



Juniper Networks MX Series 3D with Junos Trio Chipset

Performance, Scalability and Power Efficiency Validation

Introduction

In September 2009, EANTC was commissioned by Juniper Networks to validate the density, scalability, services scale and power efficiency of the new 16-port 10 Gigabit Ethernet (16x10GbE) MPC line card. Additionally, the capabilities of the new Juniper Networks® MX80 3D router were demonstrated, as well as a set of provider edge services and capabilities, all focusing on its MX Series 3D routers.

EANTC collaborated with Juniper's technical marketing engineers in creating a detailed test plan that was executed in Juniper's lab in Sunnyvale, California. EANTC was able to independently validate the new product's capabilities prior to its public introduction and found that the line card delivered the performance and scalability advertised by Juniper.

Tested Devices and Test Equipment

The breadth of the testing and the demonstration required several topologies, devices and operating system configurations to be used for the project:

- MX480 3D router with six 16x10GbE line cards running Junos 10.0-20090914.0 (pre-beta)
- MX480 3D router with a single 16x10GbE line card running Junos 10.0-20090914.0 (pre-beta)
- MX80 3D router running Junos 10.1120090821 (pre-alpha)
- Functional demonstrations were carried out with up to four Juniper MX Series 3D routers (MX240 3D, MX480 3D, and MX960 3D routers) with a Multiservices DPC running Junos 9.6R1.13



Ixia's IxNetwork 5.40 was used to conduct all testing.

EANTC Validation Highlights

Density and Scale

- 1.44 Tbit/s (unidirectional) throughput in 1/6 rack size
- 240 Gbit/s (unidirectional) IPv4 and IPv6 forwarding on a single line card

Unicast Routing Performance

- 2.4 million IPv4 and 2.4 million IPv6 prefixes

Multicast Performance

- Line-rate forwarding for (64 through 1,518 bytes) on 71 10GbE ports
- 3,300 maximum outgoing interfaces on a single line card
- 60,000 multicast groups on each of the 11 10GbE ports
- 660,000 total receiver groups on a single line card

- Line-rate performance for 6,000 L3VPNs with 2.4 million unicast VPN routes on a single line card
- Line-rate performance for 6,000 VPLS with 240,000^a MAC addresses on a single line card

- 25.34 watts per 10GbE with line rate traffic
- 3.38 watts per gigabit of throughput

Throughput and Performance

Service Scalability

Power Consumption

a. Achieved with pre-beta software in the test. Juniper stated that this is not the system maximum.



Density, Forwarding, Scale

Two groups of tests were focused on the new 16x10GbE line cards introduced by Juniper. In the first set of tests, we focused our attention on the Juniper MX480 3D router hosting a single 16x10GbE line card. We characterized the performance and scalability of the new line cards for unicast (IPv4, IPv6, and a mixture of both). We also validated multicast performance, Layer 3 VPN (L3VPN), and virtual private LAN services scalability as well as routing performance of the new line cards.

Throughput Performance on a Fully Loaded Chassis

A new line card, especially one with such impressive specifications as a 10GbE with 16 ports, requires a detailed performance analysis. We started our analysis of Juniper’s 16x10GbE line card by investigating the performance of a fully loaded MX480 3D router.

The MX480 3D router is an eight-RU (rack units) router that can accommodate six line cards, in addition to two Routing Engines (REs). We used exactly this configuration as defined in the test plan: two RE-S-2000 Routing Engines and six 16x10GbE line cards. EANTC used standard methodology as defined in the Internet Engineering Task Force (IETF) Request for Comments (RFC) 2544 for IPv4 throughput measurements and RFC 5180 for IPv6 interconnected devices. For this test we attached all the ports on the system (i.e., 96 10GbE ports) to the Ixia tester. Full-meshed traffic between all the 96 10 GbE ports was created by the Ixia tester. Since we already knew that we could expect up to 240 Gbit/s unidirectional (120 Gbit/s bidirectional) traffic from each line card, we configured the ports to send traffic at 75-percent line rate, intending to validate the advertised forwarding performance of the 16x10GbE MPC. Figure 1 depicts a sample line card configuration.




Figure 1: Line Card Test Setup Example

For both IPv4 and IPv6 packet types, we found the system capable of delivering 1.44 Tbit/s of full-mesh unidirectional traffic (720 Gbit/s bidirectional) at all tested frame sizes (from 64 bytes to 1,518 bytes) with zero packet loss.

For service providers the implications of our findings are significant. Up until now, Juniper’s line cards could accommodate up to 24 10GbE ports in the same form factor. With the introduction of the new 16x10GbE line card, we see a high-performance card that packs three times the forwarding performance and four times the port density in comparison to the previous offered configurations.

Density, Forwarding, Scale Test Highlights

- 1.44 Tbit/s unidirectional throughput
- Zero packet loss from 1 Billion+ pps

Multicast Performance Characteristics

IP multicast plays a pivotal role in most IPTV deployments, financial applications, and content distribution systems due to its highly efficient network resources usage. The tests performed in this section characterized the 16x10GbE line card multicast replication capabilities, group capacity, and forwarding performance for mixed unicast and multicast traffic.

Multicast Forwarding and Performance on a Fully Loaded Chassis.

The most challenging multicast test setup was an MX480 3D router with all 71 ports on each of six 16x10GbE cards requesting multicast traffic and with only one port as the source. In this setup, the majority of the traffic had to cross the back-plane and then further replicate on the individual line cards.

In this test, we defined the first port on the system as a multicast source, sending to 1,200 multicast groups. All remaining 71 ports joined all groups. We ran a multicast throughput test for 64-, 73-, 128-, 512-, 1,024-, and 1,518-byte frames and expected to register zero frame loss for all frame sizes on all multicast receiving ports.

We ran the tests with all the frame sizes for the duration of 120 seconds for each test iteration and validated that the MX480 3D router, with its fully populated line-card configuration, could replicate multicast traffic at line rate with zero packet loss.

The results of the test show that the new line cards are capable of high-performance multicast replication. Sending such small frame sizes as 64 and 73 bytes requires maximum multicast replication performance from a device, and the MX480 3D router had no problems delivering. For major multicast distribution centers or extremely heavy financial trading applications, these results reassure that the router can fulfill their needs.

Tests Setup. The 16x10GbE line card supports 120 Gbit/s bidirectional (240 Gbit/s unidirectional) forwarding performance. The rest of the tests described in the report focused on the 16x10GbE line card and used 12 out of 16 ports operating at line rate. Figure 2 depicts the configuration used for the tests in the following sections.

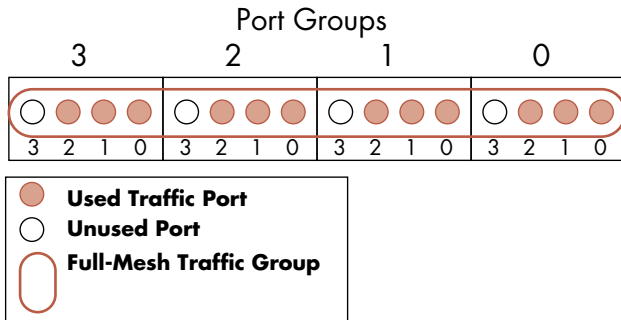


Figure 2: Single Line Card Test Setup

Mixed Multicast and Unicast Forwarding. We ran an additional multicast test using a mixture of multicast and unicast traffic. One port was used to source multicast traffic while the rest of the ports (on a single 16x10GbE line card) were sending and receiving full-mesh unicast traffic. This scenario mimics closely real-world scenarios where multicast and unicast are sent across interfaces at the same time.

We were able to verify that all interfaces could forward 50% multicast and 50% unicast traffic for all packet sizes tested. Figure 3 depicts the results for a single line card.

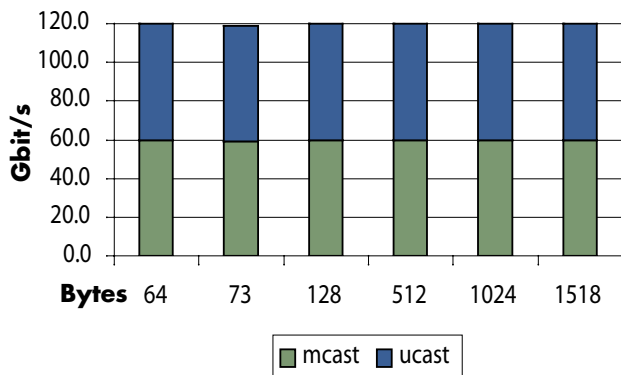


Figure 3: Mixed Multicast and Unicast Throughput

Outgoing Interfaces Scalability. Juniper positions the MX Series platforms as aggregation and services edge routers. An MX Series router, therefore, is often required to serve as the last-hop router (the router responsible for multicast traffic replication). For example, in a Triple Play deployment scenario, an MX Series router will aggregate a large number of

DSLAMs and must replicate the IPTV streams to all of them. For this reason, the number of outgoing interfaces (OIF), multicast groups, and receivers that can be supported is a critical aspect of the new line cards.

We used two 16x10GbE line cards for this test. One line card was used to source 8 multicast groups, each at 3.98 Mbit/s (roughly standard definition channel equivalent). On the second line card, we configured 300 VLANs on 11 receiving ports, each receiving all 8 groups. With 8 groups each sending 3.98 Mbit/s of traffic in each VLAN we were expecting 31.8 Mbit/s per VLAN. All together, with 300 VLANs (and their headers) on a 10GbE port, we observed 9.98 Gbit/s of traffic exiting each port while each VLAN had receivers for all 8 groups. This meant that the system was receiving 26,400 IGMP joins and replicating the 8 multicast groups to 3,300 outgoing interfaces all on a single line card.

The results met the test’s goals. We were able to verify that all multicast receivers in all VLANs on all interfaces were receiving all 8 streams with no packet loss. We reached 9.98 Gbit/s on the receiving interfaces even when facing the most aggressive packet replication scenario of 64-byte frames.

The test results demonstrated that the 16x10GbE line card is highly capable at multicast replication and is able to scale its outgoing interfaces. While the number of VLANs used were on the high side for IPTV deployments, media-rich financial applications, in which various traders are split across VLANs, could greatly benefit from such scalability seen in the test.

Multicast group scalability. The goal of the test was to measure the number of multicast groups that an MX480 3D router with two 16x10GbE line cards could support. The test goal was simple: increase the number of multicast groups while forwarding traffic without losing a single frame.

RFC 3918, which defines multicast benchmarking methodology, dictates that to find the maximum group number a system can service, the number of groups be increased as long as no loss of frames is recorded on the system under test. Instead of going through this tedious exercise, we asked Juniper to provide a target value of multicast groups that EANTC should validate.

Juniper provided an impressive target value of 60,000 multicast groups. With 11 ports, all requesting 60,000 groups, we expected to register 660,000 IGMP requests on the system and then forward packets, at various frame sizes (from 64 to 1,518 bytes) at line rate.

We sent the IGMP join requests on all receiver ports and validated that the system indeed registered all 660,000 requests. We then started the multicast traffic for all groups and let traffic run for 120 seconds. We recorded zero packets loss in all test runs for all frame sizes.

This IGMP scaling figure far exceeded today’s application requirements. The number of groups reached provide ample headroom for innovative, media-rich financial and transactional (stock ticker) services for high-end financial institutions in the next five to ten years. Such high scaling values validate that the MX480 3D router with the 16x10GbE line card is, from a multicast group scalability perspective, future-proof and provides headroom for future multicast-intensive applications.

Multicast Performance Test Highlights

- Line-rate multicast forwarding performance
- 3,300 outgoing interfaces on a single line card
- 60,000 IGMP groups with a single source
- 660,000 multicast receivers on a single line card

Services Scale

Forwarding performance, multicast replication, and scalability are the building blocks with which service providers can offer services to their customers. Once we established a baseline for the performance of the new line cards, we investigated their services scalability.

BGP scalability. We started by performing two throughput tests, each with an aggregate of 2.4 million BGP routes: one test for IPv4 and one for IPv6. Both tests used the same hardware, which included an MX480 3D router, a single 16x10GbE line card, and two Routing Engines.

We configured 12 BGP peers, one for each of the 12 ports tested, and exchanged 200,000 prefixes on every port. The IPv6 test used /64 routes. Two IPv4 routes configurations were used: one with /30 routes and the second with a more diverse set ranging from /24 to /29 routes. Once we established the BGP neighbors between the tester and router ports, we advertised the routes and sent traffic for all prefixes. We repeated all three tests for 64-, 73-, 128-, 512-, 1,024- and 1,518-byte packets. The IPv6 packets sizes were almost identical apart from the smallest packets, which were 82 bytes.

We expected the router to be able to forward traffic to all routes at all frame sizes for both IPv4 and IPv6. The

results met our expectations. For starters, the BGP peers remained stable for the duration of the tests. We were also able to send traffic at line rate for all prefixes without losing a single frame.

L3VPN scalability. We used the same hardware configuration to perform the L3VPN and virtual private LAN service (VPLS) scalability tests. We divided the 12 ports into two sets: 6 were configured as customer-facing and 6 as backbone-facing. On each of the six customer-facing ports Juniper engineers configured 1,000 VLANs, each representing a customer and each establishing an eBGP neighbor with the tester. In total we had 6,000 emulated customer connections set up for the test. Attached to the six ports we defined as upstream (i.e., network-facing), we used Ixia’s IxNetwork to emulate six provider (P) routers and six provider edge (PE) routers. Each of the emulated PE routers was configured with the corresponding Virtual Routing and Forwarding (VRF) instance to match the customer attached to the downstream ports.

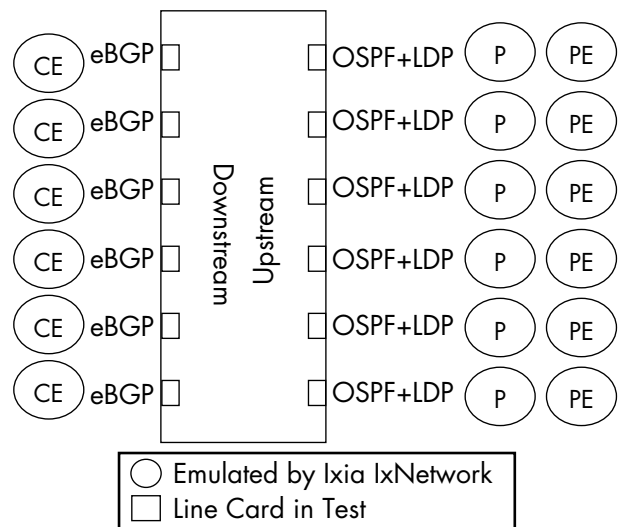


Figure 4: L3VPN Scalability Test Setup

Each of the emulated 6,000 customers advertised 200 prefixes on each end of the VRF, which equaled to 2.4 million routes. Once all eBGP sessions were established and routes were exchanged, we performed a throughput test using a common Internet mix (Imix) of frame sizes: 64 bytes, 576 bytes, 1,400 bytes and 1,470 bytes.

The router showed no hesitation in establishing the eBGP sessions and maintaining them. The 2.4 million VRF routes were learned and installed in the forwarding table after several minutes and remained unchanged for the duration of the test. We sent traffic to all VRF prefixes using the Imix described above and recorded zero packet loss from all 2.6 billion packets sent.

The test demonstrated the scalability possibilities for service providers to deploy MPLS-based L3VPNs on Juniper’s MX Series 3D platform with the 16x10GbE line card. The tested service density would allow a service provider to offer 6,000 20 Mbit/s VPN endpoints on a single line card. This demonstrates true economy of scale for deployments in heavily populated small and medium business areas, such as New York City and Hong Kong.

VPLS scalability. In recent years, Ethernet-based services have been experiencing growing success, due in part to the success of the Metro Ethernet Forum (MEF). VPLS is one way to offer multipoint-to-multipoint Ethernet connectivity to customers wishing to connect several sites together in a simple way. Since the provider network acts as an Ethernet switch in VPLS deployments, two aspects are critical for the success of the service: the number of virtual switch instances (VSIs) and MAC addresses to which the routers can scale.

This test was configured in a similar way to the L3VPN test: 6 ports were defined as customer-facing, each configured with 1,000 VLANs, and 6 ports were defined as network-facing. Each emulated customer was sourcing Ethernet traffic from 20 MAC addresses, which meant that in total 20,000 MAC addresses had to be learned on each customer-facing port and each network-facing port. This totaled 240,000 MAC addresses on a single line card.

We repeated the same procedure and configuration that was used in the L3VPN test. Again we performed a throughput test using the Imix previously described, and monitored for lost frames.

And the results did not disappoint. None of the millions of frames sent in the test were lost under these extreme traffic load and service scale conditions.

Services Scale Test Highlights

- 6,000 L3VPN instances on a single line card
- 2.4 million unique active VPN routes
- 6,000 BGP sessions
- Zero packet loss
- 6,000 VPLS instances on a single line card
- 240,000 MAC address per line card
- Zero packet loss

Both VPLS and L3VPN tests showed that network providers using Juniper’s new 16x10GbE line card can expect impressive scalability from their hardware.

Network providers are free to choose between the two services and can expect the same performance regardless of their choice.

MX80 3D Functional Validation

Juniper introduced us to a new, 2-RU platform called the MX80 3D router. The router comes built with a fixed Routing Engine and Forwarding Engine Base (FEB), as well as built-in 4x10GbE ports. The MX80 3D router also has two pluggable Media Interface Card (MIC) slots (not used in these tests). Juniper explained that the MX80 3D router is targeted at providers interested in extending their service reach through a distributed service delivery model; for example, a VPLS multi-tenant unit (MTU) platform.

As a member of the Juniper MX Series 3D family, the MX80 3D router, with its lean form factor, is also able to support L3VPNs and VPLS. In this functional demonstration we verified that the new platform is able to provide the same two services that we tested with the larger 8-RU MX480 3D router using the 16x10GbE line card.



Figure 5: Juniper MX80 3D

L3VPN Services. Much like the rest of the product family running Juniper Networks Junos® software, the MX80 3D router is also a capable router. We repeated the procedure described above with the MX480 3D router, reducing the number of customer-facing ports to two and the network-facing ports to two. These four ports are, after all, an integrated part of the device and hence users could expect to use them with the for services.

We configured 500 VLANs on each of the customer-facing ports and set up an eBGP neighbor in all of them. Each eBGP neighbor advertised 1,000 routes, which meant that in total the MX80 3D router had to support a million VPN routes. We then sent line-rate traffic to all the routes using the same Internet traffic mix defined above and verified that no frames were lost for line-rate traffic.

Despite its compact 2-RU footprint, the MX80 3D router was not phased by the number of BGP sessions it had to maintain, the packet forwarding it had to process, or the packet rate. We recorded no frame loss.

VPLS Services. We subsequently tested VPLS services on the MX80 3D router. A similar test configuration was used again, this time with 250 VLANs on each of the two customer-facing ports and with 20 MAC addresses on each VSI endpoint. This configuration totaled 500 VPLS instances and 20,000 MAC addresses in a 2-RU form factor. Juniper informed us that this does not represent the system’s maximum MAC address capacity as we did not investigate the size of the MX80 3D MAC table.

We used the same traffic mix to verify that all VPLS instances were able to switch frames and serve the 500 emulated customers. Once again, we were not disappointed. The router functioned as expected, learning the required number of MAC addresses, and switching all frames at line rate.

Both tests demonstrate Layer 2 and Layer 3 service capabilities in a compact form factor that can comfortably cater to service-rich distributed PE and MTU deployments.

MX80 3D Functional Validation Highlights

- 1,000 L3VPNs, 1 million VPN routes
- 500 VPLS instances with 20,000 MAC addresses
- Line-rate Imix forwarding for both services

Power Efficiency

The amount of energy (or power) a system uses to transport data has been a topic of growing interest in recent years. Energy efficiency translates to savings in energy costs, which leads to reduction in operating costs and in some countries, to healthy government subsidies.

Test procedures and measurement methodologies are defined by several industry organizations, such as the Alliance for Telecommunications Industry Solutions (ATIS) and by the Energy Consumption Rating (ECR) Initiative. The latter is a partnership between Juniper Networks, Ixia Corporation, and Lawrence Berkeley National Lab. For this test, we used the definition provided by the ECR document (version 1.0.4, November 2008).

The procedure defined by the ECR consists of four steps for determining the energy consumption of a system when forwarding packets. The first step is to determine the maximum "zero-packet-loss" throughput performance of the system. This step was performed by our extensive investigation of an MX480 3D system fully loaded with 16x10GbE line cards. The next three aspects of the specification define the specific system

loads that must be used while power is measured. The traffic loads specified are 100, 50, and 0 percent, measured at 20-minute interval each.

To accurately measure the power drawn by the router, we connected the four Power Entry Modules (PEMs) to a power distribution unit, which allowed us to obtain power usage measurements while running the test traffic. We used Ixia’s IxNetwork to generate test traffic and measured the power output while sending test traffic for 20 minutes. Three power draw measurements were performed, once with 720 Gbit/s full duplex (100% load), once with 360 Gbit/s full duplex (50% load), and once with no traffic load.

The results are ECR values that describe the amount of energy consumed by the device in order to move one gigabit of line-level data per second. We found that the ECR ratings per Gbit/s was 3.38 watts and that the average power consumption per 100-percent load was 2433.6 watts.

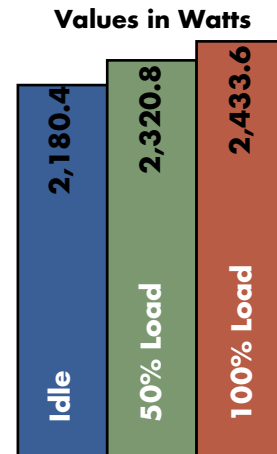


Figure 6: Full Chassis MX480 3D Power Consumption

Functional Demonstrations

Juniper’s MX Series 3D product portfolio is positioned as aggregation and edge routers in the company’s arsenal. As such, the routers must often satisfy a wide range of requirements that are specific to the network edge. These requirements range from subscriber management to supporting metro rollouts and beyond. Juniper used this opportunity to demonstrate a range of applications based on the MX Series 3D portfolio using a pre-released version of the Junos operating system and 4x10GbE line cards.

NAT Services for Business VPN Customers.

Business customers require VPN connectivity services between multiple sites while often requiring Internet access. BGP/MPLS IP VPNs deliver intra-site connectivity sometimes using private addressing schemes to conserve public addresses. Since these addresses are not Internet-routable, access to the Internet becomes an issue for business VPN customers using such addresses.

Network Address Translation (NAT) is a solution for this problem. NAT is a process that is often performed by specialized devices that modify network addresses in IP packets from one address space to another. NAT could be done on customer premises equipment (CPE) or firewalls, but why not perform the same service by the same network provider offering the VPN services?

Juniper set out to demonstrate just this idea: move the NAT solution to the network and implement it on the PE router, the same router that is the business customer’s point of entry to the network.

In this demonstration, Juniper used the Multiservices Dense Port Concentrator (MS-DPC) to perform the NAT service function. NAT service is only available with stateful firewall. EANTC configured the tester to emulate two VPN endpoints attached to PE1 and PE2 (as seen in the figure below), and another tester port was configured as an Autonomous System Boundary Router (ASBR) advertising Internet routes to the network.




Figure 7: Wholesale VPN Service with NAT

EANTC was then able to observe traffic sent between the two VPN endpoints and verify that all traffic was received with no packet loss. Once the VPN traffic was verified, we monitored the route count on the MS-DPC and verified that upon starting to send and receive traffic from the emulated Internet access port, all routes were registered.

DHCP Broadband Subscriber Access to L3VPNs.

The functionality demonstrated in this test allows service providers to offer carrier’s carrier services. When service providers lack geographical reachability, they often buy access from a second service provider. VPN customers that belong to the first provider must, therefore, traverse a different network before reaching the service provider that is actually providing them with the VPN.

In this demonstration, DHCP subscribers (emulated by the Ixia tester) were attached to a router that belonged to one service provider (PE1). PE1 served as a DHCP relay agent that intercepted the DHCP requests. PE1 then associated each incoming DHCP request (based on the port) to a pre-configured user prefix and sent a request to its own Radius server. The Radius server then authenticated the request against the service provider’s database. The authentication instructed PE1 to forward the DHCP request to an L3VPN.

In the wholesale scenario, we validated that the DHCP request was authenticated and the subscriber admitted into the VPN. The actual IP address assignment to the emulated customer came from the DHCP and Radius servers that belonged to the retail service provider, using the wholesaler’s access infrastructure.

The main focus of the test was to validate the functionality delivered by PE1. The router had to intercept the DHCP requests delivered from the 32,000 emulated active subscribers that were configured. It then had to verify that the DHCP request was allowed to enter the network, receive information from the Radius server regarding what to do with the DHCP request (in this case, forward it to the VPN), and then send the DHCP request further.

Even though we used 32,000 concurrent active subscribers in the demonstration, the router had no problem performing the actions described above and assigning VPN membership upon authentication.

Metro Ethernet Functional Demonstrations

Having participated in two demonstrations that were focused on L3VPN services, we proceeded to two functional validations that were centered around metro Ethernet rollouts. Juniper demonstrated a mechanism that can rid provider networks from spanning tree protocols and improve Ethernet resiliency, while removing the requirement for backbones to learn a large number of MAC addresses.

PBB-VPLS MAC address hiding for BGP-VPLS.

Provider Backbone Bridge (PBB) allows service providers to offer global Ethernet services for existing Layer 2 islands without running into MAC address scalability issues. In essence, what PBB brings to the table is the ability to hide the MAC addresses belonging to these Layer 2 islands from the rest of the network. In addition, PBB interworking with VPLS helps to further contain the MAC table size when multipoint access is extended beyond the metro with VPLS and/or H-VPLS.

Juniper demonstrated the ability to map encapsulate Ethernet traffic to IEEE provider backbone bridge (PBB, IEEE 802.1ah-2008) and forward the traffic into VPLS instances.

Multichassis LAG for cross-chassis resiliency.

Using multichassis link aggregation group (MC-LAG), service providers, cable companies, and enterprises can offer customers a cross-platform resiliency mechanism for their Layer 2 services. MC-LAG, combined with VPLS, removes the service provider's reliance on spanning tree protocols while providing both resiliency and loop prevention mechanisms.




Figure 8: Multichassis LAG Test Topology

Juniper demonstrated MC-LAG using the configuration described in Figure 2. A link aggregation bundle was configured on the customer edge device (a Juniper

MX960 router, in this case) with two link members each terminating different VPLS-PE routers (PE3 and PE4). The two VPLS-PE routers used the Inter-Chassis Control Protocol (ICCP) to exchange status information, in addition to the Bidirectional Fault Detection (BFD) protocol.

To verify that the MC-LAG worked as a resiliency mechanism, we disconnected the active link bundle member while sending bidirectional Ethernet traffic through the test topology. We then measured the amount of frames that were lost due to our simulated link failure. We used this measurement to deduce that the service was only briefly interrupted and successfully switched over to the alternate link member of the MC-LAG without the use of any spanning tree protocol.

In this functional demonstration EANTC validated that MC-LAG delivered on Juniper's promise of Ethernet resiliency without the use of spanning tree protocols.

Conclusions

The capabilities enabled by the Junos Trio chipset-based 16x10GbE MPC line card and MX80 3D router were very impressive. EANTC sees Juniper pushing new boundaries concurrently for Ethernet density, services, and subscriber scale, as well as power efficiency.

The fact that the Junos Trio chipset and the same Junos operating system is supported across multiple product lines enables synergies for network operators. They can now support a common set of functionalities throughout the core, aggregation, and edge parts of the network.

About EANTC



The European Advanced Networking Test Center (EANTC) offers vendor-neutral network test services for manufacturers, service providers and enterprise customers. Primary business areas include interoperability, conformance and performance testing for IP, MPLS, Mobile Backhaul, VoIP,

Carrier Ethernet, Triple play, and IP applications.

EANTC AG Einsteinufer 17, 10587 Berlin, Germany
info@eantc.com, <http://www.eantc.com/>

V1.2, 20091026