EANTC Independent Test Report

Juniper PTX Series Routers Establishes Benchmarks for Scalable Core and Peering Deployments

November 2021

EANTC

Juniper Networks

Introduction

As 400G adoption increases across cloud and service provider networks, the increased number of connections has put immense pressure on Core and Peering routers to deliver high performance, scalability, and reliability to the network. Today's market needs the flexibility and capacity of IP filters, routing scalability, and flow sampling to meet the challenging combination of requirements for core and peering uses cases, and ensure top performance, avoid bottlenecks, and guarantee service readiness on critical nodes.

To determine what is best for their networks, service providers have been carrying out extensive benchmark tests to measure and ensure high performance, scalability, and reliability. EANTC has been commissioned by Juniper Networks to independently benchmark several of their 400G-capable core and peering routers, with tests focused on verifying five major evaluation criteria to meet top requirements: Performance, Filter Scale, Forwarding Information Base (FIB) Scale, Sampling, and Media Access Control Security (MACsec).

- Performance: One of the most significant evaluation areas based on core and peering use case needs, this criterion was continuously monitored under all high throughput feature load test cases.
- Filter Scale: Peering routers are at the front line of the denial-of-service attack protection. Operators today rely on the flexibility and capacity of IP filters to mitigate the impact of denial-of-service attacks.
- FIB Scale: Routing scalability is important specifically for complex internal interconnections or external peering scenarios. As internet route tables continue to grow in a non-linear fashion, especially on the IPv6 side, service providers looking to establish multiple BGP peers or create different FIB views for their customers require a high level of FIB scale and flexibility.
- Sampling: Flow sampling has always helped network operators to gain visibility into the types of traffic transiting through their core or peering nodes. Today's challenge is to sustain flow sampling for many connections at a high rate, as traffic sampling features are required to integrates with operator's reporting systems, log servers, and security solutions.
- MACsec: As service providers and corporations consider MACsec to secure frame forwarding already at the Ethernet layer, the MACsec capability was tested and routers forwarding performance is benchmarked with traffic.

Test Highlights

- → IPv4 and IPv6 line-rate throughput of 9.6 Tbit/s with IMIX profile on 24x 400GbE ports of PTX10001-36MR
- → IPv4 and IPv6 line-rate throughput of 14.4 Tbit/s with IMIX profile on 36x 400GbE ports of PTX10008-LC1201
- → Filter match scaling up to 200,000 consecutive IPv4 prefixes or 80,000 random IPv4 prefixes with 60,000 port numbers
- → Filter match scaling up to 60,000 consecutive IPv6 prefixes or 48,000 random IPv6 prefixes with 60,000 port numbers
- → IPv4 Forward information base (FIB) scaling with 3.32 million Internet-type prefixes on four virtual routing and forwarding (VRF) instances or 4.00 million consecutive IPv4 prefixes
- → IPv6 FIB scaling with 420,000 Internet prefixes on four VRFs, or (separately) two million consecutive IPv6 prefixes with a prefix length of 64, or 900,000 consecutive IPv6 prefixes with a prefix length of 126
- → MACsec encryption with AES-XPN-128 or AES-XPN-256 ciphers across four ports with IMIX profile with 100% throughput
- → Line-rate active flow monitoring (IPFIX) on eight 400GbE interfaces with a sampling rate of 1:3300 packets/second on PTX10008-LC1201; line-rate IPFIX on seven 400GbE interfaces with a sampling ratio of 1:2850 packets/s of PTX10001-36MR

Hardware and Software

Router Type	Number of 400GbE Ports used in Test	RAM Capacity	Software Version
PTX10001-36MR	24	RE—DRAM 34990 MB (63217 MB installed)	Junos Evolved: 21.2R1
PTX10008-LC1201	36	RE—DRAM 58397 MB (63813 MB installed)	Junos Evolved: 21.2R1

Table 1: Hardware and Software Details of the PTX10001-36MR and PTX10008-LC1201

Test Scenario and Overview

Considering the use case requirements, Juniper configured five test bed topologies for the benchmark testing on two of their 400G products: the PTX10001-36MR and the PTX10008-LC1201. The PTX10001-36MR is a compact, 1RU fixed configuration router with a capability of 24 400GbE ports, and the PTX10008-LC1201 is an 8-slot, modular router with each line card capable of 36 400GbE ports.

In the test bed, Juniper used QSFP56-DD 400GE DR4 optics on the 400GbE ports of the PTX routers with a wavelength of 1311nm. Also, they configured five appliances of the Spirent N11U Chassis with eight 400G ports per appliance with version 5.19.0846. Optics QSFP-DD was used on Spirent 400GbE ports to connect to the 400GbE ports of the PTX routers via single-mode fiber optic cables.

For the RFC2544 test, Juniper connected 24 400G links directly from the PTX10001-36MR to the Spirent TestCenter as shown in Figure 1, and on the PTX10008 Juniper connected 36 400GbE ports on one line card of the PTX10008-LC1201 directly to the Spirent test center as shown in Figure 2.

For functional tests including filtering capabilities, FIB scalability, and flow sampling, Juniper built a test bed with up to eight 400GbE links of PTX10001-36MR connected directly to the Spirent TestCenter and up to 16 400GbE links of the PTX10008-LC1201 connected directly to the Spirent test center.

For the MACsec test, Juniper connected PTX10001-36MR and PTX10008-LC1201 directly via four 400G links and four 400G link connections with the Spirent test center from each PTX router as shown in Figure 3.

Tests were conducted remotely at Juniper Network's Sunnyvale headquarters in July 2021 due to the Covid-19 pandemic. EANTC supervised all test preparations, executions, and documentation efforts live and accessed test tools and routers configurations.

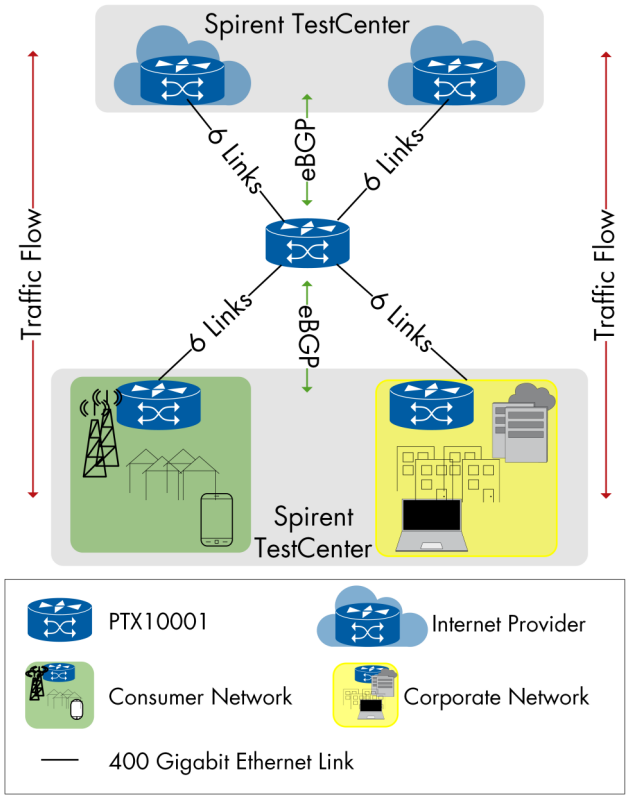

Figure 1: PTX10001-36MR RFC2544 Test Bed

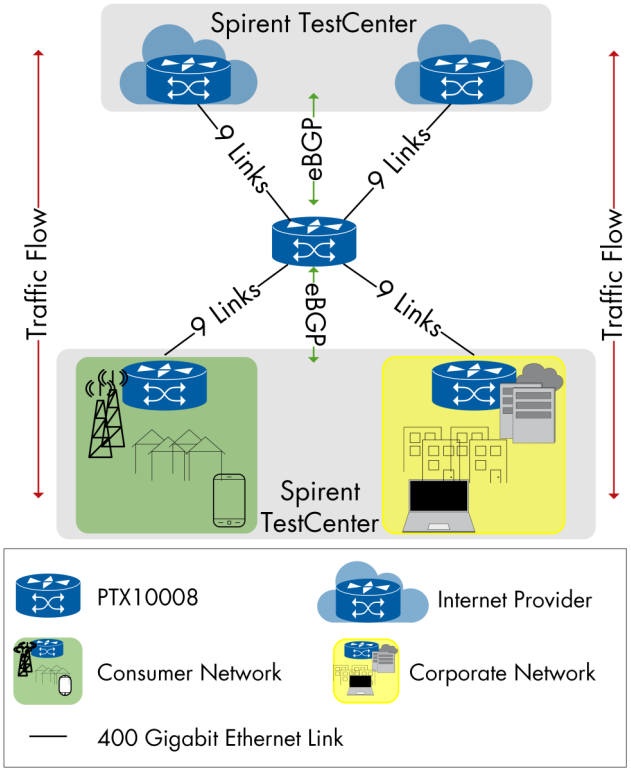

Figure 2: PTX10008-LC1201 RFC2544 Test Bed

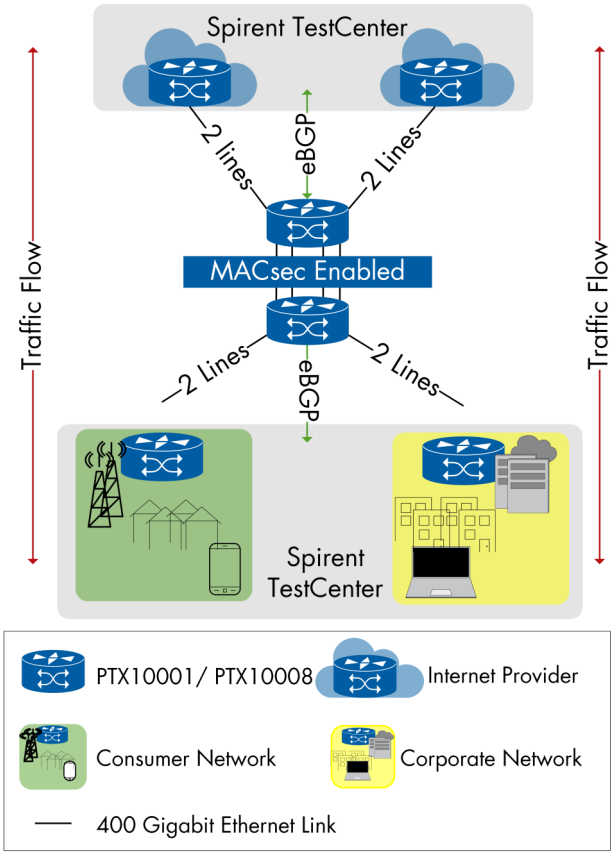

Figure 3: PTX10001-36MR and PTX10008-LC1201 MACsec Test Bed

Test Results

RFC2544 Performance Benchmarking

EANTC verified the throughput performance for IPv4 and IPv6 using RFC2544 with an IMIX profile in a complete mesh topology. During the test, Juniper connected the 400GbE PTX family router ports directly to the 400GbE ports of the Spirent TestCenter, as shown in Figures 1 and 2. For the test, Juniper configured the EBGP peers on the Spirent TestCenter ports with an equal number of routes configured across all the EBGP peers and advertised the routes in the PTX routers. Different test parameters were used on the PTX family routers, as shown in Table 3 and 4.

EANTC simulated 99.99% full mesh IPv4 and IPv6 traffic separately in two iterations for the advertised routes. In the test, we chose a 99.99% load per port i.e. 399.96Gbp/s as the remaining 40Mbp/s of bandwidth per port can be used by the control packets. We verified the throughput performance of IMIX profile for 120 seconds with three trails and for continuous 15 minutes. The used IMIX profile is shown in Table 2. Traffic sent from one port on the router to all other ports in a mesh topology using the EBGP routes advertised to the router. Additionally, EANTC measured the CPU usage during the routes learning, and it was 99.9%, while 7% was the average CPU usage during the whole traffic run. The routers passed the throughput tests, achieving 99.99% of the line rate with minimum latency values and no packet loss, as shown in Table 5 and Table 6.

Frame Size (Bytes)	Weight IPv4 IMIX	Weight IPv6 IMIX
64	3	-
78	-	3
100	26	26
373	6	6
570	5	5
1300	6	6
1518	16	16
9000	1	1

Table 2: IMIX Profiles

IP Type	Number of 400GbE Ports used in Test	Port Speed	Number of EBGP Peers	Routes per EBGP Peer	Total Number of Advertised Routes
IPv4	24	400 Gbit/s	24	166,666	4 Million
IPv6	24	400 Gbit/s	24	83,333	2 Million

Table 3: PTX10001-36MR Throughput Test Parameters

IP Type	Number of 400GbE Ports used in Test	Port Speed	Number of EBGP Peers	Routes per EBGP Peer	Total Number of Advertised Routes
IPv4	36	400 Gbit/s	36	111,111	4 Million
IPv6	36	400 Gbit/s	36	55,555	2 Million

Table 4: PTX10008-LC1201 Throughput Test Parameters

IP Type	Packet Size	Throughput per 400GbE Port, bi-directional	Packet Loss Percentage	Packet Forwarding Performance per Port	Latency (µs) min/avg/max
IPv4	IMIX (Three 120 seconds trails)	399.96 Gbp/s	0.00%	61 Mpps	4.13/7.47/11.11
IPv4	IMIX (Running for 15 minutes)	399.96 Gbp/s	0.00%	61 Mpps	4.11/7.47/11.03
IPv6	IMIX (Three 120 seconds trails)	399.96 Gbp/s	0.00%	61 Mpps	4.22/7.72/11.19
IPv6	IMIX (Running for 15 minutes)	399.96 Gbp/s	0.00%	61 Mpps	4.19/7.72/11.54

Table 5: PTX10001-36MR IMIX Profile Throughput Results

IP Type	Packet Size	Throughput per 400GbE Port, bi-directional	Packet Loss Percentage	Packet Forwarding Performance per Port	Latency (µs) min/avg/max
IPv4	IMIX (Three 120 seconds trails)	399.96 Gbp/s	0.00%	61 Mpps	4.28/6.64/8.82
IPv4	IMIX (Running for 15 minutes)	399.96 Gbp/s	0.00%	61 Mpps	4.25/6.64/8.96
IPv6	IMIX (Three 120 seconds trails)	399.96 Gbp/s	0.00%	61 Mpps	4.29/6.79/9.3
IPv6	IMIX (Running for 15 minutes)	399.96 Gbp/s	0.00%	61 Mpps	4.25/6.79/9.43

Table 6: PTX10008-LC1201 IMIX Profile Throughput Results

Filter Scale Benchmarking

In this test, Juniper demonstrated the capacity and flexibility of the PTX10001-36MR and PTX10008-LC1201 filters. Juniper configured the IPv4 and IPv6 filters on PTX routers in two profiles. Each profile has a set of match criteria parameters and actions for the matched packets. Juniper stated that PTX10001-36MR and PTX10008-LC1201 supports 42 match parameters and ten actions per matched packet, which enables service providers with a wide range of parameters to match in a packet. They also mentioned no limit in configuring all the supported match criteria parameters and actions in any combinations. In the test, Juniper configured source IP address only, source and destination IP address, or source and destination IP address with port numbers as the match criteria parameters. EANTC validated the forwarding and counter statistics actions for the configured match parameters on PTX10001-36MR and PTX10008-LC1201.

Each filter profile of the PTX routers can accommodate 200,000 consecutive IPv4 prefixes, 60,000 consecutive IPv6 prefixes, 80,000 random IPv4 prefixes, or 48,000 random IPv6 prefixes. The PTX10001-36MR and PTX10008-LC1201 performs IP prefix pattern identification to optimize the resources by compressing the consecutive IP prefixes into an aggregate IP prefix, and unique resources will be allocated per random IP prefixes. Hence the maximum filters support for the consecutive IP prefixes was more than the random IP prefixes. Juniper PTX routers also support the matching of 64,000 unique port numbers in addition to the IP prefixes per profile. In the filter profiles, Juniper PTX routers support individual ports or port ranges with consecutive port numbers. PTX10001-36MR and PTX10008-LC1201 treat the port range as a single entry and perform the match for all the port numbers in the configured port range. Each combination of a source IP address, destination IP address, and port number will be applied as a single filter on the ingress/egress port of the router. That is, each Source IP address, draws a combination with all the destination IP addresses and port numbers configured, and all combinations are applied as filters per port. For the packet match within a single filter profile, the counters for the particular filter shall increase and provide the statistics of the number of packets matched per filter. Juniper configured the third profile with counter statistics and forward actions for the no-match packets of the first two configured profiles. The counter for the number of unmatched packets was incremented, and the packets were forwarded. Figure 4 and Figure 5 show the test bed topologies for PTX10008-LC1201 and PTX10001-36MR.

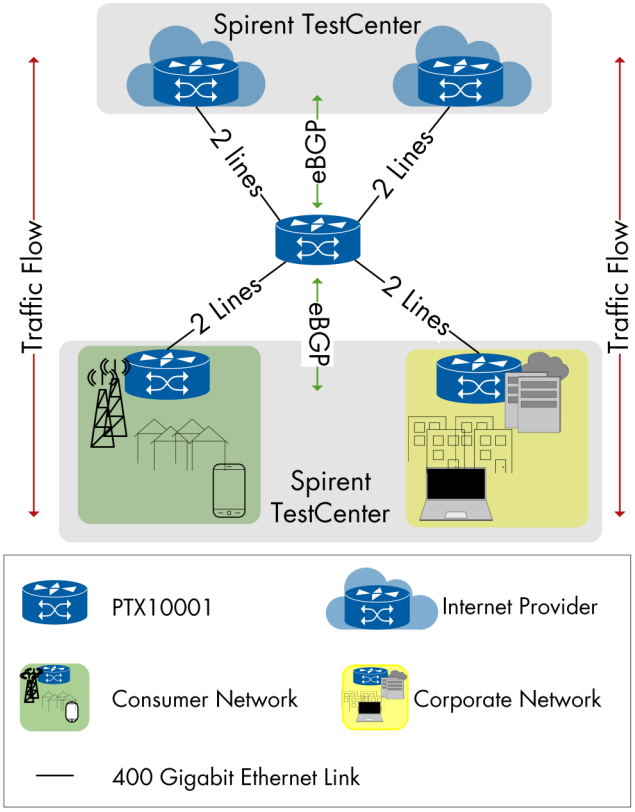

Figure 4: PTX10001-36MR Filter Scale Test Bed

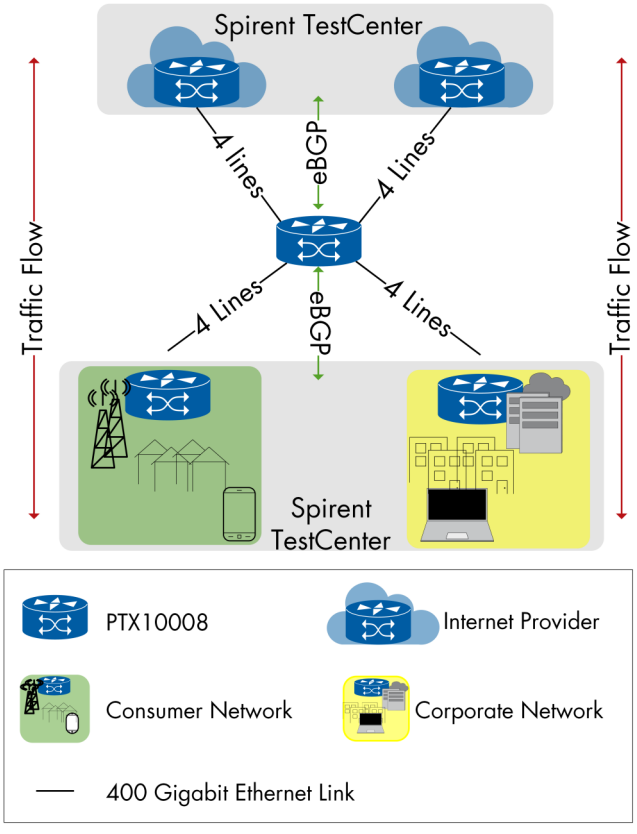

Figure 5: PTX10008-LC1201 Filter Scale Test Bed

During the test, Juniper configured two filter profiles. Each filter profile contained a set of source addresses, destination addresses, and port numbers. Also each profile had an equal number of IP addresses, an equal number of source IP addresses and destination IP addresses. They applied the filter profiles on the ingress ports of the PTX routers and used various IP addresses and port numbers, as shown in Table 7. For the test, Juniper connected and configured eight EBGP peers for the PTX10001-36MR and 16 EBGP peers for the PTX10008-LC1201 on the Spirent TestCenter ports in two different testbeds. The sixteen ports on the PTX10008-LC1201 are equally distributed across two line cards used in the test. The filter profiles were applied on four traffic ingress port of PTX10001-36MR and on eight traffic ingress ports of the first line card of PTX10008-LC1201.

EANTC simulated 320 Gbp/s uni-directional traffic using an IMIX profile as shown in Table 2 per Spirent port, presenting 80% of the line rate. EANTC chose 80% of the line rate per port for all the functional tests as it is close to the real-time usage per port.

EANTC performed the test for ten minutes and observed that the counters of the matched filter profile increased with zero packet loss. We also validated the third filter, which counted the unmatched traffic of the first two filters and the counter was incremented. For the filters with port numbers, simulated TCP traffic with a port number of extra 2 bytes for every packet size of the IMIX profile in Table 2.

IP Type	IP Prefix Type	No. of Filter Profiles Configured	No. of IP Prefixes Configured	No. of Source Addresses Configured	No. of Destination Addresses Configured	No. of Port Numbers Configured	No. of Filters Applied per Ingress Port*
IPv4	Consecutive	2	200,000	100,000	100,000	-	50,000 x 50,000 x 2
IPv4	Random	2	80,000	40,000	40,000	-	20,000 x 20,000 x 2
IPv4	Random	2	80,000	40,000	40,000	60,000	20,000 x 20,000 x 30,000 x 2
IPv6	Consecutive	2	60,000	30,000	30,000	-	15,000 x 15,000 x 2
IPv6	Random	2	48,000	24,000	24,000	-	12,000 x 12,000 x 2
IPv6	Random	2	48,000	24,000	24,000	60,000	12,000 x 12,000 x 30,000 x 2

Table 7: Filter Profiles of PTX10001-36MR and PTX10008-LC1201

* Source IP Addresses x Destination IP Addresses x Port Numbers x No. of Filter Profiles

FIB Scale Benchmarking

The FIB scale is very important to the service providers specially when it comes to peering area as the Routing Information Base (RIB) and FIB capacity can indicate how much routes can be received from the peers. Bigger RIB table will give the service provide the flexibility to receive multiple internet routing tables from different peers which will enhance the path for different destinations based on the created FIB table, which will be reflected at the end on the users experience.

Juniper configured EBGP peers on two out of eight Spirent TestCenter ports connected with four million consecutive IPv4 prefixes or two million IPv6 prefixes as shown in Figure 6. The BGP routes advertised to the PTX routers from the Spirent TestCenter EBGP peers. For the IPv4 consecutive prefixes, Juniper used different prefix length distribution as shown in Table 9, and for the IPv6 consecutive prefixes, they used /64 and /126 prefix length. 80% of the port traffic, which presents 320Gbp/s was simulated with an IMIX profile for ten minutes for each type of prefix.

For the Internet routes FIB scale validation, Juniper configured two EBGP peers per VRF on the eight Spirent TestCenter ports. In the test, they configured four VRFs with 830,000 IPv4 internet prefixes per VRF and 105,000 IPv6 internet prefixes per VRF.

The total internet routes advertised in the IPv4 test case was 3.32M prefixes and 420,000 prefixes in the IPv6 test case. 80% of the port traffic, which presents 320Gbp/s was simulated with an IMIX profile for ten minutes for each type of prefix. The FIB scale results are shown in Table 8.

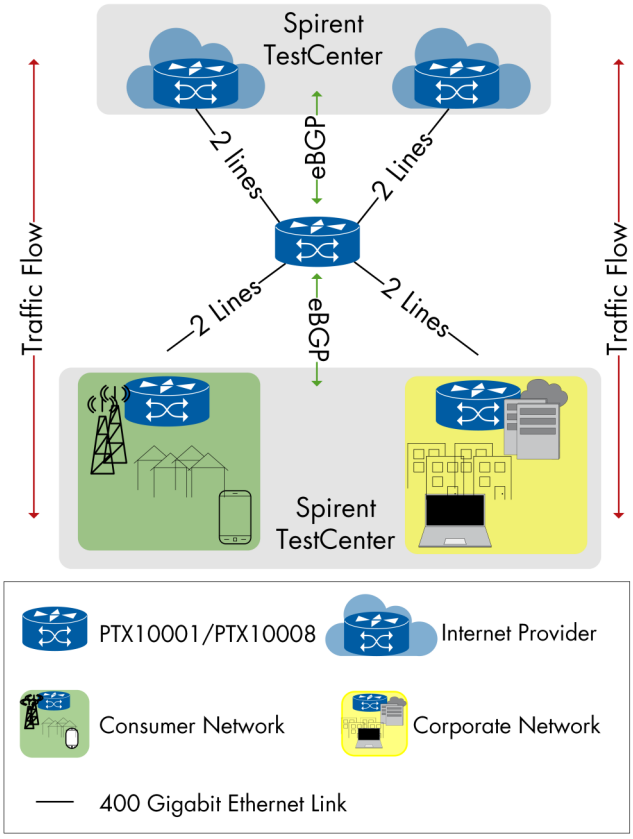

Figure 6: PTX10001-36MR and PTX10008-LC1201 FIB Scale Test Bed

EANTC tested the FIB scale with four million IPv4 consecutive prefixes, two million IPv6 consecutive prefixes with /64 prefix length, 900,000 IPv6 consecutive prefixes with /126 prefix length, 830,000 IPv4 internet prefixes and 105,000 IPv6 internet prefixes. EANTC witnessed the installation of all mentioned numbers of prefixes in FIB without impacting the PTX routers performance.

IP Type	IP Prefix Type	Packet Size	Maximum No. of installed Routes in FIB	Throughput per 400GbE Port, bi-directional	Packet Loss Percentage
IPv4	Consecutive	IMIX	4M	320 Gbp/s	0.00%
IPv4	Internet	IMIX	3.32M	320 Gbp/s	0.00%
IPv6	Consecutive IP prefix length /64	IMIX	2M	320 Gbp/s	0.00%
IPv6	Consecutive IP prefix length /126	IMIX	900,000	320 Gbp/s	0.00%
IPv6	Internet	IMIX	420,000	320 Gbp/s	0.00%

Table 8: FIB Scale Results of PTX10001-36MR and PTX10008-LC1201

Prefix Length	% Distribution	Actual Count
/20	2%	80000
/21	4%	160000
/22	14%	560000
/23	12%	480000
/24	68%	2720000

Table 9: IPv4 Prefix Length Distribution for Consecutive Prefixes

Sampling Benchmarking

Flow sampling is crucial for traffic composition analysis, validation of security policies, billing purpose and many more applications. Juniper flow (JFlow), industry wide known as IPFIX, addresses this by collecting and exporting the flow record attributes in a predefined template format to external collectors. A template defines a collection of fields, with corresponding descriptions of structure and semantics. Templates that are sent to the JFlow collector, contains structural information about the exported flow record fields and the system configuration parameters.

Juniper IPFIX implementation is distributed to line cards in a chassis based system. Export and sampling limits are specified per line card. In a fixed form factor device, sampling is implemented at the routing engine level. As per Juniper, they support both ingress and egress sampling of a physical interface or aggregated Ethernet interfaces, and allows user to customize their sampling based on various packet fields using filters, for instance user wants to sample packets only for a particular traffic pattern. During the test, EANTC validated the sampling on ingress physical interface with filter actions "Accept" and "Sample" where all the ingress packets are allowed and sampled as per the sampling ratio configured.

According to Juniper, PTX10008-LC1201 and PTX-10001-36MR are capable of sampling flows with an aggressive sampling ratio of 1:1 up to 1:7000000 based on various customer requirements. These systems can export up to 150k flows per second, where a flow record is created for every packet sampled and processed by the system. In most scenarios, the sampling ratio is less aggressive, like 1:10000 or higher.

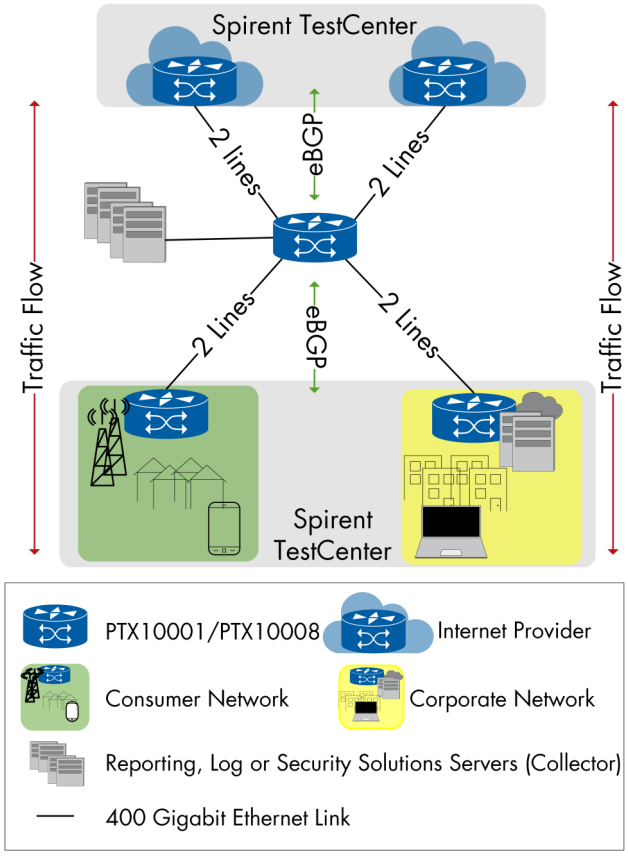

Figure 7: PTX10001-36MR and PTX10008-LC1201 Sampling Test Bed

DUT	Packets Sample Ratio	Setup Rate (Flows per Second)	Export Rate (Flows per Second)
PTX10008-LC1201	1:3300	149,126	148,282
PTX10001-36MR	1:2850	150,940	150,922

Table 10: Sampling Results of PTX10001-36MR and PTX10008LC1201

On the PTX10001-36MR, Juniper connected and configured seven EBGP peers and one JFlow collector on the traffic simulator ports, and for PTX10008-LC1201, Juniper connected and configured eight EBGP peers and one JFlow collector on the traffic simulator ports connected to the 400GbE ports of the PTX router, as shown in Figure 7, EANTC advertised two million routes to both PTX routers for the sampling test. During the test, the sampling was applied at the traffic ingress ports of the PTX routers.

In the router, EANTC monitoring flow with a sampling rate of 1:3300 packets for PTX10008-LC1201 and 1:2850 packets for PTX10001-36MR were configured. The test used the IMIX profile mentioned in Table 4 and simulated IPv4 traffic to the advertised routes for ten minutes. During the trials, EANTC simulated 99.99% of port traffic. We collected two samples of the flow information rate for every ten seconds from the routers during the traffic run. The flow information rate includes the flow packets, flows exported, and flow packets exported. The delta of the flow packets of two samples was divided by ten to calculate the setup rate, and the delta of flows exported was divided by ten to calculate the flow export rate.

Table 10 shows the sample rate and export of the flow for the PTX10008-LC1201 and PTX10001-36MR. Figure 7 shows the test bed topology for PTX10008-LC1201 and PTX10001-36MR.

MACsec Performance

MACsec uses MACsec Key Agreement (MKA), the standardized protocol to provide the required session keys and manages the required encryption keys on a MACsec connection between hosts like (PCs and servers) and switch. For the test, Juniper connected four 400GbE ports between PTX10001-36MR and PTX10008-LC1201, and four 400GbE ports between each PTX router and the Spirent TestCenter as shown in Figure 3.

MACsec was configured on the PTX10001-36MR and PTX10008-LC1201 connection ports. The encryption of traffic was initiated and terminated within the two routers. The MACsec packet header was of 24 bytes. During the test, EANTC verified the MACsec AES-XPN-128 and AES-XPN-256, using the full FIB scale i.e. four million synthetic IP prefixes with 99.99% port throughput. The test was performed with an IMIX profile shown in the Table 4 for ten minutes in two separate iterations. The throughput was achieved 99.99% without any packet loss, and latency values are as shown in the Table 11 and Table 12. The latency values were calculated when the PTX10001-36MR and PTX10008-LC1201 connected in the test bed, and traffic was passing through both routers during the test. The throughput was calculated with the interframe gap and the MACsec header.

Packet Size	Throughput per 400GbE Port, bi-directional	Packet Loss Percentage	Network Latency (µs) (min/avg/max)
IMIX	399.96 Gbp/s	0.00%	8.75/12.75/30.88

Table 11: MACsec Encryption 128 Results

Packet Size	Throughput per 400GbE Port, bi-directional	Packet Loss Percentage	Network Latency (µs) (min/avg/max)
IMIX	399.96 Gbp/s	0.00%	8.70/12.75/30.85

Table 12: MACsec Encryption 256 Results

Conclusion

EANTC verified the throughput test, filter scaling, FIB scaling, sampling and MACsec. All tests passed with the mentioned configurations in the report.

Juniper showcased significant performance and flexibility in configurations of the filters, sampling and MACsec. The stability and efficiency in PTX10001-36MR and PTX10008-LC1201 routers were remarkable. The maximum range of the filters supported per 400GbE port was impressive. The throughput with the maximum routes supported per IP type has demonstrated the efficient performance and capacity of PTX10001-36MR and PTX10008-LC1201. The flexibility in the sampling customization based on various packet fields using filters was notable, and the MACsec support with full throughput per 400GbE port was profound.

EANTC AG
Salzufer 14, 10587 Berlin, Germany
info@eantc.de, https://www.eantc.de/
[v1.0 20211109]