

IBM BNT RackSwitch G8264:

Competitive Performance Evaluation vs. Arista Networks 7148SX, Cisco Systems, Inc. Nexus 5548P, and Juniper Networks EX4500

Executive Summary

Today's data centers are serving content for fixed and mobile clients across large scale public and private cloud computing clusters. Video rendering, high frequency trading, and oil reservoir simulations are just a few examples of the complex applications being served over computing clusters connected at 10 Gigabit Ethernet (10GbE). Multi-core CPU architectures, coupled with virtualization in the data center, enable multiple workloads to run on the same machine with greater bandwidth and are driving demand for more bandwidth per server. 10GbE Top-of-Rack (ToR) switches are now being designed to support consolidated networking and storage traffic as well as server virtualization. BLADE Network Technologies, recently acquired by IBM, commissioned Tolly to benchmark the IBM BNT RackSwitch G8264 against three competing 10GbE switches for latency, throughput, power consumption, microburst buffer capacity, and price/performance.

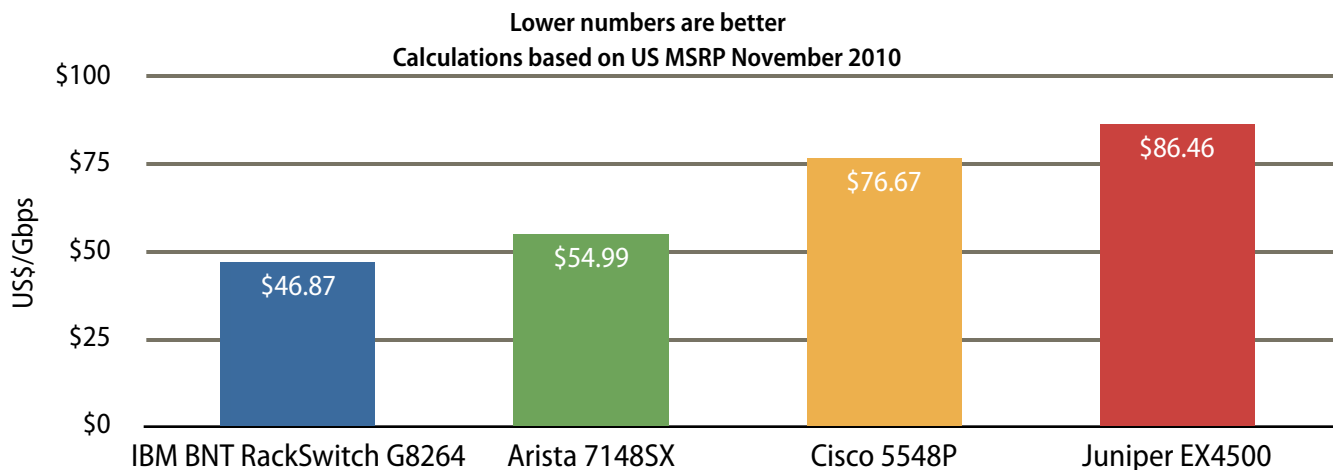
The IBM BNT RackSwitch G8264 includes a standard configuration of 48 SFP+ 10GbE ports and four 40GbE ports. These four QSFP+ ports can also be used as four 10 GbE ports each for a total of 640 Gbps of full-duplex lossless throughput, 160Gbps more than the other switches tested. The IBM BNT RackSwitch G8264 supports Converged Enhanced Ethernet (CEE)/ Fibre Channel over Ethernet (FCoE) and network virtualization, along with a robust Layer 2 and Layer 3 feature set.

The Bottom Line

The IBM BNT RackSwitch G8264 demonstrated:

- 1 Up to 11.5 times lower latency
- 2 Up to 100 times more buffer capacity
- 3 Up to 71% less power consumption
- 4 Up to 84% better price performance
- 5 160 Gbps of additional, line rate throughput via its 4x 40Gbps ports

Top-of-Rack 10GbE Switch Price/Performance Comparison



Note: Switch pricing obtained from public sources in November, 2010. Price calculated based on IBM BNT G8264 48x10GbE and 4x40GbE Switch, Arista 48Port 7148SX 10GbE Switch, Cisco Nexus 5548P 1RU Chassis 2 PS 2 Fan 32 Fxd 10GE with Nexus 5500 Module 16P 10GE, and Juniper EX4500 40-Port 10G SFP+ Switch with two 4-Port 10G SFP+ Uplinks and one 1200W AC power supply. Performance based on theoretical max throughput for all switches evaluated. See Table 1 for details.

Source: Tolly, February 2011

Figure 1



Introduction

10GbE Top-of-Rack Data Center switches are expected to provide line rate throughput, low latency, low power consumption, and be able to buffer bursts of traffic, ensuring lossless performance in networks with "bursty" traffic patterns. A switch should perform at line rate across all frame sizes. Forwarding anything less than 100% of frames when required can cause retransmission of data, potentially resulting in degraded application performance.

Even worse, depending on the type of protocol, as in multicast applications, frames can be dropped if line rate throughput is not achieved.

The IBM BNT RackSwitch G8264 consistently demonstrated 100% line-rate throughput, lower latency, and the capability to buffer up to 6 times more packets than the Arista 7148SX, up to 50 times more packets than the Cisco 5548P, and up to 100 times more packets than the Juniper EX4500, while providing an additional 16 10GbE ports, or 160Gbps more capacity than all other switches tested.



BLADE Network Technologies, an IBM Company



RackSwitch G8264

10GbE Performance Evaluation

Tested February 2011



Note: IBM BNT RackSwitch G8264 tested using 64 10GbE ports. It would appear that Arista segments its switch into port blocks, limiting its forwarding capability at line rate when traffic crosses switch segment boundaries. Juniper EX4500 latency invalid for 64-byte frames.

Source: Tolly, February 2011

Figure 2



In addition to the performance advantages, the RackSwitch G8264 consumed up to 70.6% fewer Watts per Gbps of throughput than the Arista 7148SX, up to 56.9% less Watts per Gbps than Cisco 5548P, and up to 50.7% less than the Juniper EX4500.

The IBM BNT RackSwitch G8264 also demonstrated an average of 55% better price/performance than the three competing switches.

Test Result Summary

Layer 2 Performance

The devices under test (DUTs) were tested for their performance while forwarding Layer 2 traffic. This scenario represents one of the most common deployment modes of the DUTs in data center environments.

Engineers benchmarked the DUTs in terms of throughput and latency while handling bidirectional Layer 2 traffic from 64 to 9216-bytes in size as defined by RFC 2889, RFC 2544, and Ixia's ATSS (Advanced TCL Test Suite). Throughput and latency tests

Interaction with Arista Networks, Cisco Systems, Inc. and Juniper Networks

In accordance with the process for conducting comparative tests, Tolly contacted each vendor to notify them of the evaluation and invite their participation. Management representatives of Cisco Systems, Inc. and Juniper Networks did not respond to Tolly's invitation. Representatives from Arista Networks reviewed the test plan and provided comments which were incorporated into the test methodology where applicable. After reviewing the results, Arista informed Tolly that the latency recorded for 9216-bytes was higher than expected but Arista was unable to reproduce the issue by the publication date.



For more information on the Tolly Fair Testing Charter, visit: <http://www.tolly.com/FTC.aspx>

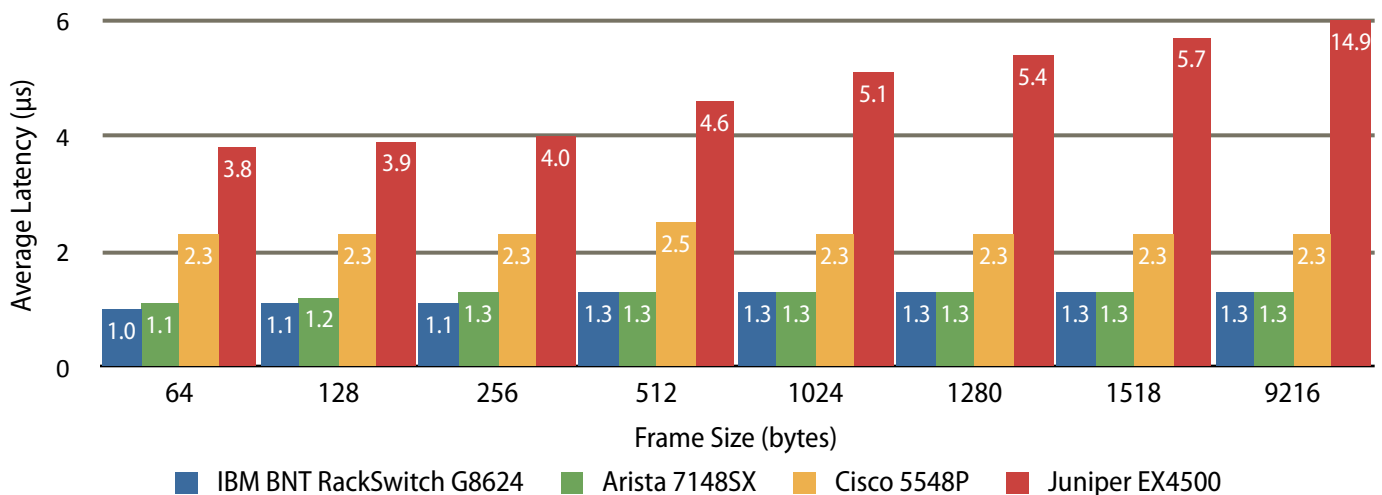
consisted of 48 10GbE ports (64 ports on the BNT RackSwitch G8264) on the DUTs configured in a full-mesh and port-pair topology.

RFC 2889 Full-Mesh Zero-loss Throughput
The RackSwitch G8264 demonstrated 100% throughput for 64, 128, 256, 512, 1024,

1280, 1518 and 9216-byte jumbo frames while utilizing all 64 10GbE ports.

As the RackSwitch G8264 also supported 40GbE ports, an additional test was run only on the G8264 where engineers configured the 48 fully meshed 10GbE ports, as well as four 40GbE ports in a dual-mesh configuration. Again, the IBM BNT

Layer 2 Multicast Latency with one Tx Port to 47 Rx 10GbE Ports
(Lower numbers are better)



Note: IBM BNT RackSwitch G8264 tested with 64 10GbE ports. All others tested with 48 Ports.

Source: Tolly, February 2011

Figure 3



RackSwitch G8264 demonstrated line rate, zero-loss throughput at all frame sizes, while delivering an average latency within the 40 GbE mesh of 853 nanoseconds. Due to the nature of QSFP+ ports, this sub-microsecond latency could also be expected when utilizing the 4x 10GbE breakout cables, as long as the traffic originated from and was destined for one of the QSFP+ ports. See Figure 6 for details.

The Arista 7148SX achieved 92.3% throughput for all frame sizes in the full-mesh testing, except for 9216-byte frames which yielded 85% throughput.



The Cisco 5548P demonstrated 100% throughput for all frame sizes from 64-bytes to 9216-bytes while using 32 fixed 10GbE ports and a 16-port 10GbE expansion module.

The Juniper EX4500 exhibited an apparent bug with their default configuration that would cause less than 1% throughput at 64-bytes unless a storm suppression command from the default configuration was removed, even though the latest available software, 10.4R1.9, was used. With this work-around applied, the Juniper EX4500 throughput was measured at 50.5% for 64-

byte frames, 88.9% through for 128-byte frames and 100% for the remaining frame sizes.

RFC 2544 Port-to-Port Throughput

Engineers benchmarked the DUTs for throughput and latency while handling bidirectional Layer 2 traffic from 64 to 9216-bytes in size as defined by RFC 2544. The port to port throughput test consisted of 48 10GbE ports (64 ports on the BNT RackSwitch G8264) on the DUTs configured in a port-pair topology with the ports




Note: Cisco 5548P did not support Layer 3 forwarding and hence was not tested. RackSwitch G8264 tested with 64 10GbE ports, all other switches tested with 48 ports. It would appear that Arista segments its switch into port blocks, limiting its forwarding capability at line rate when traffic crosses switch segment boundaries.

Source: Tolly, February 2011

Figure 4

Microburst Tolerance: Oversubscribed Link Packet Buffering



Source: Tolly, February 2011

Figure 5

arranged from the outside ports inward, i.e. 1 to 48, 2 to 47, 3 to 46, etc.

The RackSwitch G8264 and Cisco 5548P demonstrated consistent 100% throughput for 64, 128, 256, 512, 1024, 1280, 1518 and 9216 byte frames. The RackSwitch G8264 utilized all 64 10GbE ports whereas the Cisco 5548P utilized its 32 fixed 10GbE ports with the use of the optional 16-port 10Gb Ethernet expansion module.

The Arista 7148SX was only able to achieve 33.3% throughput in this port-pair test. To be fair, engineers attempted to use pseudo random MAC addresses (switch emulation) to coerce better throughput, but the performance remained constant. The Juniper EX4500 achieved 50.5% on 64-byte frames, 88.9% on 128-byte, 99.9% on 256-byte, and 100% on remaining frame sizes. See Figure 2 for the comparison of Layer 2 throughput results.

ATSS Port-to-Port Latency


The Arista, Cisco, and Juniper DUTs were tested for the cut-through latency using 48 10GbE ports connected in port-pairs, from the outside ports inward, topology

handling a load of 100% of the 10GbE line rate for Cisco and BLADE. Other devices were tested at their maximum achievable rate for each DUT based on the IxAutomate ATSS test suite. The IBM BNT RackSwitch was tested using all 64 available 10GbE ports. Despite handling 16 additional 10GbE ports worth of traffic, the BNT RackSwitch G8264

achieved the lowest latency of the switches tested.

The RackSwitch G8264 measured from 1.08 µsec for 64-byte frames to 1.41 µsec for 9216-byte frames (average latency: 1.3 µsec).

IBM BNT RackSwitch G8264 Layer 2 Throughput and Latency with four 40GbE QSFP+ ports in Port-Pairs



Note: Due to the nature of QSFP+ ports, this sub-microsecond latency can also be expected when utilizing the 4x 10GbE breakout cables, as long as the traffic originated from and was destined for one of the QSFP+ ports.

Source: Tolly, February 2011

Figure 6



The Arista 7148SX demonstrated latency ranging from 1.22 μ sec for 64-byte frames to 4.31 μ sec for 9216-byte frames (average latency: 1.9 μ sec).

The Cisco 5548P demonstrated latency ranging from 2.46 μ sec for 64-byte frames to 1.97 μ sec for 9216-byte frames (average latency: 2.12 μ sec).

The Juniper EX4500 demonstrated latency ranging from 4.08 μ sec for 128-byte frames to 16.27 μ sec for 9216-byte frames (average latency: 6.59 μ sec). See Figure 2 for the comparison of the latency results.

Layer 2 Multicast Performance

RFC 3918 One-to-Many Throughput

This test evaluated the ability of the DUTs to forward multicast traffic from one transmitting 10GbE port to the remaining 47 10GbE ports in a 1-to-47 “fan-out” configuration (for the IBM BNT RackSwitch G8264, all 64 ports were used in a 1-to-63 “fan out”), with 10 Gbps of Layer 2 unidirectional multicast traffic consisting of the standard Ethernet frame sizes between 64 bytes and 9216 bytes.

The RackSwitch G8264 and Arista 7148SX achieved 100% throughput for all frame sizes, while the Cisco 5548P achieved an average of 98.4%, and the Juniper EX4500 achieved 77.15% line rate for 64-bytes, but recovered to 100% for all other frame sizes.

RFC 3918 One-to-Many Latency

The one-to-many multicast latency tests were run using the same test port topology as the one-to-many multicast throughput tests, but while handling 95% of line rate multicast traffic.

The RackSwitch G8264 achieved 1.04 μ sec latency for 64-byte frames to 1.33 μ sec for 9216-byte frames (average 1.23 μ sec).

The Arista 7148SX achieved 1.15 μ sec latency for 64-byte frames to 1.29 μ sec for 9216-byte frames (average 1.26 μ sec).

The Cisco 5548P achieved 2.29 μ sec latency for 64-byte frames to 2.28 μ sec for 9216-byte frames (average 2.33 μ sec).

The Juniper EX4500 demonstrated much slower results with 3.82 μ sec for 64-byte frames and 14.88 μ sec for 9216-byte frames (average 5.91 μ sec).

Layer 3 Performance

RFC 2899 Full-Mesh Throughput

Similar to the Layer 2 performance tests, the switches were tested for Layer 3 performance.

To demonstrate the Layer 3 forwarding capabilities each port was configured in a separate subnet and a separate VLAN to require Layer 3 routing between ports.

The RackSwitch G8264 demonstrated consistent 100% throughput across packet sizes (64, 128, 256, 512, 1024, 1280, 1518 and 9216) across all available ports. See Figure 4. The BLADE switch utilized all 64 10GbE ports connected in full-mesh whereas the other switches were configured with 48 10GbE ports in full-mesh.

The Arista 7148SX was able to achieve an average 97.9% throughput across all packet sizes.

The Juniper EX4500 demonstrated 50.5% throughput on 64-byte packets, 88.9% on 128-byte packets, and 100% throughput on remaining packet sizes.

At the time of testing, the Cisco 5548P did not support Layer 3 forwarding and was excluded from these tests. According to Cisco, the Nexus 5548P is fully capable of Layer 3, and will be enabled in the near

future through a field upgradable daughter card.

RFC 2544 Port-Pair Throughput

The IBM BNT RackSwitch G8264 demonstrated consistent 100% throughput across packet sizes (64, 128, 256, 512, 1024, 1280, 1518 and 9216) across all available ports. See Figure 4. The BLADE switch utilized all 64 10GbE ports connected in port-pairs whereas the other switches were configured with 48 10GbE ports in port-pairs, again mapped from the outside in.

The Juniper EX4500 demonstrated 50.4% throughput on 64-byte packets, 88.9% on 128-byte packets, and 99.99% throughput on remaining packet sizes.

The Arista 7148SX was only able to achieve 33.3% throughput across all packet sizes. The poor performance of the 7148SX in the port-to-port testing stems from their multi-ASIC switch design, which creates a bottleneck between groups of ports, as demonstrated by the port-pairs test. Better performance may be achieved if the ports were aligned in sequence, but this would represent a far less likely network scenario.

ATSS Port-to-Port Latency

The DUTs were tested for Layer 3 cut-through latency at their highest sustained throughput with the port-pairs connected as per the Layer 2 testing. Each port was configured in a separate subnet and a separate VLAN to require Layer 3 routing between ports. The IBM BNT RackSwitch G8264 exhibited a latency ranging from 1.08 μ sec for 64-byte packets to an 1.41 μ sec for 9216-byte packets. See Figure 4.

The Arista 7148SX exhibited 1.29 μ sec latency for 64-byte packets to 4.53 μ sec for 9216-byte packets (average 2.00 μ sec).

The Juniper EX4500, being a store and forward switch, achieved from 4.96 μ sec




Figure 7

latency for 64-byte packets to 16.36 µsec for 9216-byte packets (average 6.47 µsec).

Microburst Tolerance

Microbursts are defined as sub-second periods of time when major bursts of network usage occur, causing utilization of network interfaces to become temporarily oversubscribed. This can possibly result in packet loss depending on the network device’s capacity to buffer the excess packets.

These conditions are typically caused when many devices transmit to a single device at the same time; the Microburst can even occur in a period of nanoseconds and result in dropped packets. Detection of such instances may also be nearly impossible to identify without the use of specialized network monitoring tools.

Microburst tests were run at the maximum achievable throughput rate of each packet size for each DUT using only three ports.

The IBM BNT RackSwitch G8264 demonstrated the capability to buffer 45,200 64-byte packets over 60 second intervals and 1,000 packets at a size of 9216-bytes.

The Arista 7148SX was capable of buffering 6,375 64-byte packets and 335 9216-byte packets.

The Cisco 5548P was able to buffer from 2,260 packets for 64-byte packets to 105 packets for 9216-byte packets.

The Juniper EX4500 was able to buffer only 445 64-byte packets and 25 9216-byte packets. See Figure 5 for more details.

The combination of a single ASIC design and a shared memory architecture make the IBM BNT switch a model for high performance packet replication as well as microburst absorption. Microbursts typically occur when multiple traffic flows are destined for a single egress port.

On the IBM BNT switch, every port had access to all the packet buffers in the switch.

This contrasts with multi-chip designs where the packet buffers are available only within a given chip and thus unavailable to a congested port on other chips.

Price/Performance Comparison

The price/performance ratio of the DUTs was calculated in terms of US dollars per Gbps (US\$/Gbps), using the maximum throughput measured by the engineers.

For this comparison, engineers configured the DUTs as similarly as possible. The Cisco Nexus 5548P was outfitted with the optional 16-port 10GbE uplink module in addition to the 32 base 10GbE ports.

The list price of the RackSwitch G8264 was obtained from the BLADE sales channel. The prices of the Cisco Nexus 5548P, Arista 7148SX, and Juniper EX4500 were obtained from an invoice of the DUTs, omitting the reseller discount to arrive at the list price.

The IBM BNT RackSwitch G8264 offered 57.5% better price/performance than the Cisco 5548P with a cost of \$46.87 per Gbps.



Systems Under Test Details

Vendor	Product	Optional Modules	Software/ Hardware Version	CDW Price (USD)	Stated Throughput (Gbps)	Price/Performance (USD/Gbps)
BLADE Network Technologies, an IBM Company	IBM BNT RackSwitch G8264	None	6.4.3.0	29,999.00	640	\$46.87
Arista Networks	Arista 7148SX	None	EOS 4.6.1	26,395.00	480	\$54.99
Cisco Systems, Inc.	Nexus 5548P Chassis	16 port 10GbE Module	NXOS 5.0(2)N1(1)	36,800.00	480	\$76.67
Juniper Networks	Juniper EX4500	2x 4 port 10GbE Uplink 1200W AC Power Supply	10.4R1.9	41,500.00	480	\$86.46

Note: All systems were priced equipped with at least 48 10GbE ports and redundant power supplies. MSRP listed as selling price on CDW website as of November 18, 2010.

Source: Tolly, February 2011

Table 1

The Arista 7148SX required no additional equipment and was priced at its list switch price.

The Juniper EX4500 solution required the purchase of two 4-port 10G SFP+ Uplink modules in addition to a 1200W AC power supply in order to match the base 48 10GbE ports.

None of the price/performance measurements include cost of the vendor-specific transceivers or cables. Instead, the four DUTs were tested using standard BLADE branded 10GbE SFP+ Direct Attach Cables, demonstrating the ability of Arista, Cisco, and Juniper to interoperate with the BLADE DACs.

Power Consumption

Engineers benchmarked the power consumption of the two DUTs at idle, and while handling 50% and 100% Layer-2 traffic across all available ports in full-mesh topology.

- The RackSwitch G8264 demonstrated 56.9% fewer watts per Gbps than the Cisco 5548P, 70.6% fewer watts/Gbps

than the Arista 7148SX, and 50.7% fewer Watts/Gbps than the Juniper EX4500.

- The RackSwitch G8264 consumed an average of 227 Watts at idle, 236 Watts at 50%, and 240 Watts at 100% traffic load.
- The Arista 7148SX consumed an average of 467 Watts at idle, 538 Watts at 50%, and 611 Watts at 100%.
- The Cisco 5548P consumed an average of 395 Watts at idle, 406 Watts at 50% load, and 417 Watts at 100% load.
- The Juniper EX4500 consumed an average of 361 Watts at idle, 363 Watts at 50%, and 364 Watts at 100%.
- In terms of annual cost to operate the equipment continuously on a 24 hour a day basis for 365 days per year, the RackSwitch G8264 costs an average of \$211.42 per year, versus an average of \$538.69 for the Arista 7148SX, \$367.56 for the Cisco Nexus 5548P, and \$321.39 for the Juniper EX4500. (All prices based on \$0.1007 per kilowatt hour, the 2010 National Commercial average.)
- The RackSwitch G8264 Watts per Gbps of throughput (lower is better) measured an average of 0.375 Watts per

Gbps, compared to an average of 0.868 Watts per Gbps on the Cisco 5548P, 1.27 on the Arista 7148SX, and 0.76 Watts per Gbps on the Juniper EX4500.

Test Environment

Devices Under Test

The IBM BNT RackSwitch G8264 is designed with 64 10GbE ports, for 640 Gbps full duplex lossless throughput. The IBM BNT RackSwitch G8264 supports Converged Enhanced Ethernet/Fibre Channel over Ethernet and network virtualization along with a robust Layer 2 and Layer 3 feature set and was running software version 6.4.3.0.

The Arista device under test consisted of 48 10GbE ports, and running software version EOS 4.6.1.

The Cisco device under test consisted of 32 10GbE ports and add-on 16-port 10GbE module, running software version NXOS 5.0 (2)N1(1).

The Juniper device under test consisted of 40 10GbE ports and two add-on 4-port 10GbE modules, and was running software version JunOS 10.4R1.9



Throughput and latency test traffic was generated using up to 64 10GbE ports on an Ixia Optixia XM12 chassis with IxAutomate Application version 6.90GA SP1 equipped with eight, 8-port 10GbE Line Cards, and four 1-port 40GbE Line Cards.

In all cases, the configuration of the DUTs had Spanning Tree and Flow Control disabled to demonstrate the pure forwarding capabilities of the switch and to avoid forwarding delays. Test durations consisted of 60 second intervals for each frame size under test.

Test Methodology

Layer 2 Performance Test

All performance benchmarks were completed using the RFC 2544, RFC 2889, and ATSS benchmarking suites for throughput and latency.

RFC 2544 & 2889 Throughput Tests

The throughput tests were run with both full-mesh and port-to-port traffic patterns to provide the most stressful and real-world data patterns for our benchmark. Port-Pair traffic patterns were used for latency tests to derive the most accurate latency values and pairs were matched from first to last port, i.e. 1 to 48, 2 to 47, 3 to 46, and so on. In order to test latency, engineers also used the lowest max throughput capability of the DUTs to ensure no packet loss would occur and skew the latency results.

ATSS Port-to-Port Latency Test

The IBM BNT RackSwitch G8264, Arista 7148SX, and Cisco 5548P operate in a Cut-Through packet forwarding mode, while Juniper uses a store and forward architecture. FIFO, also known as First In-First Out, represents the test tools setting for testing the switch in a Cut-Through mode by measuring the first point of a packet's ingress to the switch and the first point of the same packet's egress from the switch.

Using the "Calibrate Latency" feature present in IxAutomate, engineers were able to easily eliminate the latency induced by the presence of the test equipment, as well as the cabling used.

Layer 2 Multicast Forwarding

RFC3918 Multicast Throughput and Latency Tests

All multicast tests were performed using the Ixia Optixia XM12 chassis. The multicast performance tests were measured using the RFC3918 multicast benchmarking suite for throughput and latency.

The test consisted of one multicast transmitting port sending to all other ports using the same single IGMP group (225.0.0.1) membership for each port. Default configurations were used on all DUTs, with the exceptions detailed above.

Layer 3 Performance

To demonstrate the performance of the IBM BNT switch under the most stressful network conditions, full-mesh and port-to-port traffic patterns were used to derive the zero loss throughput of the switch. The Cisco 5548P did not support Layer 3 features and was hence not tested. The BLADE switch was tested using the standard RFC 2544 and 2889 throughput tests with Layer 3 bi-directional traffic.

To demonstrate the Layer 3 forwarding capabilities each port was configured in a separate subnet and a separate VLAN to require Layer 3 routing between ports.

RFC 2544 & 2889 Throughput Tests

Each DUT was tested using Layer 3 bi-directional traffic using the same methods outlined in the Layer 2 testing.

ATSS Port-to-Port Latency Test

Engineers measured the latency of the DUTs in a port-pair configuration using all available ports on each switch, much like the Layer 2 test.

Microburst Tolerance Test

Microburst tests were run at the maximum achievable throughput rate of each packet size for each DUT using only three ports. To derive the maximum throughput rate of each packet size Ixia's IxAutomate RFC 2889 full-mesh test was run for a 60 second duration per packet size.

The IxAutomate RFC 2889 Many-to-Many test was used with the "Peak-Load" option, which is designed to burst a total of (n) packets above the maximum throughput rate to force the DUT to buffer the excess packets. The test consisted of three ports on the XM12 chassis connected to each DUT using ports 1, 24, and 48. The number of packets sent during the "Peak-Load" burst was a configurable option that was varied manually at each packet size until the zero drop rate was determined for 64, 512, 1024, 1518, and 9216-byte packet sizes.

Power Consumption Tests

Engineers used a Watts up? PRO power meter to measure the power consumption of all DUTs in similarly-equipped configurations while idle and while handling 50% and 100% line-rate bidirectional, full-mesh traffic of 64, 512 and 1518-byte frames. At each load, traffic was maintained for 60 seconds and the average power consumption recorded. Tests were repeated three times and results averaged. Since all the switches were equipped with dual power supplies, engineers utilized a 1 to 2 breakout cable, which aggregated the power draw to the power meter.



About Tolly...



The Tolly Group companies have been delivering world-class IT services for 20 years. Tolly is a leading global provider of third-party validation services for vendors of IT products, components and services.

You can reach the company via E-mail at sales@tolly.com, or via telephone at +1 561.391.5610.

Visit Tolly on the Internet at: <http://www.tolly.com>

Test Equipment Summary

The Tolly Group gratefully acknowledges the providers of test equipment/software used in this project.

Vendor	Product	Web
Ixia	Optixia XM12 4x HSE 40GE QSFP-01 & 8x 10Ge LSM XMR8S line cards IxOS 5.70.6 EA SP2 IxAutomate 6.90 GA SP1	 http://www.ixiacom.com
Watts up?	Watts up? PRO	 http://wattsupmeters.com

Terms of Usage

This document is provided, free-of-charge, to help you understand whether a given product, technology or service merits additional investigation for your particular needs. Any decision to purchase a product must be based on your own assessment of suitability based on your needs. The document should never be used as a substitute for advice from a qualified IT or business professional. This evaluation was focused on illustrating specific features and/or performance of the product(s) and was conducted under controlled, laboratory conditions. Certain tests may have been tailored to reflect performance under ideal conditions; performance may vary under real-world conditions. Users should run tests based on their own real-world scenarios to validate performance for their own networks.

Reasonable efforts were made to ensure the accuracy of the data contained herein but errors and/or oversights can occur. The test/audit documented herein may also rely on various test tools the accuracy of which is beyond our control. Furthermore, the document relies on certain representations by the sponsor that are beyond our control to verify. Among these is that the software/hardware tested is production or production track and is, or will be, available in equivalent or better form to commercial customers. Accordingly, this document is provided "as is", and Tolly Enterprises, LLC (Tolly) gives no warranty, representation or undertaking, whether express or implied, and accepts no legal responsibility, whether direct or indirect, for the accuracy, completeness, usefulness or suitability of any information contained herein. By reviewing this document, you agree that your use of any information contained herein is at your own risk, and you accept all risks and responsibility for losses, damages, costs and other consequences resulting directly or indirectly from any information or material available on it. Tolly is not responsible for, and you agree to hold Tolly and its related affiliates harmless from any loss, harm, injury or damage resulting from or arising out of your use of or reliance on any of the information provided herein.

Tolly makes no claim as to whether any product or company described herein is suitable for investment. You should obtain your own independent professional advice, whether legal, accounting or otherwise, before proceeding with any investment or project related to any information, products or companies described herein. When foreign translations exist, the English document is considered authoritative. To assure accuracy, only use documents downloaded directly from Tolly.com.

No part of any document may be reproduced, in whole or in part, without the specific written permission of Tolly. All trademarks used in the document are owned by their respective owners. You agree not to use any trademark in or as the whole or part of your own trademarks in connection with any activities, products or services which are not ours, or in a manner which may be confusing, misleading or deceptive or in a manner that disparages us or our information, projects or developments.