality validations needed for today's use case scenarios required by service providers, large enterprises, and hyperscalers. These requirements have tightened compared with previous router generation tests, and include:

- A much more differentiated ecosystem of mobile and fixed network applications;
- Much stricter real-time SLAs for enterprise and mobile traffic;
- More efficient network operations, including provisioning, troubleshooting, and optimization;
- Much better use of redundant infrastructure for load balancing;
- An evolved set of standards in Segment Routing over IPv6 (SRv6) and Software-Defined Networking (SDN).

Due to these tightened requirements, each new line card generation must provide an incremental throughput upgrade and an extensive feature set improvement.

The test plan used in this project was developed with the whole industry in mind. It is not specific to Huawei solutions. When service providers issue RfPs for core and service edge routers, we recommend evaluating the topics below—no matter which vendors are involved.

## Test Environment Setup

All tests were conducted in a lab scenario representing a minimal network design required to execute all tests. Four NetEngine 8000 series X4 and X8 routers were included in the test bed. The lab scenario was constructed to enable the execution of all test cases, including redundancy failover.

Figure 1 below shows the details of the test bed. Two routers were equipped with the new 19.2T line card under test, facilitating 800GigabitEthernet (800GbE) links. Two additional routers were equipped with 4T line cards using 400GbE and 100GbE links. All routers were configured with IS-IS and BGP. Redundancy failover tests were carried out between the routers marked with the green triangle in the diagram.

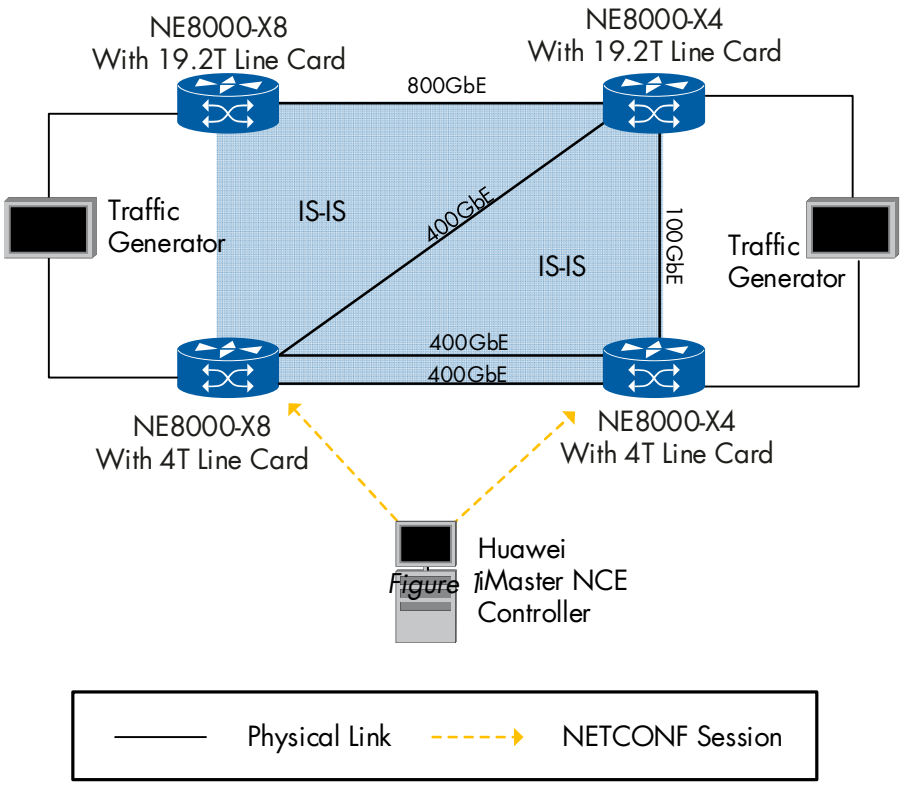


*Figure 1: Test Bed Setup*

Some of the test cases included network management aspects. We used the Huawei iMaster NCE-IP for these test cases. The NCE-IP managed the two routers using 4T line cards; the 19.2T line card was not involved in the network management tests. Finally, Spirent TestCenter traffic generators were connected to all four routers.

## Hardware Details

The table below lists all hardware and software versions of equipment used in the test. All routers under test and the iMaster NCE-IP management solution used production software available to customers.

| Device | Software | Line Cards |
|---|---|---|
| NetEngine 8000 X8 | V800R023C10 | 19.2T LPU and 4T LPU |
| NetEngine 8000 X4 | V800R023C10 | |
| iMaster NCE | V100R022C10 | |

*Table 1: Network Element and Controller Versions*

As traffic generator, we used two Spirent TestCenter units provided by Huawei. The 400GbE-enabled system used software version 5.13, the 100GbE system used software 5.11.

## Traffic Parameters

As shown in Table 2 below, EANTC defined a specific IPv4 and IPv6 traffic mix (IMIX) based on the CAIDA UCSD Internet statistics. This mix uses ten individual frame sizes varying from 78 to 9018 bytes, producing an average frame size of 553.2 bytes.

This IMIX presents a more significant challenge to network devices than a standard IMIX, as it more accurately reflects real-world Internet traffic patterns regarding the selected frame sizes and their respective proportions.

| Frame Size | Weight |
|---|---|
| 78 Bytes | 5243 |
| 150 | 861 |
| 314 | 273 |
| 486 | 233 |
| 575 | 230 |
| 970 | 127 |
| 1028 | 151 |
| 1518 | 2882 |
| 1865 | 1 |
| 9018 | 1 |

*Table 2: IMIX Traffic Definition*

## Test Results

The main section of the report describes all test cases in detail: Their rationales, setups, test steps, expected results, and verdicts. These descriptions help to understand advanced configuration use cases – specifically related to SRv6 practical network designs, power efficiency aspects, and router access security.

Good news first: The Huawei solution passed all functional test cases. Obviously, the performance tests have more fine-grain results, which also reached or exceeded Huawei's claims.

Now, on to the details:

## Forwarding Performance of 19.2T Line Card

Forwarding performance tests are fundamental for understanding and ensuring router efficiency, capacity, and reliability, critical for maintaining a robust and high-performing network infrastructure.

The test aimed to assess a router's most common stress points, including ingress packet buffer capacity, packet classification efficiency, and overall performance regarding look-ups and processing speed. Understanding these stress points was crucial for developing a targeted test methodology focused on these aspects.

The best setup would have been to connect all line card ports one-on-one with traffic generator ports. However, given the large number of 400GbE interfaces, the available traffic generator ports were insufficient. Additionally, no traffic generators with 800GbE ports were available for this project.

As a result, we implemented a daisy chain setup to fully utilize all 42 ports of the 19.2T line card. This involved a configuration where six ports from the traffic generator were connected to the line card under test. Virtual Routing and Forwarding (VRF) instances were configured to manage and route traffic through this setup. Each VRF included two ports and was set up with static routes to enable controlled traffic forwarding between these VRFs. Physical links were used to connect the VRFs to complete the daisy chain configuration.

Huawei configured the Spirent test equipment to simulate 60,000 dual-stack devices using six 400GbE ports. This setup generated the same number of MAC addresses for IPv4 ARP entries and an equally populated IPv6 neighbor table, each containing 60,000 entries.

**☑ The 19.2T line card achieved a throughput of 87,223,737 packets per second (pps) per port. The line card reached an impressive line rate throughput of 19.2 Tbit/s.**

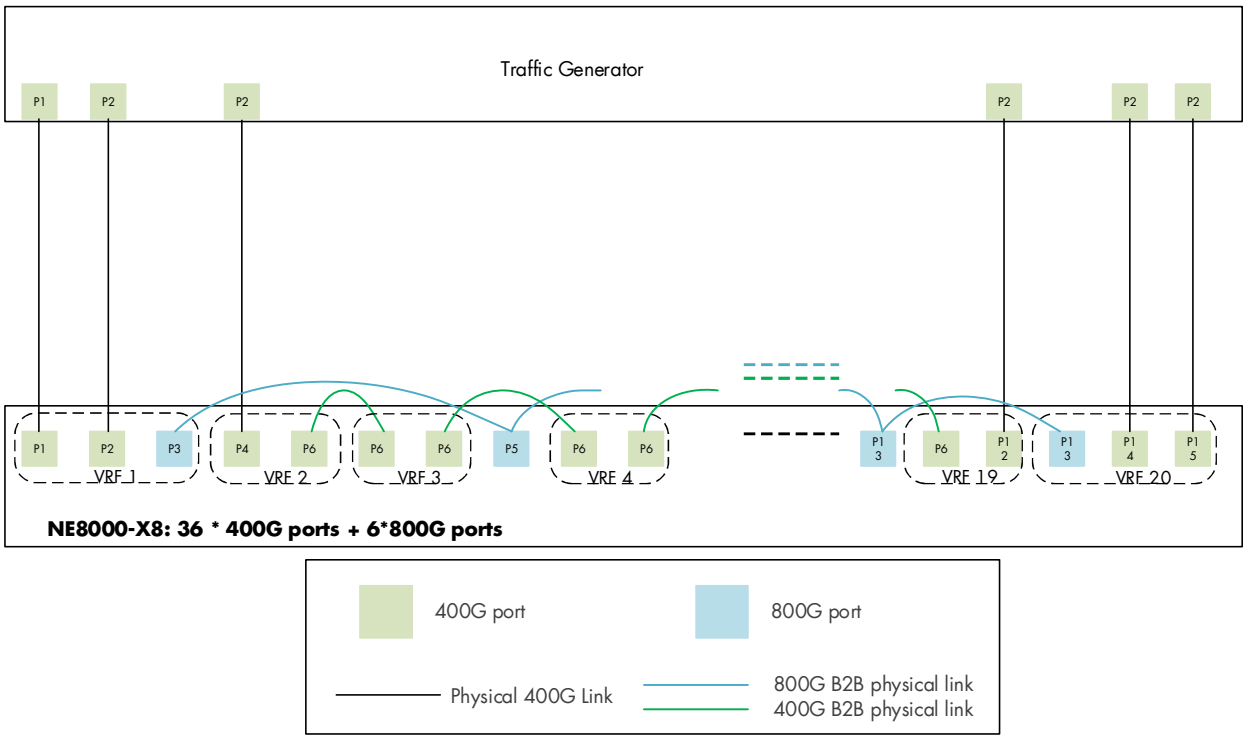


*Figure 2*

InUti/OutUti: input utility/output utility

| Interface | PHY | Protocol | InUti | OutUti |
|---|---|---|---|---|
| 400GE1/0/1 | up | up | 99.99% | 99.99% |
| 400GE1/0/2 | up | up | 99.99% | 99.99% |
| 400GE1/0/3 | up | up | 100% | 99.99% |
| 400GE1/0/4 | up | up | 99.99% | 99.99% |
| 400GE1/0/6 | up | up | 99.99% | 99.99% |
| 400GE1/0/7 | up | up | 100% | 99.99% |
| 400GE1/0/8 | up | up | 100% | 100% |
| 400GE1/0/9 | up | up | 99.99% | 99.99% |
| 400GE1/0/11 | up | up | 100% | 99.99% |
| 400GE1/0/12 | up | up | 100% | 100% |
| 400GE1/0/13 | up | up | 100% | 100% |
| 400GE1/0/14 | up | up | 99.99% | 99.99% |
| 400GE1/0/15 | up | up | 99.99% | 99.99% |
| 400GE1/0/16 | up | up | 99.99% | 99.99% |
| 400GE1/0/17 | up | up | 99.99% | 99.99% |
| 400GE1/0/18 | up | up | 99.99% | 99.99% |
| 400GE1/0/19 | up | up | 99.98% | 99.98% |
| 400GE1/0/20 | up | up | 99.99% | 99.99% |
| 400GE1/0/22 | up | up | 100% | 100% |
| 400GE1/0/23 | up | up | 99.99% | 99.99% |
| 400GE1/0/24 | up | up | 100% | 99.99% |
| 400GE1/0/25 | up | up | 100% | 99.99% |
| 400GE1/0/27 | up | up | 99.99% | 99.99% |
| 400GE1/0/28 | up | up | 99.99% | 99.99% |
| 400GE1/0/29 | up | up | 100% | 100% |
| 400GE1/0/30 | up | up | 100% | 99.99% |
| 400GE1/0/32 | up | up | 100% | 99.99% |
| 400GE1/0/33 | up | up | 99.99% | 99.99% |
| 400GE1/0/34 | up | up | 99.98% | 99.98% |
| 400GE1/0/35 | up | up | 99.99% | 99.99% |
| 400GE1/0/36 | up | up | 99.99% | 99.99% |
| 400GE1/0/37 | up | up | 100% | 99.99% |
| 400GE1/0/38 | up | up | 100% | 100% |
| 400GE1/0/39 | up | up | 99.99% | 99.99% |
| 400GE1/0/40 | up | up | 99.99% | 99.99% |
| 400GE1/0/41 | up | up | 99.99% | 99.99% |
| 800GE1/0/0 | up | up | 100% | 99.99% |
| 800GE1/0/5 | up | up | 99.99% | 99.99% |
| 800GE1/0/10 | up | up | 99.99% | 99.99% |
| 800GE1/0/21 | up | up | 100% | 100% |
| 800GE1/0/26 | up | up | 99.99% | 99.99% |
| 800GE1/0/31 | up | up | 99.99% | 99.99% |

*Figure 3*

Figure 3 shows the full utilization of all the line card ports (36*400GbE ports + 6*800GbE ports).

The test scope included IPv4 and IPv6, as well as a dual-stack scenario where IPv4 and IPv6 traffic were split 50:50. Each of these test cases — IPv4, IPv6, and dual-stack — was run three times, with each iteration lasting for five minutes while using the IMIX developed by EANTC.

It was confirmed that there was no packet loss across all tests, an indicator of robust data handling and forwarding capabilities. Additionally, the device's memory and CPU utilization remained stable throughout the test durations. However, an interesting observation was noted regarding latency. Results from Spirent indicated that the latency for two ports was slightly higher than the other four. This difference was attributed to these two ports being connected to the 400GbE ports on Huawei's 19.2 line card and then routed through a snake topology, involving more hops than the others. The latency slightly increased during the test due to Spirent packet generator micro-bursts; these overloaded the fully saturated ports already running at line rate.

Huawei clarified this aspect, noting that Spirent generated minor bursts intermittently throughout the test duration while the line rate was utilized at 100%. This factor likely contributed to the slightly increased latency observed.

## 800GbE Port Service Functions

Huawei asked us to validate a range of modern transport services typically required by operators. It's important to note that the 800GbE ports support terminating these services, not only their transit. Typically, so-called "provider routers" implement limited service functions on high-speed ports to simplify the line card hardware. Huawei has chosen to support service termination from day one, although these 800GigabitEthernet ports are unlikely to be used for service termination today. 800GbE is the latest core transport technology now, but it will become more of a standard type of interface a few years ahead.

**☑ The 19.2T line card proved support of Ethernet VPN (EVPN) service termination for L3VPNs, for virtual private wire service (VPWS) over traffic-engineered Segment Routing (SR-TE) and VPWS over SRv6 best effort mode (SRv6-BE). Tests were run using 400GbE and 800GbE ports on the 19.2T line card.**

### EVPN VPWS

EVPN VPWS (Ethernet VPN Virtual Private Wire Service) is a standard network technology that provides point-to-point Ethernet services over a Multiprotocol Label Switching (MPLS) or Segment Routing network.

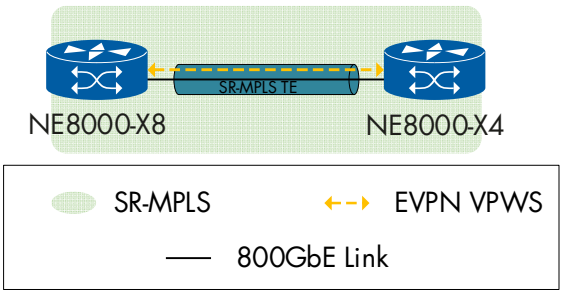


*Figure 4*

The Huawei team prepared the underlay, configuring IS-IS, and segment routing. Additionally, Huawei configured a BGP EVPN neighbor relationship between DUT1 and DUT2 to use the extended reachability information to implement control-plane (instead of data-plane) MAC address learning and advertisement.

Both devices were configured with an EVPL instance and an EVPN VPWS instance and associated with an SR-TE tunnel policy.

Figure 5 confirms the service status was up over the 800 GbE interface (interface connected to Spirent) using tunnel-type SR-TE.

```
[~NE8000-X8-1]disp bgp evpn evpl
Total EVPLs: 1       1 Up      0 Down

EVPL ID : 501
State : up
EVPL Down Cases : --
Evpl Type : none
Interface : 400GE6/0/13.50
Interface Status : up
Bridge-domain : --
Ignore AcState : disable
Local MTU : 9600
Local Control Word : false
Local Redundancy Mode : all-active
Local DF State : primary(-)
Local ESI : 0000.0000.0000.0000.0000
Local SID : ::
Remote Redundancy Mode : all-active
Remote Primary DF Number : 1
Remote Backup DF Number : 0
Remote None DF Number : 0
Peer IP : 100.0.0.3
 Origin Nexthop IP : 100.0.0.3
 Remote SID: ::
 DF State : primary
 Eline Role : primary
 Remote MTU : 9600
 Remote Control Word : false
 Remote ESI : 0000.0000.0000.0000.0000
 Tunnel info : 1 tunnels
  NO.0  Tunnel Type : sr-te, Tunnel ID : 0x000000000300000001
Last Interface UP Timestamp: 2024-01-16 06:52:36+00:00
Last Designated Primary Timestamp: 2024-01-16 06:52:36+00:00
Last Designated Backup Timestamp: --
Last EVPL UP Timestamp: 2024-01-16 06:53:04+00:00
Last EVPL DOWN Timestamp: --
[~NE8000-X8-1]
```

*Figure 5*

Finally, the test was verified by generating and analyzing bi-directional EPVN traffic at an 80% line rate. The absence of packet loss was a key indicator of a stable and well-configured service.

## EVPN L3VPN

EVPN can utilize an MPLS L3VPN for data plane traffic forwarding, a feature commonly referred to as EVPN L3VPN. In a network structured around EVPN L3VPN, the VPN routes are propagated using IP prefix advertisement routes, making it capable of delivering both Layer 2 and Layer 3 VPN services and enhancing the versatility of network operations.

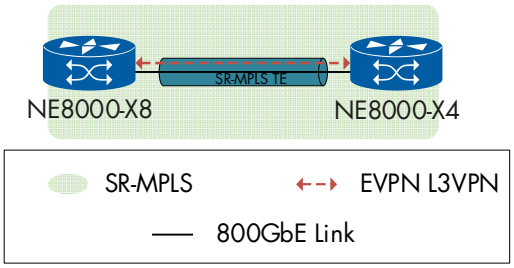


*Figure 6*

We configured two L3VPN instances on each router and enabled the VPN instance to advertise EVPN IP prefix routes.

An MP-IBGP peer relationship was established between the devices, and it was confirmed that the VPN routes were successfully incorporated into the EVPN routing table. A tunnel policy was also configured on the PEs, ensuring that EVPN L3VPN is selectively utilized for an SR-MPLS TE tunnel as the designated public tunnel.

Figure 7 shows that the router learned the VPN instance route from the traffic generator to the remote-emulated device.

We generated IPv4 bi-directional traffic over these services with 80% of the 400GbE ports without packet loss.

## EVPN VPWS over SRv6-BE

In our test, we explored the implementation of EVPN VPWS SRv6 BE (Best Effort mode). SRv6 BE utilizes the optimal SRv6 path determined by the IGP's (Interior Gateway Protocol) SPF (Shortest Path First) algorithm, simplifying the network configuration by using a single service SID (Segment ID) for packet forwarding across links. One of the key advantages of this setup is its straightforward configuration process, which eliminates the need for an external controller, relying instead on the IGP's SPF algorithm to calculate the paths.

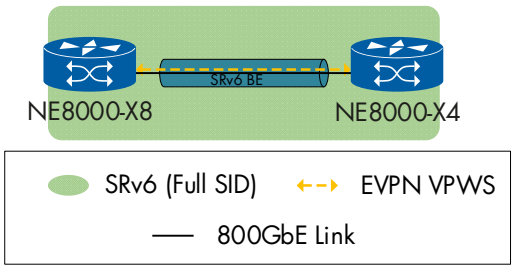


*Figure 8*

In our test, as depicted in Figure 8, we demonstrated a single-homed EVPN VPWS setup. The BGP EVPN address family was pre-deployed with specific instructions to use a VPN SID and SRv6 Best Effort forwarding. Additionally, we configured the system to generate the locator and End automatically.DX2 SID. The End.DX2 SID is critical in handling the decapsulation and Layer 2 cross-connect to an outbound interface, streamlining the operation within the EVPN VPWS environment.

```
[~X4-0]disp ip routing-table vpn-instance l3evpn
Route Flags: R - relay, D - download to fib, T - to vpn-instance, B - black hole route
------------------------------------------------------------------------------
Routing Table : l3evpn
         Destinations : 6        Routes : 6

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      70.10.1.0/24  IBGP    255  0           RD  100.0.0.4       Tunnel1
      70.20.1.0/24  Direct  0    0           D   70.20.1.1       400GE4/0/8.70
      70.20.1.1/32  Direct  0    0           D   127.0.0.1       400GE4/0/8.70
    70.20.1.255/32  Direct  0    0           D   127.0.0.1       400GE4/0/8.70
      127.0.0.0/8   Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
[~X4-0]disp ip routing-table vpn-instance l3evpn-2
Route Flags: R - relay, D - download to fib, T - to vpn-instance, B - black hole route
------------------------------------------------------------------------------
Routing Table : l3evpn-2
         Destinations : 6        Routes : 6

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      71.10.1.0/24  IBGP    255  0           RD  100.0.0.4       Tunnel1
      71.20.1.0/24  Direct  0    0           D   71.20.1.1       400GE4/0/8.71
      71.20.1.1/32  Direct  0    0           D   127.0.0.1       400GE4/0/8.71
    71.20.1.255/32  Direct  0    0           D   127.0.0.1       400GE4/0/8.71
      127.0.0.0/8   Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
[~X4-0]
```

*Figure 7*

```
[~HUAWEI]disp bgp evpn peer

 BGP local router ID : 61.6.11.1
 Local AS number : 65001
 Total number of peers : 1                 Peers in established state : 1

  Peer                    V          AS  MsgRcvd  MsgSent  OutQ  Up/Down       State  PrefRcv
  2001::2                 4       65001       19       19     0 00:13:16 Established        1
[~HUAWEI]disp bgp evpn evpl
Total EVPLs: 1       1 Up      0 Down

EVPL ID     : 1
State : up
EVPL Down Causes : --
Evpl Type : none
Interface : 400GE6/0/13.1
Interface Status : up
Bridge-domain : --
Ignore AcState : disable
Local MTU : 1500
Local Control Word : false
Local Redundancy Mode : all-active
Local DF State : primary(-)
Local ESI : 0000.0000.0000.0000.0000
Local SID : 2001:DB8:100::1:0:3F
Remote Redundancy Mode : all-active
Remote Primary DF Number : 1
Remote Backup DF Number : 0
Remote None DF Number : 0
Peer IP : 2001::2
 Origin Nexthop IP : 2001::2
 Remote SID: 2001:DB8:200::1:0:3F
 DF State : primary
 Eline Role : primary
 Remote MTU : 1500
 Remote Control Word : false
 Remote ESI : 0000.0000.0000.0000.0000
 Tunnel info : 1 tunnels
  NO.0   Tunnel Type : srv6-be, Tunnel ID :
Last Interface UP Timestamp: 2024-01-16 07:50:22+00:00
Last Designated Primary Timestamp: 2024-01-16 07:50:22+00:00
Last Designated Backup Timestamp: --
Last EVPL UP Timestamp: 2024-01-16 07:51:55+00:00
Last EVPL DOWN Timestamp: --
[~HUAWEI]disp se
[~HUAWEI]disp seg
[~HUAWEI]disp segment-routing ipv6 local-sid end-dx2 fo
[~HUAWEI]disp segment-routing ipv6 local-sid end-dx2 forwarding

                My Local-SID End.DX2 Forwarding Table
                -------------------------------------

SID        : 2001:DB8:100::1:0:3F/128          FuncType    : End.DX2
EVPL ID    : 1                                 LocatorID   : 5
LocatorName: dut1                              SidCompress : NO
Flavor     : NO-FLAVOR
UpdateTime : 2024-01-16 07:50:47.383

Total SID(s): 1
[~HUAWEI]
```

*Figure 9*

We verified that the service was operational with the correct tunnel type and Function Type (End.DX2), while the ESI value of 0 indicated a single-homing configuration. Additionally, the traffic generated over the VPWS service experienced no packet loss.

## SRv6 & VPN Service Tests

As the next step, we focused on advanced segment routing service functionality. EANTC has tested Huawei's implementations of SRv6, the segment routing over the IPv6 family of standards, in multiple engagements previously (SDN interoperability events 2023 and before [1]; Huawei Intelligent IP & Cloud Network Solution 2020/2021 [2], etc.). This time, we focused on complementary tests of segment ID scaling and compression and legacy MPLS interworking functions.

### SRv6 Maximum SID Depth

Maximum SID Depth (MSD) represents the limit on the number of SIDs that can be included in the Segment Routing Header (SRH) of an IPv6 packet. This parameter is crucial for ensuring that the packet headers do not exceed the processing capabilities of the network devices along the path.

Figure 10 shows the topology used to verify MSD SID depth while setting up SRv6 policies on the ingress nodes. On the first device under test (DUT1), the policy directed traffic towards DUT2 through a precisely defined path, looping multiple times between DUT2 and DUT3, utilizing 128-bit non-compressed End SIDs to test the network's handling of extensive SRv6 paths.

In contrast, DUT2 was configured with a policy targeting DUT1 through DUT3.

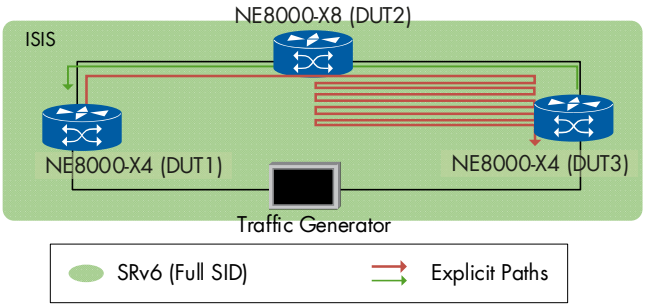


*Figure 10*

By utilizing port mirroring on the routers, we successfully intercepted the traffic for analysis. We examined the IS-IS LSP packets to verify the advertised maximum SID depth values within the Router Capability (Figure 11).

```
Node Maximum SID Depth (t=23, l=12)
    MSD Type: Base MPLS Imposition (1)
    MSD Value: 10
    MSD Type: Maximum Segments Left (41)
    MSD Value: 10
    MSD Type: Maximum End Pop (42)
    MSD Value: 11
    MSD Type: Unknown (43)
    MSD Value: 10
    MSD Type: Maximum H.Encaps (44)
    MSD Value: 10
    MSD Type: Maximum End D (45)
    MSD Value: 11
```

*Figure 11*

Additionally, the packets generated from DUT1 to DUT2 carried the completed list of locator SID of the explicit path in their Routing Header (Figure 12).

```
Internet Protocol Version 6, Src: 2001:db8:1::1, Dst: 2001:db8:300::333
    0110 .... = Version: 6
  > .... 0000 0000 .... .... .... .... = Traffic Class: 0x00 (DSCP: C
    .... 0100 1010 1001 1100 0100 = Flow Label: 0x4a9c4
    Payload Length: 678
    Next Header: Routing Header for IPv6 (43)
    Hop Limit: 255
    Source Address: 2001:db8:1::1
    Destination Address: 2001:db8:300::333
  Routing Header for IPv6 (Segment Routing)
      Next Header: IPIP (4)
      Length: 22
      [Length: 184 bytes]
      Type: Segment Routing (4)
      Segments Left: 10
      Last Entry: 10
      Flags: 0x00
      Tag: 0000
      Address[0]: 2001:db8:200::1:0:2
      Address[1]: 2001:db8:200::222
      Address[2]: 2001:db8:300::333
      Address[3]: 2001:db8:200::222
      Address[4]: 2001:db8:300::333
      Address[5]: 2001:db8:200::222
      Address[6]: 2001:db8:300::333
      Address[7]: 2001:db8:200::222
      Address[8]: 2001:db8:300::333
      Address[9]: 2001:db8:200::222
      Address[10]: 2001:db8:300::333
```

*Figure 12*

☑ **The NetEngine 8000 supports SRv6 routing headers with up to ten uncompressed segment IDs (SIDs), which is sufficient for very complex policy routing scenarios.**

### E-Tree over SRv6 Policy with 16-bit Compression

E-Tree services construct point-to-multipoint EVPN services, for example for multicast traffic. On the segment routing transport layer, E-Tree services face specific implementation challenges related to the root and leaves of the tree communication structure.

Generally, these challenges have been solved and implemented successfully in SRv6. In this test, Huawei asked us to verify that the E-Tree service over SRv6 can also be implemented with compressed segment IDs. As described in detail in our 2023 multi-vendor SDN interoperability report [1], SID compression reduces the overhead of SRv6 segment information in each packet significantly.

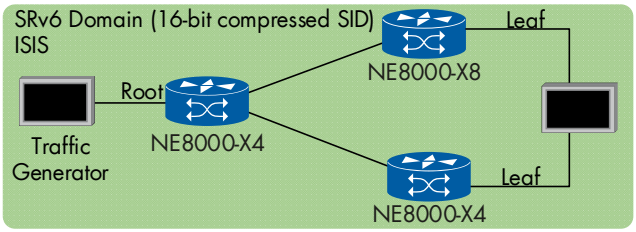


*Figure 13*

We established a 16-bit compression SRv6 policy targeting the respective routers and orchestrated an EVPN E-tree service over this SRv6 policy, carefully defining the Root and Leaves according to our network setup (Figure 13). We utilized an SR policy with a segment list comprising multiple compressed SIDs to optimize our test further.

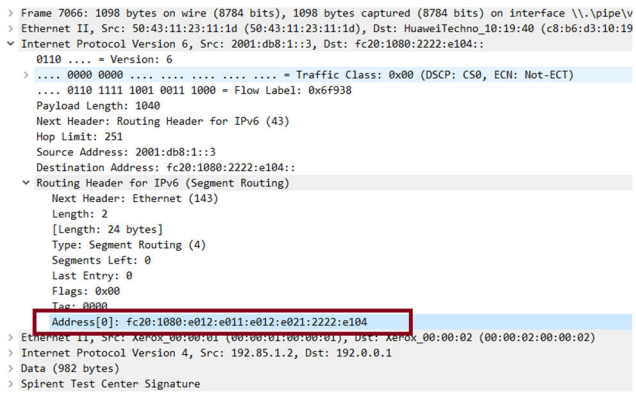


*Figure 14*

This configuration allowed us to demonstrate, through packet capturing, the effectiveness of the compressed 16-bit SID list in streamlining the SRv6 header, thereby enhancing the efficiency of our network routing (Figure 14).

In our test scenario, we verified the behavior of the E-tree service by conducting a series of targeted traffic tests:

**Unicast Traffic between Root and Leaves:** To validate the fundamental E-tree functionality, we initiated unicast traffic between the root (DUT1) and the leaves (DUT2 and DUT3). This test ensured that the root could send and receive unicast traffic to and from each leaf. The successful transmission of these packets confirmed that the E-tree service was correctly routing unicast traffic from the root to the leaves and vice versa.

**Unicast/Broadcast Traffic between Leaves:** A critical aspect of the E-tree service is that leaves should not communicate directly. To validate this, we sent unicast/broadcast traffic between the leaves (DUT2 and DUT3), and as expected, the traffic was blocked (Figure 15).

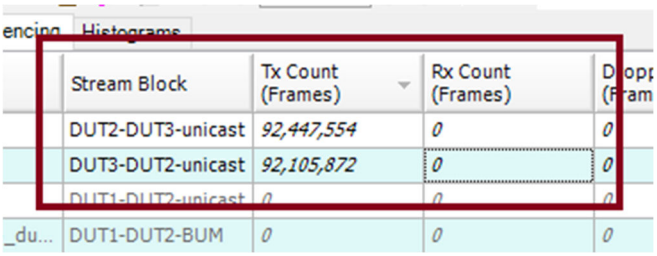


*Figure 15*

**Broadcast, Unknown Unicast, and Multicast (BUM) Traffic:** We also tested the behaviour of Broadcast, Unknown Unicast, and Multicast (BUM) traffic within the E-Tree topology. Sending BUM traffic from the root allowed us to observe its distribution to all leaves. On the other hand, sending BUM traffic from a leaf should result in the traffic being forwarded to the root and not to other leaves, adhering to the E-Tree principles.

| Stream Block | Tx Count (Frames) | Rx Count (Frames) |
|---|---|---|
| DUT1-DUT2-BUM | 145,102,547 | 290,205,094 |
| DUT2-DUT1-BUM | 580,322,448 | 580,322,448 |

*Figure 16*

We generated bidirectional BUM traffic between a root node and a leaf node. The receivers in this setup were the other nodes present in the topology. We observed that the BUM traffic originating from the root was received by both leaf nodes, as indicated by the doubled count of received frames. However, when the BUM traffic was sent from the leaf node, it was only

received at the root port and not at the other leaf node (Figure 16).

☑ **EANTC validated that the NetEngine 8000 router supports E-Tree services over SRv6 with compressed SIDs.**

### L3VPN over MPLS Interworking EVPN L3VPN over SRv6

While the future belongs to Segment Routing-driven networks, MPLS technology will not disappear any time soon. The coexistence of MPLS and SR networks requires advanced interworking solutions. One of the main goals is to interconnect VPN services between the MPLS and SR domains without an IP router or Ethernet "stitching" entity in between. In this case, EANTC tested the interconnection of MPLS L3VPN services with SRv6 EVPNs.

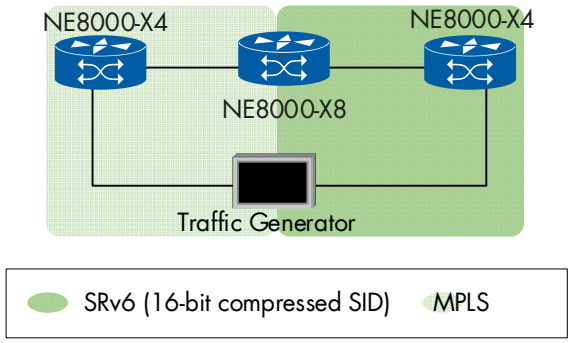


*Figure 17*

We used three devices in our test setup (Figure 17). ISISv4 and MPLS LDP were established on DUT1 and DUT2 to facilitate the foundational MPLS infrastructure, followed by configuring an L3VPN service over MPLS LDP on these devices.

The configuration process moved forward with the setup of ISISv6 on DUT2 and DUT3, which enabled the propagation of the SRv6 locator through the ISISv6 protocol. SRv6 policies incorporating 16-bit compression were then implemented on DUT2 and DUT3. This step was followed by establishing an EVPN L3VPN over SRv6 policy between DUT2 and DUT3.

The border node (DUT2) was enabled to advertise the routes re-originated in the BGP VPNv4/VPNv6 address family to the BGP VPN IPv6 peer and vice versa (Figure 18).

```
ipv4-family vpnv4
 undo policy vpn-target
 peer 100.0.0.2 enable
 peer 100.0.0.2 reflect-client
 peer 100.0.0.2 import reoriginate
 peer 100.0.0.2 advertise route-reoriginated evpn mac-ip
 peer 100.0.0.2 advertise route-reoriginated evpn ip
 l2vpn-family evpn
 undo policy vpn-target
 peer 100.0.0.4 enable
 peer 2001:DB8:1::2 enable
 peer 2001:DB8:1::2 reflect-client
 peer 2001:DB8:1::2 advertise encap-type srv6
 peer 2001:DB8:1::2 import reoriginate
 peer 2001:DB8:1::2 advertise route-reoriginated vpnv4
```

*Figure 18*

We examined the routing table for the ingress nodes and confirmed that each endpoint was accessible either via an SRv6 policy (L3EVPN policy) or through an MPLS LDP tunnel.

The last step was generating bidirectional traffic over the configured services, and we confirmed zero packet loss (Figure 19).

| Tx Port | Rx Port | Stream Block | Sig Count (Frames) | Rx Count (Frames) |
|---|---|---|---|---|
| DUT3 | DUT1 | DUT3-DUT1 | 609,784,409 | 609,784,409 |
| DUT1 | DUT3 | DUT1-DUT3 | 610,202,637 | 610,202,637 |

*Figure 19*

☑ **Service-level interworking was confirmed between MPLS L3VPNs and EVPN L3VPNs over Segment Routing v6.**

### EVPN VPLS over SRv6 Coexisted with Martini VPLS over MPLS

VPN VPLS over SRv6 and Martini VPLS over MPLS are two distinct methodologies for implementing multi-point virtual private network services, each leveraging different underlying technologies. Martini VPLS over MPLS is a traditional approach that utilizes Label Distribution Protocol (LDP) for signaling.

We confirmed that the VPLS over MPLS sessions were active between two sets of device pairs. In contrast, the VPLS over the SRv6 session was operational between the other device pair using the SRv6 policy with 16-bit compression. Each device successfully populated its table with the MAC addresses of all nodes, demonstrating effective MAC address learning across all routers.

We initiated bi-directional traffic across the EVPN services spanning all nodes and verified the absence of packet loss, ensuring seamless connectivity (Figure 21).

☑ **The function of service-level interworking between MPLS "Martini" VPLS and EVPN SRv6 VPLS services was confirmed. 16-bit compressed Segment IDs were used on the SRv6 side of the interworking service.**

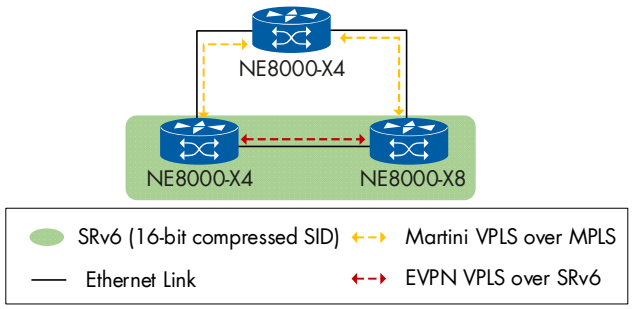


*Figure 20*

| Stream Block | Rx Count (Frames) | Sig Count (Frames) | T |
|---|---|---|---|
| DUT1-DUT2 | 224,738,119 | 224,738,119 | 1 |
| DUT1-DUT3 | 224,738,119 | 224,738,119 | 1 |
| DUT2-DUT1 | 224,508,454 | 224,508,454 | 1 |
| DUT2-DUT3 | 224,508,453 | 224,508,453 | 1 |
| DUT3-DUT1 | 224,601,882 | 224,601,882 | 1 |
| DUT3-DUT2 | 224,601,881 | 224,601,881 | 1 |

*Figure 21*

### EVPN VPWS over SRv6 Policy with 32-bit Compression

We conducted this test for VPWS over SR-TE with 32-bit compression SID, driven by the specific needs of some service providers and operators. Their strict address planning requirements consider using 32-bit compression (instead of 16-bit compression) as a convenient and logical choice for efficient network management and service delivery.

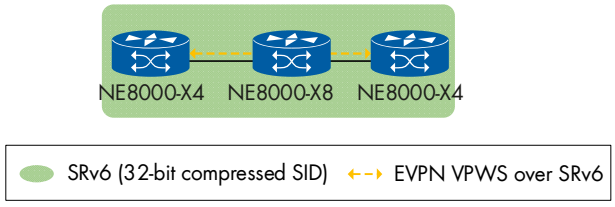


*Figure 22*

**☑ The NetEngine 8000 also supports 32-bit compressed Segment IDs in SRv6, validated with EVPN point-to-point services.**

## 19.2T Line Card Power Consumption

In an era where network infrastructure is expanding and energy efficiency is increasingly emphasized, testing the power consumption of a router is critical.

Routers, especially in large networks, can consume significant amounts of electricity. Understanding and managing power consumption is crucial for controlling operational costs, particularly in data centers and large enterprise environments.

We reused the topology from the forwarding performance test (Figure 1) while using a clamp-on ampere meter to perform this test. The precision of this basic test device was limited and not precisely defined; we assume a measurement error of ± 2 %. A more precise power meter was unavailable since the test was not conducted at the EANTC lab.

To reflect a realistic operating scenario for the router, we incorporated a substantial background load of routing information, adding 4 million IPv4 BGP routes and 2 million IPv6 routes. Additionally, Huawei set the NetEngine 8000 fans to a fixed speed of 65% to guarantee a high and reproducible load from the fans. The ambient temperature in the lab was between 25 and 28 degrees Celsius during the test; the humidity in the lab could not be measured, but the environmental humidity outside was 25 % at the time of the test.

First, we measured the power usage of the router without the 19.2T line card – only measuring the essential parts of the router (processing module, fans, power supplies, etc.). Then, we measured the power consumption when the 19.2T line card was inserted and configured but idle otherwise. The 19.2T line card alone took 1.920 Watts ± 40 W.

Then, we generated a bidirectional traffic dual-stack IPv4:IPv6 ratio of 50:50 at 99.5% load for six ports. Utilizing a daisy chain topology, we effectively increased this to a total throughput of 19.2 Tbit/s.

As expected, the power consumption increased compared to the basic reading without traffic. At full load, the power consumption was 290 Watts ± 6 W higher (approximately 4.3 %).

We also tested the line card at 50% load; power usage was 150 W ± 3 W (approximately 2.3%) higher than the idle scenario.

**☑ The 19.2T line card took between 1,920 W (idle) and 2,210 W (full throughput) ± 40W of electrical energy.**

## Green and Energy-Saving Router Mechanisms

It is crucial to optimize hardware components such as the 19.2T line card to forward data with higher energy efficiency. Smaller chipset structures are a vital contribution to this goal, but there are technological limits. Two other factors can help reduce router operations' carbon footprint further: Optimizing software algorithms and creating advanced operational procedures. The two test cases in this section focus on the latter. Huawei asked us to verify two configuration options that help to switch off hardware components during low utilization periods.

### Switch Fabric Unit Warm-Backup Redundancy

In this test, we utilized the Huawei NetEngine 8000 X8 -2, featuring 8 LPUs (Line Processing Units), 8 SFUs (Switch Fabric Units), one DC power unit, and three fans, part of the NetEngine 8000 x8 series. In this router type, the 8 LPUs are interconnected to facilitate traffic exchange, as depicted in Figure 23. This setup aims to route traffic directly between the LPUs, managed by the SFUs, which were the focal point of this test. Note that Huawei used "4T" LPUs in this test, which have eight 100GigabitEthernet ports plus eight 400GigabitEthernet ports (total nominal capacity of 4 Terabit/s). EANTC tested the performance of these line cards in 2018 [3].

Huawei implements energy-saving features in its network components, focusing on reducing power consumption dynamically based on service requirements. Components that are not in use or when no services are active are either shut down or switched to sleep mode, especially during periods when users are offline.

To mirror real-world conditions, we configured 20,000 flows evenly split between native IPv4 and IPv6. No SRv6 or other advanced services were running on the router during this test case. We employed 16 ports, each at 100 Gbps, for the previously used *snake* topology (see forwarding performance test at the beginning of this report). This way, the router was loaded with a total of 1.6 Tbit/s traffic – compared to its nominal maximum of 320 Tbit/s in this configuration (8 x 40 x 100GbE). Effectively, the router ran at 0.5 % of its nominal forwarding capacity.

Traffic generation began from ports in LPU 1, sequentially connecting to LPUs 2 through 8. At this time, we recorded the regular power consumption of the whole router via the power meter used in the previous test (see 19.2T Line Card Power Consumption). The total amount of power, 9010 W ± 180 W, served as a reference, with the goal of reducing it without impairing the routing and forwarding services.

Given the SFU's role in linking LPUs, Huawei engineers configured the SFU-Warm-Backup Function. With the deep energy-saving mode activated in the router configuration, the router randomly selected SFU-1 for Warm-Backup and kept the other six SFUs in basic mode. We measured the power consumption again and observed a 66 Watts drop (± 1 W).

To confirm the router's ability to adapt when certain components are not in use, we conducted a test where one SFU was intentionally (manually) shut down while another SFU1 was already in sleep mode. Upon shutting down the extra SFU, SFU1 successfully exited its warm-up backup status, demonstrating the router's efficient response and adjustment to changes in component utilization.

☑ **Huawei demonstrated 66 Watt power savings through the "SFU Warm-Backup" sleep function in a router fully deployed with eight SFUs, running at a 0.5 % nominal load.**

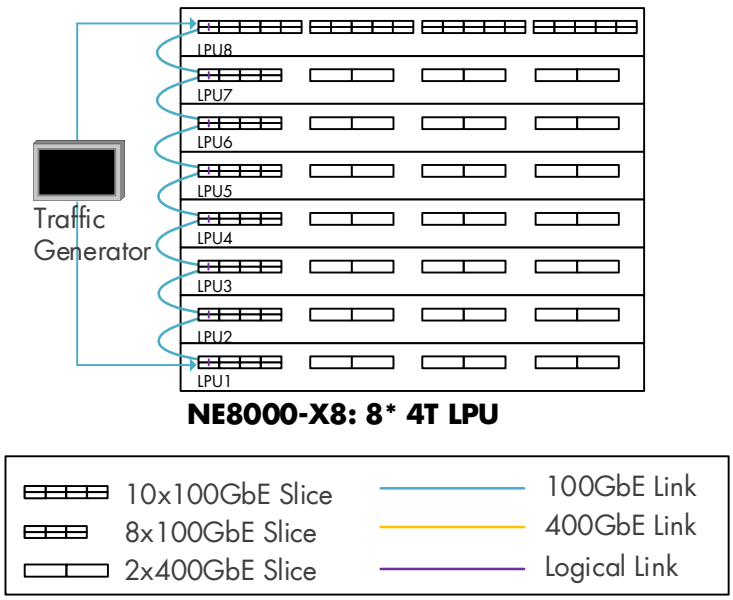


*Figure 23*

### 4T Line Card Slice Sleep

In a Line Processing Unit (LPU), each slice, essentially a chipset, plays a crucial role in forwarding traffic and is responsible for a set of ports. Huawei explained to EANTC that the 4T LPU (40x 100GbE ports) implements four slices, each servicing ten ports. Note that this test did not validate the new 19.2T line card.

Huawei further explained that a "slice sleep" mode has been implemented in the 4T LPUs to save energy for unused ports. To enter sleep mode, all ports connected to a slice must remain inactive for at least 30 minutes. That said, Slice 0 is an exception and cannot sleep due to its role in managing basic port-to-port communication.

To demonstrate the slice sleep mode, we measured the total power consumption at the start of the test. Before implementing the sleep mode, the router's power consumption was 9010 Watts ±180 W.

Huawei then configured slice sleep mode and manually shut down all the ports of one LPU(4T) of NE8000-X8. Subsequently, we waited 30 minutes for the slice sleep mode to take effect.

Once three slices had successfully entered sleep mode, the power meter indicated a reading of 8790 Watts ±176 W.

**☑ Huawei demonstrated a power-saving capability of approximately 220 W from one 4T LPU through the "slice sleep" mode when all LPU ports were shut down, translating to a substantial 25% reduction in power consumption at the LPU level.**

To benefit from this feature, operators should consider a) using the first ten ports of the 4T LPU for management service since the first slice cannot go to sleep mode anyway, and b) manually shutdown unused ports in chunks of ten so that the slice sleep function can kick in.

## 19.2T Line Card FIB/RIBv4 and FIB/RIPv6 Scalability

The Forwarding Information Base (FIB) and the Routing Information Base (RIB) are critical components in the architecture of network routers. Their size and scalability directly affect the router's ability to handle many routes efficiently.

In our testing environment, shown in Figure 24, we configured the setup by directly connecting the router to a traffic generator back-to-back. We started initiating background traffic, comprising IPv4 and IPv6 flows, each contributing 21.2 Gbps and totaling 100 flows. Then, we established an external BGP (eBGP) session between the tester and the router.

Various IPv4 and IPv6 routes were advertised during this session, featuring prefix lengths ranging from /8 to /32. To mimic real-world scenarios, these prefixes included a combination of consecutive and random addresses, comprehensively assessing the router's capabilities in handling routing scenarios.

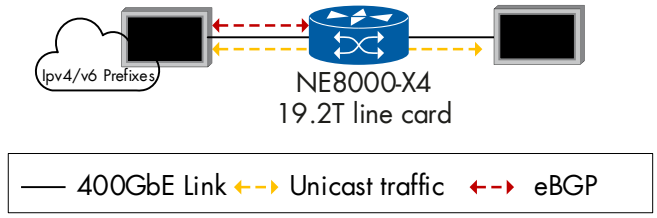


*Figure 24*

The test successfully pushed the FIB to its limits, achieving a maximum count of 4,194,304 entries for IPv4 routes and 2,097,152 entries for IPv6 routes. To confirm the proper installation of all routes in the FIB, bidirectional traffic was generated from the BGP-advertised routes, reaching 99 % of the port's capacity, equivalent to 792 Gbps. The FIB showcased its robust capacity to handle this maximum load seamlessly without incurring any packet loss while keeping the memory utilization at a moderate 40 %.

We also intended to test the RIB's potential capacity. However, since the device doesn't have a limit on the number of entries in the RIB, there is no hard limit on the RIB scale. The EANTC team decided to advertise approximately 20 times more routes than the FIB capacity for IPv4 and IPv6, respectively. We verified that these routes were successfully installed in the RIB, as shown in Figure 25.

```
~X4-0]display BGP routing-table statistics
Total Number of Routes: 80000000
~X4-0]display BGP IPV6 routing-table statistics
Total Number of Routes: 40000000
~X4-0]
```

*Figure 25*

**☑ The 19.2T line card showed a FIB capacity of 4.2 Million IPv4 routes and 2.1 Million IPv6 routes (simultaneously) and confirmed RIB capacity of up to 80 Million IPv4 and 40 Million IPv6 routes.**

## Key Configuration Verification & Second Authentication

The NetEngine 8000 implements various security features with the Huawei iMaster NCE controller. Huawei invited us to witness some of these security functions.

The first function we verified is called Security Situational Awareness. It refers to the proactive approach in cybersecurity where the focus is not just on preventing attacks but, more importantly, on early identification and detection of potential threats. It involves being constantly aware of the security environment and potential vulnerabilities, enabling a quicker and more effective response to attacks.

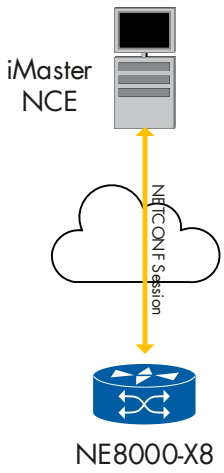


*Figure 26*

This testing process involved the Huawei iMaster NCE controller establishing a connection with the NE8000-X8 router. The NCE identified a range of key configurations as potential risk items, which it evaluated on the routers. These include settings such as *Telnet* activation and the choice of authentication algorithms.

The interface displayed each item's name (Configuration name), the type of feature (e.g., MGMT, Channel MGMT, or IP Routing MGMT), and any insecure configurations, specifying the nature of the risk. A column for the baseline value suggests improvements from the NCE and another for operational suggestions related to the item.

During the test, we reviewed the Baseline Suggestion for the IP-blocking feature in SSH, which was marked as risky. The reason for insecurity was that this feature was turned off on the node (Figure 27).

After enabling it on the NetEngine 8000 (manually), we accessed the NCE controller GUI and clicked 'check now' to update the NE8000's status and features. Subsequently, the feature's status transitioned from risky to normal following the configuration adjustment on the NE8000 (Figure 28).

Next, we confirmed that the NE8000-X4 supports an important feature: second-authentication. Huawei has identified specific commands as highly risky due to their potential to disrupt ongoing services and lead to failures in the storage system.

We tested this by attempting to execute high-risk commands, such as 'fast reboot' or 'undo mpls.' In these instances, a warning was triggered, prompting us to re-enter the login password of the current user as a form of secondary authentication. Note that "secondary" does not mean "multi-factor" in this context. The secondary authentication is merely a protection against attackers taking over an existing admin session, for example, by physically accessing an unlocked workstation or hacking into an active administration session.

**☑ The NetEngine 8000, in conjunction with iMaster NCE, demonstrated risk assessment of specific administrative actions, enabling password re-verification for high-risk actions.**

| Check Item Name | Check Result | Feature Name | Feature Type | Insecure Conf... | Check Baseline V... | Risk Level | Check Item Type | Last Checked |
|---|---|---|---|---|---|---|---|---|
| IP block feature for SSH server | Risky | SSH | Channel Management | DISABLE | ENABLE | High | Insecure Configuration | 2024-01-05 16:30:40 |

*Figure 27*

| Check Item Name | Check Result | Feature Name | Feature Type | Insecure Conf... | Check Baseline V... | Risk Level | Check Item Type | Last Checked |
|---|---|---|---|---|---|---|---|---|
| IP block feature for SSH server | Normal | SSH | Channel Management | -- | ENABLE | High | Insecure Configuration | 2024-01-18 18:10:32 |

*Figure 28*

## Intrusion Detection and Security Orchestration

Routers' management interfaces are traditionally protected from attacks by limiting IP-level access. In the optimal case, the management interface of a router cannot be attacked because hackers won't have network-level access. However, that's a somewhat idealistic assumption – in many cases, hackers will get access from within an organization, even from a network domain deemed secure and protected, by successfully attacking a weak device in such domain – for example, the laptop of an administrator or a forgotten old Windows server. Once inside the secure domain, the router must protect the management interface as a last resort against attacks, even if external firewalls have failed.

Huawei demonstrated several security mechanisms implemented on the NetEngine 8000 to protect the system independent of external firewalls.

### Multi-target Brute Force Attack & Abnormal Network Login

One way to break into a router is to hack its administrative password. Traditionally, administrators use weak passwords to simplify their lives, and advanced methods such as multi-factor authentication are either not supported by the router vendor or deemed too uncomfortable or may even create an operational risk (if a router needs to be recovered after Internet access failure, which might make multi-factor authentication unavailable).

Since the selection of passwords is difficult to influence, the primary hacking method must be prevented in another way. Huawei NetEngine 8000 implements a brute-force attack prevention. Attackers are not allowed to try out an unlimited amount of passwords quickly.

To test brute-force attack prevention, we utilized the NCE's feature of whitelisting IP addresses for Admin access. Any login attempts to Huawei Network Elements (NEs) from non-whitelisted IP addresses were logged in the NCE as potential attack attempts. During our test, we made six unsuccessful login attempts with incorrect passwords to two devices and succeeded on the seventh attempt on one of the routers only.

By doing this, we will generate two alarms for each node. However, the options for managing the user who successfully logged in will only be displayed for the device where the login was successful.

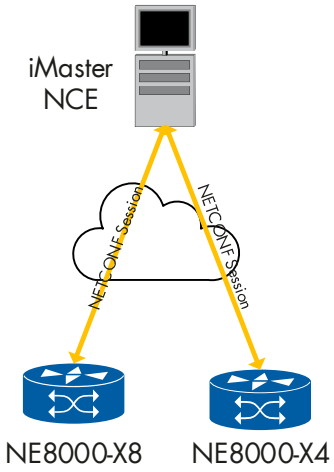


*Figure 29*

Upon detecting this activity, the NCE presented several response options:

**Whitelist the Attacker's IP:** This option involves adding the IP address of the attacker to the whitelist, which is generally not advisable in a real-world scenario but can be used for testing purposes.

**Orchestrator Handling:** This is a more sophisticated response where the NCE employs a specific playbook tailored to the detected threat. This playbook may include checks such as verifying if the user behavior is typical, examining if the account has AAA or SNMP user privileges, and other security checks (Figure 30).

Following this analysis, the NCE may block the user and require a password reset directly within the NCE system, enhancing security measures against such brute-force attacks.

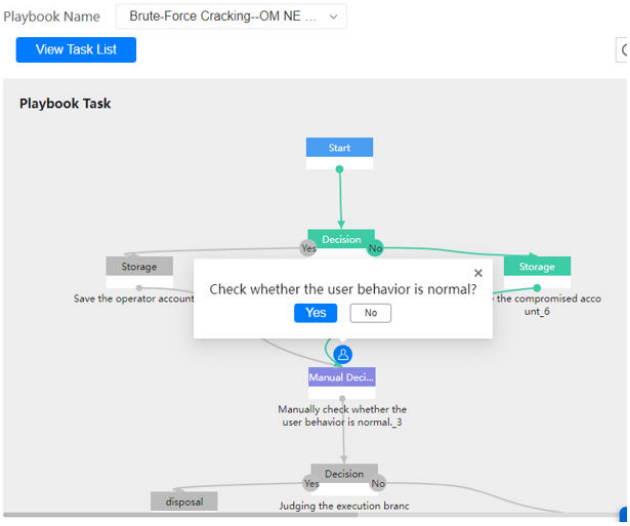


*Figure 30*

Next, we verified that the NCE can recognize abnormal login behavior, such as detecting an attacker using an unusual IP for device login. Usually, it requires 30 days to establish a user's typical behavior in iMaster NCE, including login times and IP addresses, but we modified this duration to only 24 hours.

On the first day, we logged into the router using a user account five times. This was part of creating a behavioral profile for this user. According to the system's

protocol, the user's IP address gets logged after more than five login attempts, and the login time is recorded after more than ten attempts.

Then, the next day, we attempted to log in to the router using a different IP address. As a result of this deviation from the established user behavior, the NCE issued an informational warning, demonstrating its ability to detect and alert unusual login activities (Figure 31).

**☑ The NetEngine 8000, in conjunction with iMaster NCE, protects the management interface against attackers by whitelists and login behavioral analysis.**

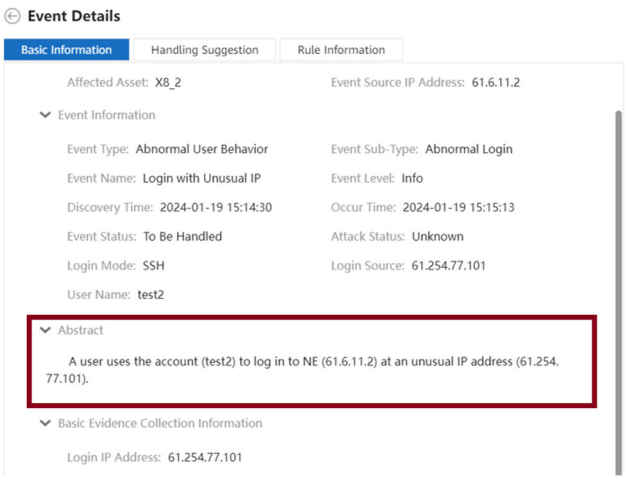


*Figure 31*

## BGP Graceful Degradation

BGP degradation protection is designed to prevent the router's performance from deteriorating due to excessive BGP route processing. When a router handles many BGP routes or frequent changes in these routes, it can strain the router's resources, leading to degradation in performance. This protection involves memory utilization monitoring, route limiting, selective route acceptance, and other measures.

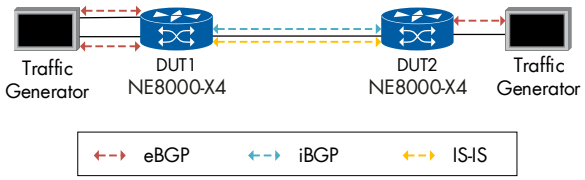


*Figure 32*

In this test, we prepared the setup between the two nodes with IS-IS neighborship and BGP sessions. We then set up External BGP, each with a Spirent port, and started advertising BGP routes to achieve around 60% memory usage. On the first device under test (DUT1), the routes were a mix of Public BGP (20 IPv4 and IPv6 peers; each peer advertised 1M routes) and the VPN BGP peers (178 IPv4 and IPv6 peers each advertising 3000 routes), and these VPN peer routes were to stay on DUT1 and were not to be advertised to DUT2.

The DUT2 had two peers with the tester 1 EBGPv6 and 1 EBGPv4; each one advertised 1M routes were advertised from DUT2 to DUT1.

We utilized a Huawei script to measure the route learning rate, recording data every 100 ms and measuring 292,816 routes/second. The total number of routes learned was 22,534,001 on DUT1, while memory utilization reached 59 %.

Later, Huawei implemented a device memory usage threshold to discard new routes when memory utilization hit 85 % automatically. This threshold command restricts the rate of route learning once the device memory reaches 60 % and proceeds to reject any additional routes upon reaching 85 % memory load.

Before we started the overloading, we generated IMIX traffic between the test tool, DUT1, and DUT2, respectively, using all 22 Million public routes utilizing 80% of the line rate of 2x100GbE ports. The background traffic reached 160.18 Gbit/s.

We started the overloading by establishing the overloading peers from the tester to the VPN instance on DUT1, which should advertise 70 M 50:50 v4/v6. We observed that the memory usage increased, and the new limited rate was 97,477 routes/s until the MPU memory usage reached 85%; then, the device stopped learning any new routes. We noticed no packet loss during the entire test, and the latency measurement for the background traffic didn't show any distinct differences.

Finally, we experimented with modifying the AS path on DUT2 and implementing a manual policy on DUT1 to deny peers that advertised overload routes. This was a demonstration of manually protecting the device's memory from overload. Following these adjustments, the number of routes in the VPN instance decreased, and the memory usage returned to around 60%.

## BGP Long-Term Memory Overload Control

We reused the setup in Figure 32 and replicated the overloading situation on DUT1 (NE8000-X4).

The test was designed to simulate overload for 40 minutes, a period configured to activate overload protection mechanisms. Typically, this protection would last 24 hours in real-world scenarios, providing a day-long buffer for maintenance teams to address issues once memory utilization reaches the critical 85% threshold.

Upon reaching this overload threshold, the router signaled an alarm. After the 40-minute alarm period, the overload escape timer expired, resetting BGP peers.

Notably, the router was programmed to selectively reset connections, preserving those peers pre-configured to be exempt from resets due to low memory. As a result, these exempted peers remained active and were never reset, demonstrating the router's handling of network resources under strain.

The automatic reset of the overload peers brought memory utilization below 85%, prompting some peers to re-establish until manual intervention was applied. We then manually halted some BGP overload peers from the Spirent, waiting for memory stabilization to assess further action. Eventually, the memory stabilized at 76%, with all BGP sessions on DUT1 active and stable.

**☑ The NetEngine 8000 software protects against BGP processing overload by monitoring memory usage, stopping BGP route learning in advance, and resetting peers when needed.**

## Security Functionality Tests

In our series of tests, we implemented a Committed Access Rate (CAR) to explore a vital use case centered around network security and efficiency. The primary scenario under investigation is one where the router experiences a surge in packet traffic. This situation often requires routing a significant number of packets to the CPU for processing, a process that, while routine, can make the router susceptible to various attacks.

In all the subsequent tests, we focused on the 19.2T LPU of the NE8000 x4. We utilized both ISIS and BGP protocols, covering IPv4 and IPv6. The tests were conducted under significant load conditions, using 80 % of the 400GbE port's capacity, and included a substantial amount of background traffic, amounting to 20,000 flows.

### Invalid Packet Filtering

We focused on the NE8000 X4's ability to filter invalid packets and withstand common cyber threats. Note that this family of tests was carried out on a separate test bed provided by Huawei with an NE8000-X4. Two ports of the NE8000-X4 were connected to a traffic generator that created network and transport layer (L3 and L4) attacks.

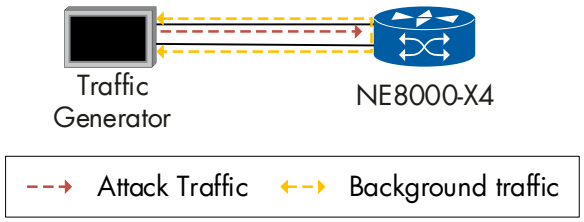


*Figure 33*

We started by sending IPv4 and IPv6 background traffic at 80% of the port speed. This step was important for understanding how the router handles high-volume traffic during an attack. In our scenario, we could send up to 400 Gbit/s traffic bidirectionally since the setup was limited to two ports. Attack traffic had to be sent from one port only. In a real network scenario, this would represent an attack coming from a single port only, which may or may not be representative. Transport layer attacks that are difficult to counter are called "distributed denial of service" (DDoS) attacks, which indicate that they typically come from multiple ports. In our tests, we were invited to validate only the pure functionality of the protection mechanisms.

The Huawei team enabled the software features to counter well-known port attacks, smurf attacks, zero-load packet attacks, and TCP SYN floods.

```
abnormal-packet-defend enable
udp-packet-defend enable
tcpsyn-flood enable
fragment-flood enable
ipv6-abnormal-packet-defend enable
ipv6-udp-packet-defend enable
ipv6-tcpsyn-flood enable
```

*Figure 34*

Next, we started sending simulated attack traffic, including TCP SYN Floods for IPv4 and IPv6, smurf attacks, malicious traffic sent to well-known ports (BGP, FTP, TELNET..), and zero-data packets. The NE8000

| Slot/Intf | Attack-Type | Total-Packets | Passed-Packets | Dropped-Packets |
|---|---|---|---|---|
| 4 | Tcpip-defend | 301469832 | 160389 | 301309443 |
| | Abnormal-packet | 104338629 | 0 | 104338629 |
| | Udp-packet | 0 | 0 | 0 |
| | Tcpsyn-packet | 197131203 | 160389 | 196970814 |
| | Fragment-packet | 0 | 0 | 0 |

*Figure 35*

| Index | CarID | Packet-Info | Passed Packets | Dropped Packets |
|---|---|---|---|---|
| 401 | 124 | ICMP-BROADCAST-ADDRESS-ECHO | 0 | 58126109 |

*Figure 36*

command-line interface allowed us to observe the router's real-time response.

Figure 35 & Figure 36 showed that the router could defend against all attack types chosen in this test case. The abnormal packet attacks included zero-data packets on IPv4 and IPv6 flows, attacks on well-known router ports, smurf attacks, and ICMP-Broadcast-address-echo packets.

We confirmed this protection worked with minimal impact on CPU (max 9 %) and memory usage (max 25 %) without dropping any 0.32 Terabit/s bidirectional background traffic.

### Port Level Isolation

The test evaluated a router's Port-level CAR (Committed Access Rate) feature, explicitly targeting the port level within the Line Processing Unit (LPU).

We configured dual stack ISIS and BGP along with 20K flows split evenly between IPv4 and IPv6. This setup aimed to create a realistic and robust testing environment.

The primary objective was to verify the Port-level CAR, essentially assessing the port-level throughput control in the LPU. We generated 320 Gbit/s of background traffic on 400GbE ports, targeting BGP and ISIS protocols with attack traffic for IPv4 and IPv6 at 10 Gbps per attack stream.

The background traffic flowed smoothly for 5 minutes during the test without packet loss, and latency and jitter values were monitored. When attack traffic was introduced, it specifically targeted one of the ports (port 13) on the chip, and we observed that while this port failed to establish any peerings, the other port (port 8) on the same chip remained unaffected and successfully established peerings. After halting the attack traffic, both protocols could establish peering again without issues.

Figure 37 shows that the background traffic experienced no packet loss throughout the test while generating the attack traffic simultaneously.

This demonstrated the router's resilience, seamlessly maintaining protocol operations on one port while another was under direct attack.

| Stream Block | Tx Count (Frames) | Rx Count (Frames) |
|---|---|---|
| ISIS_port | 629,336,410 | 0 |
| BGP_neighbor | 629,336,535 | 0 |
| isis_port_ipv6 | 629,336,472 | 0 |
| bgp_neighbor_ipv6 | 629,336,595 | 0 |
| bg-ipv4_1 | 19,312,602,757 | 19,312,602,757 |
| bg-ipv6_1 | 19,312,602,756 | 19,312,602,756 |
| bg-ipv4_2 | 19,251,020,343 | 19,251,020,343 |
| bg-ipv6_2 | 19,251,020,342 | 19,251,020,342 |

*Figure 37*

### Neighbor-level Isolation

Micro-isolation protection is a network security feature designed to safeguard the establishment of network sessions, particularly for protocols like ISIS, BGP, LDP, and OSPF as shown in Figure 38.

The primary function of micro-isolation protection is to prevent the monopolization of interface bandwidth resources by an excessive flow of protocol-specific packets.

```
[~X4-0]micro-isolation protocol-car ?
  bgp             Border Gateway Protocol(BGP)
  igmp            Specify Internet Group Management Protocol
  isis            Intermediate System to Intermediate System (ISIS)
  ldp             Specify Label Distribution Protocol (LDP) configuration
                  information
  ldp-tcp         LDP TCP Packet type
  ldp-udp-local   LDP UDP_local Packet type
  mld             Specify Multicast Listener Discovery
  ospf            Open Shortest Path First (OSPF)
  ospfv3          OSPF version 3 for IPv6
  pcep-ipv4       PCEP IPv4
  pcep-ipv6       PCEP IPv6
  rsvp-te         Specify RSVP-TE configuration information
```

*Figure 38*

We enabled this feature during our test and generated traffic attacks using ISIS and BGP ports while generating background traffic. Upon initiating attack traffic for BGP and ISIS (both v4 and v6), the statistics for dropped and passed packets started increasing, confirming the effectiveness of micro-isolation CAR.

We successfully established BGP and ISIS neighbors from the tester amidst an ongoing attack and background traffic. This demonstrated the router's ability to identify and allow legitimate traffic while under attack, thereby maintaining stable peering connections.

The effectiveness of this feature was further validated when we disabled it and observed that the router, facing an attack and without the protective feature active, could not establish new sessions. This indicates that the router, in the absence of the security feature, rejects all connections during an attack.

### Session-level Isolation

Two BGP neighbors between NE8000-X4 and the Tester were successfully established. For precise session identification and control, we used key characteristics (protocol, source, and destination IP addresses; source and destination ports for BGP; protocol, source, and destination MAC addresses; ISIS interface, and the index for ISIS) to create distinct BGP and ISIS sessions. One session from each protocol pair was targeted for attack, using matching key characteristics.

Continuous attack traffic was generated, targeting the first BGP session using its specific key characteristics. This led to the targeted BGP session going down and subsequently re-establishing. After three minutes, we checked the state of all BGP neighbors, confirming the attack's impact on the sessions. The new BGP session exhibited new CAR statistics with only Green packets, indicating normal traffic.

In contrast, the attacked ISIS session went down and could not be reestablished due to the identical keys.

### Board-level Protocol Isolation

We initiated the attack traffic at 10 Gbps for both BGPv4 and BGPv6. We monitored the CLI and noted increased dropped and passed packets in the BOARD CAR statistics for BGPv6 and the new BGP instance (NEW_BGP).

Next, we successfully established 50 ISIS peers from the tester and one BGP peer each for IPv4 and IPv6.

The neighbor CAR effectively protected the BGP sessions against the BGP attack. Thanks to protocol-level isolation on the same board, the ISIS protocol was shielded from the BGP attack from the same port. For both BGPv4 and BGPv6, we observed an increase in both dropped and passed packets. However, there were zero dropped packets for ISIS, with an increase only in passed packets.

**☑ The NetEngine 8000 implements versatile security functions for the major control plane protocols BGP, ISIS, OSPF, and LDP. Packet content validation, port-, neighbor-, session- and board-level protocol isolation can be combined to stop advanced attacks on the routing control plane.**

## Conclusion

During an extensive range of tests of the Huawei NetEngine 8000 at Huawei labs in Beijing in January 2024, the EANTC team successfully confirmed all of Huawei's claims of line card performance, service scale, energy efficiency, manageability, and router security. Most of the test cases focused on the new 19.2T line card, where we confirmed its impressive throughput performance, SRv6, EVPN service support, and low energy consumption.

Router security and service continuity are two additional major test areas that Huawei invited us to validate on the NetEngine 8000 with the iMaster NCE controller. The initial functional observations are promising, confirming that Huawei aims to help network operators protect the router management ports without relying on external firewalls only.

In summary, EANTC was impressed with the range of test areas and the efforts of Huawei's product management teams to improve all aspects of the NetEngine 8000 X series core and aggregation routers.

### References

[1] EANTC multi-vendor MPLS SDN Interoperability Test 2023 (https://eantc.de/events/mpls-sdn-interop-2023/)

[2] EANTC test of Huawei Intelligent IP & Cloud Network Solution, 2020/2021 (https://eantc.de/test-reports/huawei-intelligent-ip-cloud-network-solution/

[3] EANTC test of Huawei 4T LPU, 2018 (https://eantc.de/wp-content/uploads/2023/11/EANTC-Huawei-LPU4T-line-card-Marketing-Report-final.pdf)

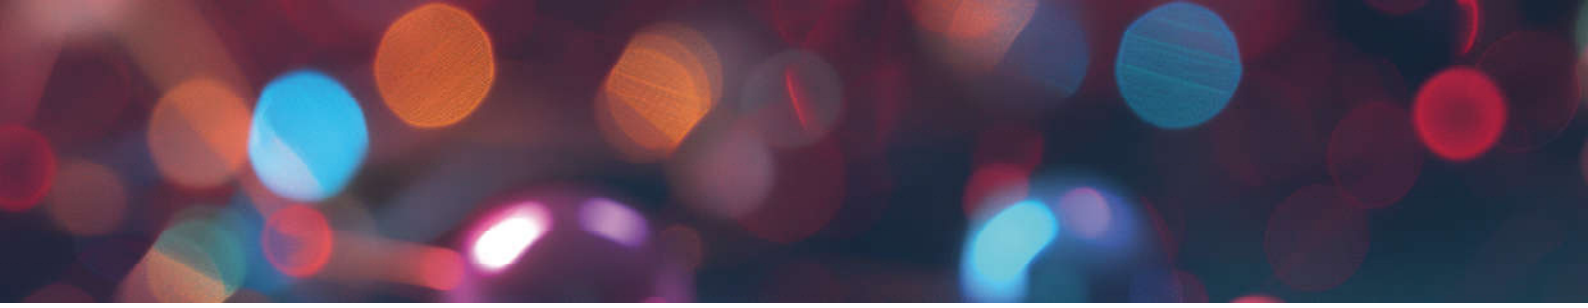



EANTC AG
Salzufer 14, 10587 Berlin, Germany
info@eantc.de, https://www.eantc.de/
[v1.1 20240221]