EANTC Independent Test Report

Juniper versus Cisco

Core and Peering Benchmarking

November 2021

Introduction

The telecom market is dynamic and challenging. It forced the big players to invest massively into their research and development and focus on equipping the market with the next generation of technology. The areas of improvement and investment could be achieving new scaling, adding new capability, or coming up with new features.

Juniper commissioned EANTC to independently validate the achieved capability and scalability of one of their core and peering PTX routers compared with Cisco's competitive product. In this test, EANTC compared the PTX10001-36MR fixed form factor device with a capability of 24 400 GbE ports with Cisco 8201 as it has the same number of rack-unit and 400GbE ports. Tests were conducted at Juniper's Sunnyvale headquarters remotely in August 2021.

EANTC and Juniper agreed to focus on the most relevant capabilities in the core and peering router deployments: throughput performance, filtering capabilities, routing scalability, and flow sampling.

Test Highlights

- → IPv4 and IPv6 line-rate throughput of 9.6 Tbit/s with IMIX profile on 24x 400 GbE ports of PTX10001-36MR and Cisco 8201
- → IPv4 Forward information base (FIB) scaling with 3.32M Internet-type prefixes on four virtual routing and forwarding (VRF) instances or 4M consecutive IPv4 prefixes on PTX10001-36MR and IPv4 Forward information base (FIB) scaling with 2.4M Internet-type prefixes on four virtual routing and forwarding (VRF) instances or 3.2M consecutive IPv4 prefixes on Cisco 8201
- → IPv6 FIB scaling with 420,000 Internet prefixes on four VRFs, or (separately) 2M consecutive IPv6 prefixes with a prefix length of 64, or 900,000 consecutive IPv6 prefixes with a prefix length of 126 on PTX10001-36MR and IPv6 FIB scaling with 30,848 Internet prefixes on four VRFs, or (separately) 1.3M consecutive IPv6 prefixes with a prefix length of 64, or 15,000 consecutive IPv6 prefixes with a prefix length of 126 on Cisco 8201
- → Filter match scaling up to 200,000 consecutive IPv4 prefixes or 80,000 random IPv4 prefixes with 60,000 port numbers on PTX10001-36MR and up to 2,480 consecutive IPv4 prefixes or 1,200 random IPv4 prefixes with 64,000 port numbers on Cisco 8201
- → Filter match scaling up to 200,000 consecutive IPv4 prefixes or 80,000 random IPv4 prefixes with 60,000 port numbers on PTX10001-36MR and up to 2,480 consecutive IPv4 prefixes or 1,200 random IPv4 prefixes with 64,000 port numbers on Cisco 8201
- → Line-rate active flow monitoring on seven 400 GbE interfaces with a sampling ratio of 1:2850 packets/s of PTX10001-36MR-36MR and Cisco 8201

Hardware and Software

Router Type	Number of 400GbE Ports used in Test	Software Version
PTX10001-36MR	24	Junos Evolved: 21.2R1
Cisco 8201	24	Cisco IOS XR Software, Version 7.3.1

Table 1: Hardware and Software Details of the PTX10001-36MR and Cisco 8201

Test Bed

EANTC configured four test bed topologies for the benchmark testing to cover performance and functional test cases on the PTX10001-36MR versus Cisco 8201. Juniper used Optics QSFP56-DD 400G DR4 on the 400 GbE ports of the PTX10001-36MR and QSFP-DD Optics on the 400 GbE ports of the Cisco 8201 router with a wavelength of 1311nm for all optics.

EANTC configured five appliances of the Spirent N11U chassis with eight 400G ports per appliance with version 5.19.0846. Optics QSFP-DD were used on Spirent 400 GbE ports to connect to the 400 GbE ports of the PTX10001-36MR and Cisco 8201 routes via single-mode fiber optic cables.

For the RFC2544 test, Juniper connected 24 400G links directly from the PTX10001-36MR and Cisco 8201 to the Spirent TestCenter separately, as shown in Figure 1. For functional tests like filtering capabilities, routing scalability, and flow sampling, Juniper built a test bed with up to eight 400 GbE links of PTX10001-36MR and Cisco 8201 connected directly to the Spirent TestCenter.

The test was conducted remotely due to the Covid-19 pandemic. EANTC supervised all test preparations, executions, and documentation efforts live and accessed test tools and nodes configurations.

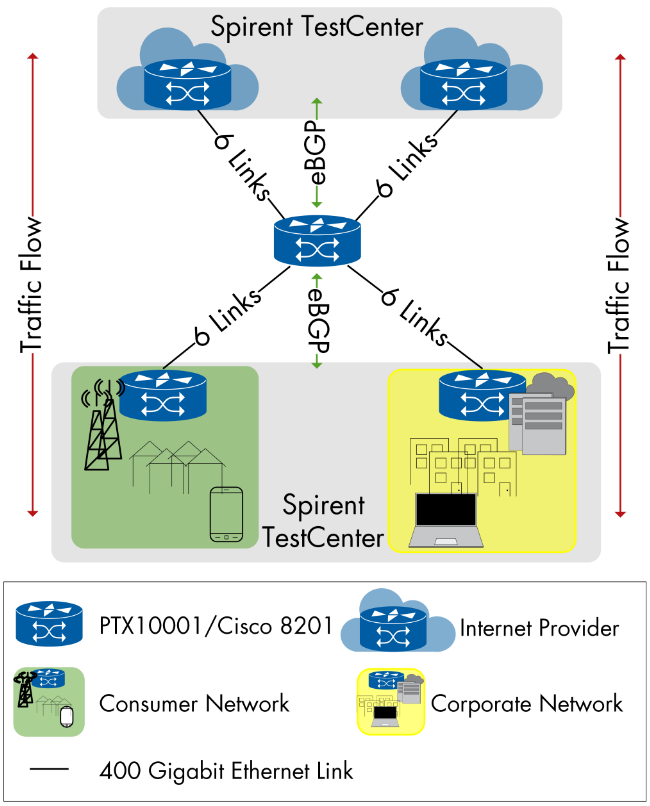

Figure 1: PTX10001-36MR and Cisco 8201 Throughput Test Bed

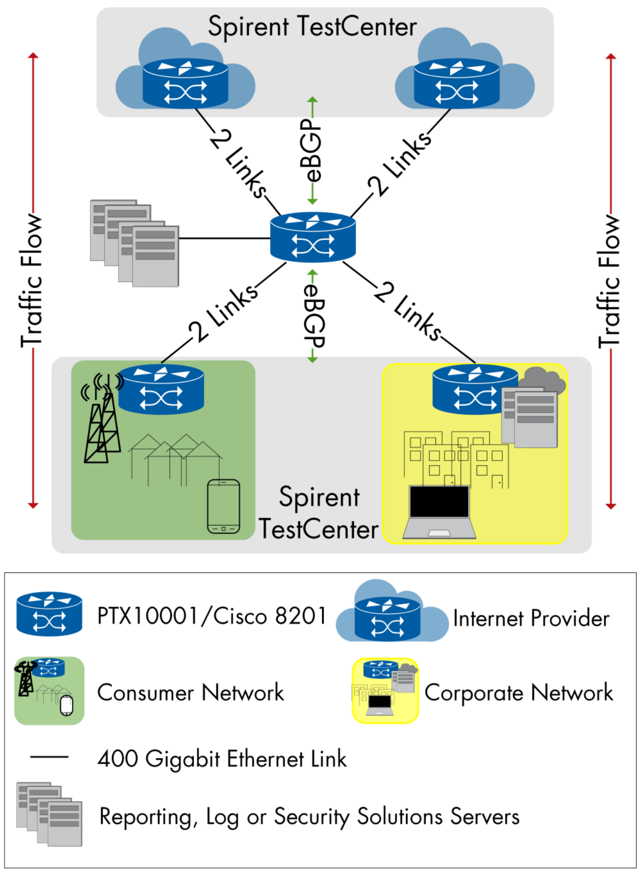

Figure 2: PTX10001-36MR and Cisco 8201 Functional Test Bed

Throughput Test

EANTC validated the throughput tests of the Juniper PTX10001-36MR and Cisco 8201 using the RFC2544 standards. In the test, Juniper connected the 24 400GbE ports of PTX10001-36MR and Cisco 8201 routers to the 400GbE Spirent TestCenter ports in two separate testbeds. EANTC configured EBGP peers to advertise routes from the Spirent TestCenter ports to the routers.

Initially, 4M IPv4 routes were advertised to both the routers from the EBGP peers of the Spirent TestCenter. With 4M IPv4 routes, hardware resource utilization was exceeded, and errors were observed on Cisco 8201. Later EANTC performed the RFC2544 test on Cisco 8201 with 2M IPv4 routes to monitor the behavior with half of the actual routes to be tested in comparison to Juniper PTX10001-36MR. For the IPv6 routes, EANTC advertised 2M routes and chose 99.99% load per port, which is 399.96Gbp/s, as the control packets can use the remaining 40Mbp/s of the port's bandwidth. IPv4 and IPv6 traffic with IMIX profiles mentioned in Table 2 were used for the test.

Two traffic runs for IPv4 and IPv6 were performed with the first run consisting of three trails for 120 seconds per trail, and the second run included a single trail for 15 minutes. The CPU usage was more than 100% for the Network Processing Unit (NPU) and the Forward Information Base (FIB) main processors on Cisco 8201, and on the PTX router, CPU usage was 99.9%. After the FIB learning, the CPU usage on both routers dropped to less than 50% and remained stable. The traffic was sent to all the advertised routes in a mesh topology, i.e., traffic was sent from each Spirent TestCenter port to all other ports of the router.

The pass or fail test results of RFC2544 on Cisco 8201 were inconsistent. We observed packet loss on Cisco 8201, and the tests failed with different traffic loads from 399.96Gbp/s up to 10.11Gbp/s for IPv4 or 72.03Gbp/s for IPv6 traffic in different iterations of the RFC2544 test. The test passed only with the 10.11Gbp/s IPv4 or 72.03Gbp/s IPv6 traffic load per port with 120 seconds iterations. On Juniper, PTX10001-36MR 399.96Gbp/s throughput was achieved without any traffic loss. The IPv4 and IPv6 throughputs and latency results of the Juniper PTX10001-36MR and Cisco 8201 are mentioned in Tables 3-6.

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 Routes	Traffic Load Intended per 400 GbE Port	Throughput Test per 400 GbE Port, Bi-direction Verdict	Packet Loss %	Packet Forwarding Performance per Port	Latency (μs) min/avg/max
IPv4	4M	399.96 Gbp/s	Passed	0.00%	61.00 Mpps	4.13/7.47/11.11
IPv6	2M	399.96 Gbp/s	Passed	0.00%	61.00 Mpps	4.22/7.72/11.19

Table 3: IMIX Three 120 seconds trails Throughput Results—Juniper PTX10001-36MR

IP Type	Number of Routes	Traffic Load Intended per 400 GbE Port	Throughput Test per 400 GbE Port, Bi-direction Verdict	Packet Loss %	Packet Forwarding Performance per Port	Latency (μs) min/avg/max
IPv4	4M	399.96 Gbp/s	Failed	-	-	-
	2M	399.96 Gbp/s	Failed	41.64%	-	-
	2M	10.18 Gbp/s	Passed	0.00%	1.56 Mpps	1.94/2.35/3.44
IPv6	2M	399.96 Gbp/s	Failed	-	-	-
	500,000	399.96 Gbp/s	Failed	0.22%	-	-
	500,000	72.05 Gbp/s	Passed	0.00%	11.05 Mpps	1.94/2.40/3.63

Table 4: IMIX Three 120 seconds trails Throughput Results—Cisco 8201

IP Type	Number of Routes	Traffic Load Intended per 400 GbE Port	Throughput Test per 400 GbE Port, Bi-direction Verdict	Packet Loss %	Packet Forwarding Performance per Port	Latency (μs) min/avg/max
IPv4	4M	399.96 Gbp/s	Passed	0.00%	61.00 Mpps	4.11/7.47/11.03
IPv6	2M	399.96 Gbp/s	Passed	0.00%	61.00 Mpps	4.19/7.72/11.54

Table 5: IMIX Running for 15 minutes Throughput Results—Juniper PTX10001-36MR

IP Type	Number of Routes	Traffic Load Intended per 400 GbE Port	Throughput Test per 400 GbE Port, Bi-direction Verdict	Packet Loss %	Packet Forwarding Performance per Port	Latency (μs) min/avg/max
IPv4	4M	399.96 Gbp/s	Failed	-	-	-
	2M	399.96 Gbp/s	Failed	41.35%	-	-
	2M	10.18 Gbp/s	Passed	0.00%	1.56 Mpps	1.94/2.35/3.46
IPv6	2M	399.96 Gbp/s	Failed	-	-	-
	500,000	399.96 Gbp/s	Failed	0.06%	-	-
	500,000	3.99 Gbp/s	Passed	0.00%	0.61 Mpps	1.93/2.20/3.14

Table 6: IMIX Running for 15 minutes Throughput Results—Cisco 8201

Route Scaling

A comparison test was also performed on the route scale capabilities of the Juniper PTX10001-36MR and the Cisco 8201. Different iterations and combinations of the IP prefix scale were tested to derive the maximum IP prefix scale values per router. During the test, we observed that Juniper PTX10001-36MR supports a maximum of 4M consecutive IPv4 prefixes and 2M consecutive IPv6 prefixes. On Cisco 8201, a maximum of 3.2M consecutive IPv4 prefixes and 1.3M consecutive IPv6 were observed. EANTC chose the Internet route table of 830,000 IPv4 BGP routes or 105,000 IPv6 BGP routes, which were available at the time of the test plan.

For a comparison test of IPv4 internet prefix scaling, EANTC configured four VRFs with 830,000 IP prefixes per VRF and advertised 3.32M routes to PTX10001-36MR. For Cisco 8201, when advertised four VRFs with 3.32M routes, the maximum internet IPv4 prefix routes learned was 2.7M. Of the 2.7M routes, 2.4M routes were learned with complete 830,000 routes of the three VRFs, and the remaining 0.3M routes were learned from the 830,000 routes of the fourth VRF. The 0.3M routes learned are not the same routes in every iteration, and the traffic simulation for the routes learned from the fourth VRF weren't feasible as different routes were learned in every iteration. Hence the test was performed using three VRFs, achieving 2.4M routes less than the maximum route learning capacity of the Cisco 8201.

For the IPv6 Internet IP prefix scaling test, EANTC configured four VRFs with 105,000 routes per VRF and advertised 420,000 routes to the PTX10001-36MR. Whereas on Cisco 8201, the IPv6 internet IP prefix scaling test was performed with a single VRF as the maximum number of IPv6 routes processed by Cisco 8201 was 15,000. Multiple hardware resources were utilized and observed the "Out of resource" processing error when advertised more than 30,848 IPv6 internet routes. Different routes advertised on Juniper PTX-10001-36MR and Cisco 8201 during the test are shown in Figure 3.

Juniper connected eight 400 GbE ports of the Juniper PTX10001-36MR and eight 400 GbE ports of the Cisco 8201 routers to the Spirent TestCenter in two different testbeds, as shown in Figure 2. For the consecutive IP prefixes scale test, EANTC configured two EBGP peers on the Spirent TestCenter. For the Internet IP prefix, we configured four EBGP peers for both Juniper PTX10001-36MR and Cisco 8201. The IPv4 and IPv6 consecutive prefixes scaling and Internet prefixes scaling was performed in four different iterations. The routes were advertised to the routers from the EGBP peers of the Spirent TestCenter, and the traffic was simulated for the advertised routes. For the test, 80% of port speed which is 320Gbit/s, was simulated. 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. IPv4 or IPv6 IMIX profiles mentioned in Table 2 for traffic simulation were used during the test. In order to observe the route learning capacity with different IPv6 prefixes, the IPv6 consecutive IP prefix lengths of /64 and /126 were tested in two different iterations. With the consecutive and internet IP prefixes scaling, the CPU usage of the BGP processor on the Juniper PTX10001-36MR router achieved 99.9% and on the Cisco 8201 router resulted in 100%. Traffic results of both routers observed are as shown below in Table 7.

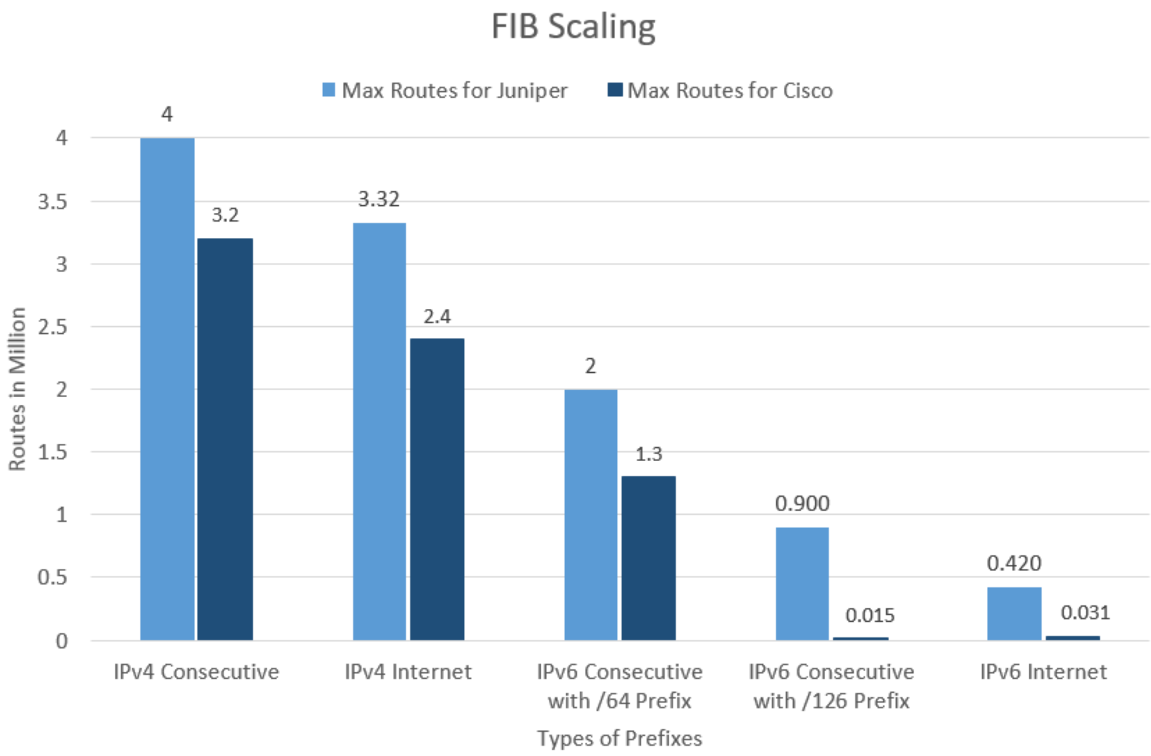

Figure 3: Maximum IPv4 and IPv6 FIB Scale

IP Prefix Type	Juniper PTX10001		Cisco 8201	
	Traffic Received per 400GbE Port	Packet Loss Percentage	Traffic Received per 400 GbE Port	Packet Loss Percentage
IPv4 Consecutive	320.00 Gbp/s	0.00%	319.88 Gbp/s	0.03%
IPv4 Internet	320.00 Gbp/s	0.00%	316.00 Gbp/s	0.97%
IPv6 Consecutive with /64 prefix	320.00 Gbp/s	0.00%	320.00 Gbp/s	0.00%
IPv6 Consecutive with /126 prefix	320.00 Gbp/s	0.00%	320.00 Gbp/s	0.00%
IPv6 Internet	320.00 Gbp/s	0.00%	320.00 Gbp/s	0.00%

Table 7: FIB Scale Traffic Test Results

Filter Scaling

In this test, EANTC observed the flexibility and capability of the filters on Juniper PTX10001-36MR and Cisco 8201, industry-wide known in Cisco as an Access Control List (ACL). For the Juniper PTX10001-36MR, the IPv4 and IPv6 filters are configured in profiles. Each profile has a set of match criteria parameters and actions for the matched packets.

Juniper stated that PTX10001-36MR 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. Juniper also mentioned no limit in configuring all the supported match criteria parameters and actions in any combinations. In the test, EANTC configured source IP address only, source and destination IP address, or source and destination IP address with port numbers as the match criteria parameters and validated the forwarding and counter statistics actions for the configured match parameters on PTX10001-36MR.

Each filter profile of the PTX10001-36MR 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 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 according to maximum filters, support for the consecutive IP prefixes was more than the random IP prefixes. Juniper PTX10001-36MR also supports the matching of 64,000 unique port numbers in addition to the IP prefixes per profile. In the filter profiles, Juniper supports individual ports or port ranges with consecutive port numbers. PTX10001-36MR treats the port range as a single entry and performs 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. EANTC configured the third profile with counter statistics and forward action 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.

In Cisco 8201, we understand that there are two general types of ACL, Standard, and Extended ACL. Standard ACL verifies only the source IP address of the packets, whereas Extended ACL demonstrates more attributes such as destination address, specific IP protocols, User Datagram Protocol (UDP), or Transmission Control Protocol (TCP) port numbers, and Differentiated Services Code Point (DSCP). In addition to these conventional types, there is another feature called Hybrid ACL; it can be used in the extensive scale network. EANTC configured source IP address only, source and destination IP address, or source and destination IP address with port numbers as the match criteria parameters. We validated the forwarding and counter statistics actions for the configured match parameters using the conventional ACL type.

In Cisco 8201, we observed the slicing architecture in which all the router ports are mapped to the six slices with four ports per slice. ACLs are configured per slice and are applied to all the ports of the particular slice. In order to maintain the same test bed and test configuration, ACLs on Cisco were configured on a single slice to which four 400GbE ports were mapped. During the test, observed for both consecutive and random IP prefixes, the behavior and a maximum number of Access Control Entries (ACE) that can be configured per ACL were the same. From the tests, maximum of 2480 IPv4 ACEs and 956 IPv6 ACEs can be configured per slice. For every IPv4 ACL with a port range of any number, 13 Ternary content addressable memory (TCAM) resources were utilized, resulting in shorting of total ACEs that can be configured per slice. But for the IPv6 ACL with port range, one hardware resource was utilized. Each ACL can be configured with a set of match parameters and actions to be performed on the matched packets. During the test, we observed Cisco 8201 could support configuration for two actions and three-match parameters. EANTC verified the accept action with match criteria parameters Source IP address, Destination IP address, and Port number. The ACL profiles on Cisco were configured with source and destination IP address combination or with source and destination IP address with port numbers. On Cisco 8201, port numbers can be configured in either port number per ACE or port range with consecutive port numbers per ACE. The port range can be configured from 1-65,535 per ACE. During the test, eight 400 GbE ports of Cisco 8201 and Juniper PTX10001-36MR were connected to the Spirent TestCenter with filters applied to the traffic ingress ports of the routers. IPv4 or IPv6 IMIX traffic profiles mentioned in Table 2 were used.

80% of port speed which is 320Gbit/s traffic, was simulated during the test. 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. The test simulated uni-directional traffic among four traffic pairs with the source and destination IP addresses and TCP port numbers configured in the filter profiles. For the packet match on the ingress port of the router, the filter counters were incremented.

Traffic was received successfully without any loss, and all the filter counters were incremented with a stable CPU usage of less than 15%. The filters configured during the test are shown in Table 8 and Table 9.

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

Table 8: Maximum Filter Scaling Statistics—Juniper PTX10001-36MR

IP Type	IP Prefix Type	No. of IP Prefixes Configured per Slice	No. of Source Addresses Configured	No. of Destination Addresses Configured	No. of TCP Port Numbers Configured	No. of ACEs Applied per Ingress Port
IPv4	Consecutive	2480	2480	2480	-	2480
IPv4	Random	2480	2480	2480	-	2480
IPv4	Random	1200	1200	1200	64,000	65,200
IPv6	Consecutive	956	956	956	-	956
IPv6	Random	480	480	480	-	480
IPv6	Random	480	480	480	64,000	64,480

Table 9: Maximum Filter Scaling Statistics—Cisco 8201

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

Sampling

Flow sampling is crucial for traffic composition analysis, validation of security policies, accounting 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 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 allow user to customize their sampling based on various packet fields using filters. 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.

In Cisco 8201, the sampled packets were exported to the cache initially from the NPU and then to the CPU for the processing of the packets and to create the flow records. The sampling with flow cache entries requires a timer at which the entries can be sent to the collector after the timer expiry. EANTC chose a flow cache timeout rate limit of 11,200 with a NetFlow ingress policer of 133Mbps bandwidth.

For the test, EANTC configured seven EBGP peers and one flow collector on the traffic simulator ports, connected with eight 400GbE ports of PTX10001-36MR and Cisco 8201 in two different test bed, as shown in Figure 2. EANTC advertised 2M routes to both routers for the sampling test. The sampling was applied at the traffic ingress ports of the PTX10001-36MR and Cisco 8201 router. The flow was monitored with a sampling rate of 1:2850 packets for benchmarking; in real-time, the sampling ratio is less aggressive, at 1:10000 or higher.

EANTC used the IMIX profile mentioned in Table 2 and simulated IPv4 traffic to the advertised routes for ten minutes. We simulated 99.99% of port traffic. For Juniper, we collected two samples of the flow information rate for every ten seconds from the router 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. In Cisco 8201, the counters of the cache entries and flows exported were collected from the console of the router. The cache entries counter gives the number of flow records cached, and the flows exported counter provides the numbers of flows sent to the collector. The timestamp of the cache entries counter was recorded after the start of traffic and until the cache table size was complete, which is 1M. Once the cache table was complete, the delta of the time stamp was calculated in seconds, and the supported cache table size was divided with the delta of the timestamp to define the number of flow records cached per second, i.e., setup rate. Two samples of exported flows were collected in two different timestamps. The delta of the flows exported counter from two samples is divided by the delta of the timestamps to define the export rate. The setup rate and export rate with two different sample ratios are shown in below Table 10.

Packets Sample Ratio	Juniper PTX10001		Cisco 8201	
	Setup Rate (Flow Packets per Second)	Export Rate (Flow Packets per Second)	Setup Rate (Flow Packets per Second)	Export Rate (Flow Packets per Second)
1:2850	150,940	150,922	55,555	11,200

Table 10: Sampling Results

Conclusion

During the comparison tests between the Juniper PTX10001-36MR and Cisco 8201, EANTC verified the throughput test, filter scaling, FIB scaling, and sampling.

PTX10001-36MR showcased stability and scalability in throughput performance with 4M routes compared to Cisco 8201 throughput performance with 2M routes. FIB scaling capabilities on the PTX10001-36MR with 4M consecutive IP prefix routes and 3.32M internet routes were remarkable with stable CPU usage compared to FIB scaling on Cisco 8201. The maximum range of the filters supported per 400GbE port in PTX10001-36MR was notable compared with the number of conventional ACLs supported by Cisco 8201. With the same packet sample ratio configured, PTX10001-36MR has sampled the flows two times greater than that on Cisco 8201.

Overall, the Juniper PTX10001-36MR passed all the tests which were designed for the real-time scaled network scenarios.

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