Open Industry Network Performance & Power Test *for* Cloud Networks

Evaluating 10/40 GbE Switches

Fall 2011 Edition

A Report on the Extreme Networks BlackDiamond® X8 and Summit® X670V Data Center Switches

November, 2011

Table of Contents

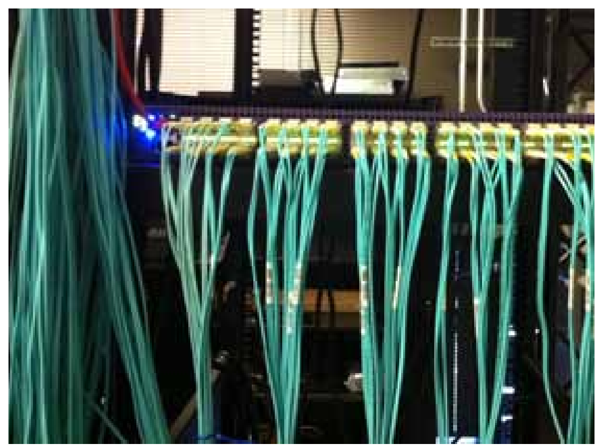

Introduction

To assist IT business leaders with the design and procurement of their private or public data center cloud fabric, the Lippis Report and Ixia have conducted an open industry evaluation of 10GbE and 40GbE data center switches. In this report, we review the test results of the Extreme Networks BlackDiamond® X8 and Summit® X670V Data Center Switches. These tests were conducted at the Ixia iSimCity Santa Clara, CA and Casabas, CA laboratories. The resources available for this test at Ixia's iSimCity are out of reach for nearly all corporate IT departments with test equipment on the order of $9.5M, devices under test on the order of $2M, plus costs associated with housing, power and cooling the lab. It's our hope that this report will remove performance, power consumption and latency concern from the purchase decision, allowing IT architects and IT business leaders to focus on other vendor selection criteria, such as post sales support, platform investment, vision, company financials, etc.

The Lippis test reports based on independent validation at Ixia's iSim City, communicates credibility, competence, openness and trust to potential buyers of 10GbE and 40GbE data center switching equipment as the tests are open to all suppliers and are fair, thanks to RFC and custom-based tests that are repeatable. The private/public data center cloud 10GbE and 40GbE fabric test was free for vendors to participate and open to all industry suppliers of 10GbE and 40GbE switching equipment, both modular and fixed configurations. The tests took place during the week of October 10 and October 31st 2011. Ixia supplied all test equipment needed to conduct the tests while Leviton provided optical SPF+ connectors and optical cabling and Siemon provided copper and fiber optic QSFP+ cables and transceivers for 40GbE connections.

Each 10GbE supplier was allocated lab time to run the test with the assistance of an Ixia engineer. Each switch vendor configured their equipment while Ixia engineers ran the test and logged the resulting data.

Extreme Networks BlackDiamond® X8 Core Switch

Extreme Networks launched its new entry into the enterprise data center and cloud computing markets with its new line of BlackDiamond® X series switches in May of 2011. This line of Top of Rack and Core switches pushes the envelope on performance, power consumption, density and data center network design.

Extreme Networks submitted their new BlackDiamond® X8 (BDX8) core switch into the Fall Lippis/Ixia industry test. The BDX8 is the highest density 10GbE and 40GbE core switch on the market. It's capable of supporting 768 10GbE and 192 40GbE ports in the smallest packaging footprint of one third of a rack or only 14.5RU. This high port density core switch in a small footprint offers new flexibility in building enterprise core networks and cloud data centers. For example, a single BDX8 could support 768 servers connected directly at 10GbE creating a non-blocking flat layer 2 network. But this is but one option in many designs afforded by the BDX8's attributes.

From an internal switching capacity point of view, the BDX8 offers 1.28 Tbps slot capacity and more than 20 Tbps switching fabric capacity; enough to support the 768 ports of 10 GbE or 192 ports of 40 GbE per chassis. In addition to raw port density and footprint, the BDX8 offers many software features. Some notable software features include automated configuration and provisioning, and XNV™(ExtremeXOS Network Virtualization), which provides automation of VM (Virtual Machine) tracking, reporting and migration of VPPs (Virtual Port Profile). These are but a few software features enabled through ExtremeXOS® network operating system. The BlackDiamond X8 is also key to Extreme Network's OpenFabric architecture that enables best-in-class choice of switching, compute, storage and virtualization solutions.

Extreme Networks BlackDiamond® X8 Core Switch Test Configuration*

	Hardware	Software Version	Port Density
Device under test	BSX8 www.extremenetworks.com	15.1	352
Test Equipment	Ixia XM12 High Performance Chassis	IxOS 6.10 EA IxNetwork 6.10 EA	
Ixia Line Cards	Xcellon Flex AP10G16S 16 port 10G module		
	Xcellon Flex Combo 10/40GE AP 16 port 10G / 4 port 40G		
	http://www.ixiacom.com/		
Cabling	10GbE Optical SFP+ connectors. Laser optimized duplex lc-lc 50 micron mm fiber, 850nm SPF+ transceivers www.leviton.com		
	Siemon QSFP+ Passive Copper Cable 40 GbE 3 meter copper QSFP-FA-010 www.siemon.com		
	Siemon Moray Low Power Active Optical Cable Assemblies Single Mode QSFP+ 40GbE optical cable QSFP30-03 7 meters www.siemon.com		

* Extreme BDx8 was configured in 256 10GbE ports and 24-40GbE ports or an equivalent 352-10GbE ports

Extreme Networks BlackDiamond® X8 Core Switch

The BDX8 breaks all of our previous records in core switch testing from performance, latency, power consumption, port density and packaging design. In fact, the Extreme Networks BDX8 was between 2 and 9 times faster at forwarding packets over a wide range of frame sizes than previous core switch measurements conducted during Lippis/Ixia iSimCity test of core switches from Alcatel-Lucent, Arista and Juniper. In other words, the BDX8 is the fastest core switch we have tested with the lowest latency measurements. The BDX8 is based upon the latest Broadcom merchant silicon chip set.

For the fall Lippis/Ixia test we populated the Extreme Networks BlackDiamond® X8 with 256 10GbE ports and 24 40GbE ports, thirty three percent of its capacity. This was the highest capacity switch tested during the entire series of Lippis/Ixia cloud network test at iSimCity to date.

We tested and measured the BDX8 in both cut through and store and forward modes in an effort to understand the difference these latency measurements offer. Further, latest merchant silicon forwards packets in store and forward for smaller packets while larger packets are forwarded in cut-through making this new generation of switches hybrid cut-through/store and forward devices. The Extreme Networks BlackDiamond X8 is unique in that its switching architecture forward packets in both cut-through and store-and-forward modes. As such, we measure the BlackDiamond X8 latency with both measurement methods. The purpose is not to compare store-and-forward vs cut through, but to demonstrate the low latency results obtained independent upon forwarding method. During store-and-forward testing, the latency of packet serialization delay (based on packet size) is removed from the reported latency number by test equipment. Therefore, to compare cutthrough to store-and-forward latency measurement, packet serialization delay needs to be added to store-and-forward latency number. For example, in a store-and-forward latency number of 800ns for a 1,518 byte size packet, the additional latency of 1240ns (serialization delay of a 1518 byte packet at 10Gbps) is required to be added to the store-and-forward measurement. This difference can be significant. Note that other potential device specific factors can impact latency too. This makes comparisons between two testing methodologies difficult.

Extreme Networks BDX8 RFC 2544 Layer 2 Latency Test in Cut-Through Mode

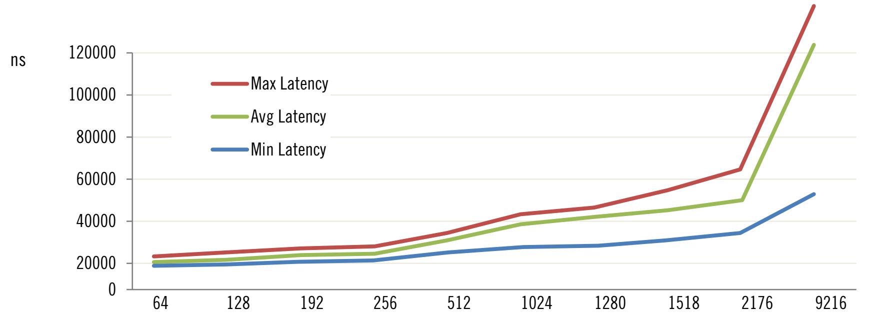

Max Latency	2510	2657	2862	2997	3542	4642	4877	5277	6402	18262
Avg Latency	2324.921	2442.854	2579.339	2673.136	3137.229	3919.9	4169.825	4396.996	5087.025	12755.718
Min Latency	2100	2162	2222	2277	2522	2897	2977	3062	3317	5982
Avg. delay Variation	8.314	5.403	9.179	5.386	7.636	7.146	5.632	9.314	4.986	9.09

The Extreme Networks BDX8 was tested across all 256 ports of 10GbE and 24 ports of 40GbE. Its average cut through latency ranged from a low of 2324 ns or 2.3 µs to a high of 12,755 ns or 12.7 µs at jumbo size 9216 Byte size frames for layer 2 traffic. Its average delay variation ranged between 5 and 9 ns, providing consistent latency across all packet sizes at full line rate; a welcome measurement for converged I/O implementations

Extreme Networks BlackDiamond® X8 Core Switch

The Extreme Networks BDX8 was tested across all 256 ports of 10GbE and 24 ports of 40GbE. Its average store and forward latency ranged from a low of 2273 ns or 2.2 μs to a high of 12,009 ns or 12 μs at jumbo size 9216 Byte size frames for layer 2 traffic. Its average delay variation ranged between 5 and 9 ns, providing consistent latency across all packet sizes at full line rate.

For layer 3 traffic, Extreme Networks BDX8's measured average cut through latency ranged from a low of 2,321 ns at 64Bytes to a high of 12,719 ns or 12.7 μs at jumbo size 9216 Byte size frames. Its average delay variation for layer 3 traffic ranged between 4 and 9 ns, providing consistent latency across all packet sizes at full line rate; again a welcome measurement for converged I/O implementations.

Extreme Networks BDX8
RFC 2544 Layer 2 Latency Test
in Store and Forward Mode

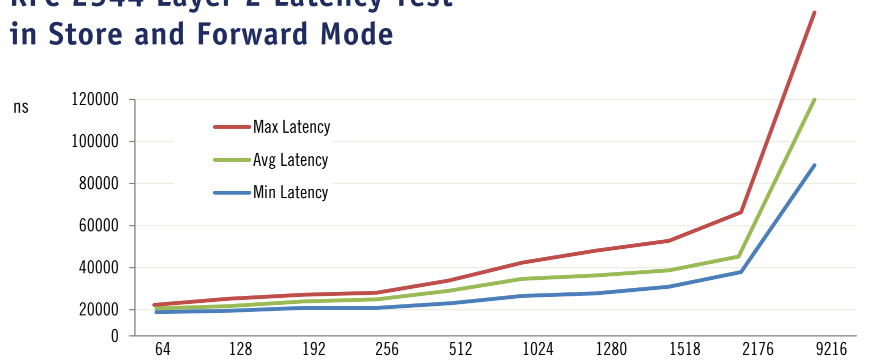

	64	128	192	256	512	1024	1280	1518	2176	9216
Max Latency	2482	2642	2797	2910	3442	4547	4990	5430	6762	21317
Avg Latency	2273.521	2351.750	2436.686	2493.575	2767.593	3331.511	3588.750	3830.261	4542.814	12095.05
Min Latency	2040	2100	2160	2180	2417	2780	2980	3140	3660	8960
Avg. delay Variation	8.314	5.403	9.179	5.386	7.636	7.146	5.632	9.314	4.986	9.079

Extreme Networks BDX8
RFC 2544 Layer 3 Latency Test
In Cut-Through Mode

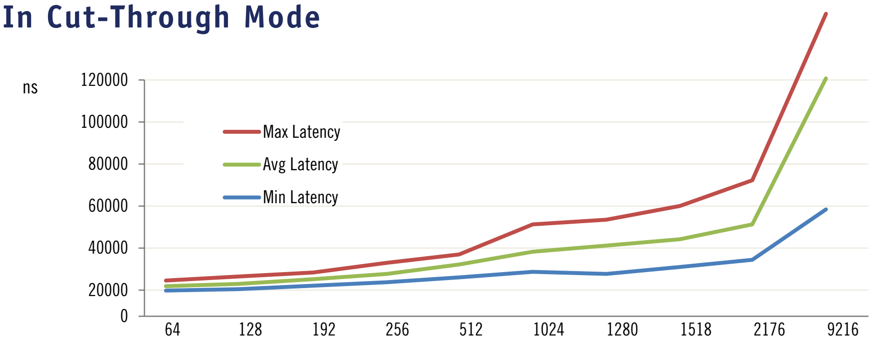

	64	128	192	256	512	1024	1280	1518	2176	9216
Max Latency	2550	2722	2902	3162	3770	5102	5430	6010	7290	21970
Avg Latency	2321.164	2440.482	2577.779	2668.411	3132.839	3898.596	4162.443	4406.857	5126.400	12719.646
Min Latency	2100	2162	2230	2302	2547	2870	2987	3070	3302	5997
Avg. delay Variation	9.982	5.454	9.2	5.411	7.593	7.164	5.529	9.321	4.982	9.118

Extreme Networks BlackDiamond® X8 Core Switch

For layer 3 traffic, Extreme Networks BDX8's measured average store and forward latency ranged from a low of 2,273 ns at 64Bytes to a high of 11,985 ns or 11.9 µs at jumbo size 9216 Byte size frames. Its average delay variation for layer 3 traffic ranged between 4 and 9 ns, providing consistent latency across all packet sizes at full line rate; again a welcome measurement for converged I/O implementations.

The above latency measures are the lowest we have measured for core switches by a large degree. The BDX8 forwards packets ten to six times faster than other core switches we've tested.

The Extreme Networks BDX8 demonstrated 100% throughput as a percentage of line rate across all 256-10GbE and 24-40GbE ports. In other words, not a single packet was dropped while the Extreme Networks BDX8 was presented with enough traffic to populate its 256-10GbE and 24-40GbE ports at line rate simultaneously for both L2 and L3 traffic flows measured in both cut-through and store and forward; a first in these Lippis/Ixia test.

Extreme Networks BDX8
RFC 2544 Layer 3 Latency Test
in Store and Forward Mode

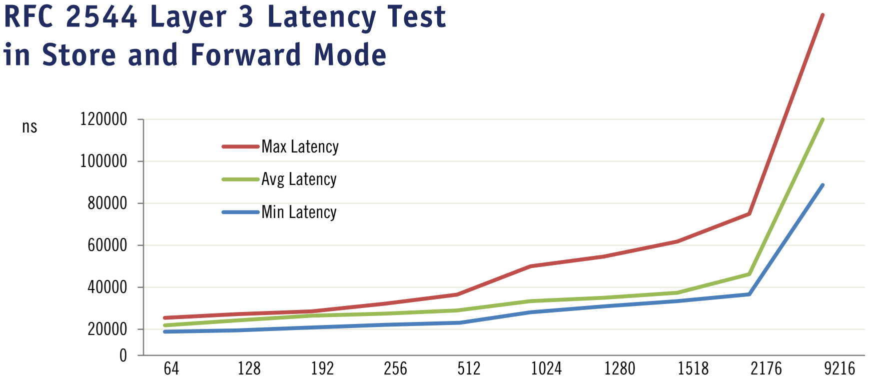

	64	128	192	256	512	1024	1280	1518	2176	9216
Max Latency	2542	2717	2937	3117	3657	5062	5537	6130	7677	25070
Avg Latency	2273.346	2353.996	2432.400	2492.879	2766.289	3324.279	3581.018	3826.468	4504.139	11985.946
Min Latency	2040	2100	2140	2220	2400	2780	3000	3160	3640	8960
Avg. delay Variation	9.982	5.454	9.2	5.411	7.593	7.164	5.529	9.321	4.982	9.118

Extreme Networks BDX8
RFC 2544 L2 & L3 Throughput Test

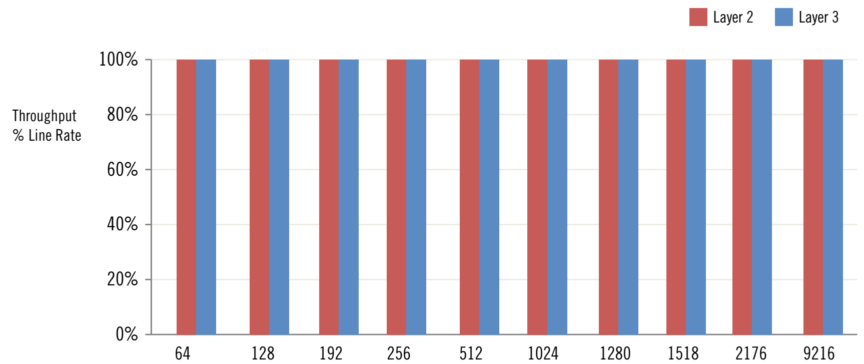

	64	128	192	256	512	1024	1280	1518	2176	9216
Layer 2	100	100	100	100	100	100	100	100	100	100
Layer 3	100	100	100	100	100	100	100	100	100	100

Extreme Networks BlackDiamond® X8 Core Switch

The Extreme Networks BDX8 demonstrated 100% of aggregated forwarding rate as percentage of line rate during congestion conditions. We tested congestion in both 10GbE and 40GbE configurations. A single 10GbE port was flooded at 150% of line rate. In addition, a single 40GbE port was flooded at 150% of line rate; another first in these Lippis/Ixia test.

The Extreme Networks BDX8 did not use HOL blocking which means that as the 10GbE and 40GbE ports on the BDX8 became congested, it did not impact the performance of other ports. Back pressure was detected. The BDX8 did send flow control frames to the Ixia test gear signaling it to slow down the rate of incoming traffic flow.

The Extreme Networks BDX8 demonstrated 100% aggregated throughput for IP multicast traffic measured in cut through mode with latencies ranging from a 2,382 ns at 64Byte size packet to 3,824 ns at 9216Byte size packets.

Extreme Networks BDX8 RFC 2889 Congestion Test

150% of Line Rate into a single 10GbE

Frame (Bytes)	Agg Forwarding Rate (% Line Rate)	Head of Line Blocking	Back Pressure	Agg Flow Control Frames	
				LAYER 2	LAYER 3
64	100	no	yes	8641332	9016074
128	100	no	yes	5460684	5459398
192	100	no	yes	6726918	6399598
256	100	no	yes	7294218	5798720
512	100	no	yes	8700844	5081322
1024	100	no	yes	5887130	5342640
1280	100	no	yes	5076130	5076140
1518	100	no	yes	4841474	4840594
2176	100	no	yes	5268752	5268552
9216	100	no	yes	5311742	5219964

150% of Line Rate into a single 40GbE

Frame (Bytes)	Agg Forwarding Rate (% Line Rate)	Head of Line Blocking	Back Pressure	Agg Flow Control Frames LAYER 2	Agg Flow Control Frames LAYER 3
64	100	no	yes	11305858	11303930
128	100	no	yes	9679780	9682464
192	100	no	yes	9470606	9466194
256	100	no	yes	29566584	10193602
512	100	no	yes	23892956	8909568
1024	100	no	yes	22794016	8915370
1280	100	no	yes	24439074	7944288
1518	100	no	yes	21749118	9093160
2176	100	no	yes	20897472	9101282
9216	100	no	yes	12938948	8977442

Extreme Networks BDX8
RFC 3918 IP Multicast Test Cut Through Mode

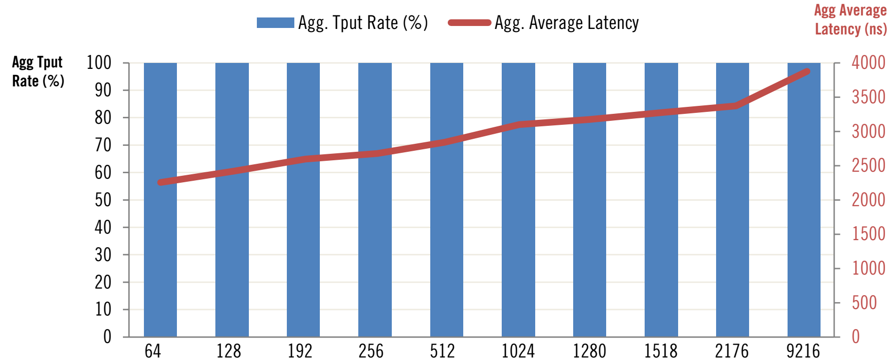

	64	128	192	256	512	1024	1280	1518	2176	9216
Agg Tput Rate (%)	100	100	100	100	100	100	100	100	100	100
Agg. Aver. Latency	2382.767	2448.613	2519.792	2578.477	2835.100	3117.849	3143.043	3161.283	3217.358	3824.054

Extreme Networks BlackDiamond® X8 Core Switch

The Extreme Networks BDX8 performed very well under cloud simulation conditions by delivering 100% aggregated throughput while processing a large combination of east-west and north-south traffic flows. Zero packet loss was observed as its latency stayed under 4.6 µs and 4.8 µs measured in cut through and store and forward modes respectively. This measurement too breaks all previous records as the BDX8 is between 2 and 10 times faster in forwarding cloud based protocols under load.

Extreme Networks BDX8 Cloud Simulation Test

Traffic Direction	Traffic Type	Cut-Through Avg Latency (ns)	Store and Forward Avg Latency (ns)
East-West	Database_to_Server	2574	3773
East-West	Server_to_Database	1413	1339
East-West	HTTP	4662	4887
East-West	iSCSI-Server_to_Storage	1945	2097
East-West	iSCSI-Storage_to_Server	3112	3135
North-South	Client_to_Server	3426	3368
North-South	Server_to_Client	3576	3808

The Extreme Networks BDX8 represents a new breed of cloud network core switches with power efficiency being a core value. Its WattsATIS/port is 8.1 and TEER value is 117. To put power measurement into perspective, other core switches are between 30% and 2.6 times less power efficient. In other words, the BDX8 is the most power efficient core switch tested to date in the Lippis/Ixia iSimCity lab.

Extreme Networks BDX8 Power Consumption Test

$\text{Watts}_{\text{ATIS}}$/10GbE port	8.1
3-Year Cost/$\text{Watts}_{\text{ATIS}}$/10GbE	$29.61
Total power cost/3-Year	$10,424.05
3 yr energy cost as a % of list price	1.79%
TEER Value	117
Cooling	Front to Back, Reversible

Note that the total number of BDX8 10GbE ports to calculate Watts/10GbE were 352, as 24-40Gbs ports or 96 10GbE ports were added to the 256-10GbE ports for Watts/ATIS/port calculation. The 40GbE ports are essentially four 10GbE ports from a power consumption point of view.

While these are the lowest Watts/10GbE port and highest TEER values observed for core switches, the Extreme Networks BDX8's actual Watts/10GbE port is actually lower; we estimate approximately 5 Watts/10GbE port when fully populated with 756 10GbE or 192 40GbE ports. During the Lippis/Ixia test, the BDX8 was only populated to a third of its port capacity but equipped with power supplies, fans, management and switch fabric modules for full port density population. Therefore, when this full power capacity is divided across a fully populated BDX8, its WattsATIS per 10GbE Port will be lower than the measurement observed.

The Extreme Networks BDX8 power cost per 10GbE is calculated at $9.87 per year. The three year cost to power the Extreme Networks BDX8 is estimated at $10,424.05 and represents less than 1.7% of its list price. Keeping with data center best practices, its cooling fans flow air front to back.

Discussion:

The Extreme Networks BDX8 seeks to offer high performance, low latency, high port density of 10 and 40GbE and low power consumption for private and public cloud networks. This architecture proved its value as it delivered the lowest latency measurements to date for core switches while populated with the equivalent of 352 10GbE ports running traffic at line rate. Not a single packet was dropped offering 100% throughput a line rate for the equivalent of 352 10GbE ports. Even at 8 Watts/10GbE port, the Extreme Networks BDX8 consumes the least amount of power of core switches measured to date at the Lippis/Ixia industry test at iSimCity. Observations of the BDX8's IP Multicast and congestion test performance were the best measured to date in terms of latency and throughput under congestion conditions in both 10GbE and 40GbE scenarios.

The Extreme Networks BDX8 was found to have low power consumption, front-to-back cooling, front-accessible components and an ultra compact form factor. The Extreme Networks BDX8 is designed to meet the requirements for private and public cloud networks. Based upon these Lippis/Ixia test, it achieves its design goals.

Extreme Networks Summit® X670V Top-of-Rack Switch

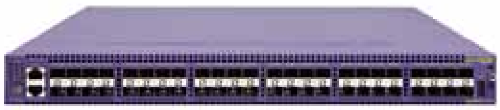

Extreme Networks launched its new entry into the enterprise data center and cloud computing markets with its new line of Summit® X series of switches in May of 2011. This line of Top of Rack switches pushes the envelope on performance, power consumption, density and data center network design.

The Summit® X670 series are purpose-built Top-of-Rack (ToR) switches designed to support emerging 10 GbE enabled servers in enterprise and cloud data centers. The Summit X670 is ready for 40GbE uplinks thanks to its optional 40 GbE support. In addition the Summit X670 provides for both 1 and 10GbE physical server connections to ease transition to high-performance server connections while providing unique features for virtualized data center environments. The ExtremeXOS® operating system, with its high availability attributes provides simplicity and ease of operation through the use of one OS everywhere in the network.

Extreme Networks submitted their new Summit® X670V ToR switch into the Fall Lippis/Ixia industry test. The X670V is a member of the Summit® X family of ToR switches. The X670V is capable of supporting 48-10GbE and 4-40GbE ports in a small 1RU footprint. The X670V in combination with Extreme Networks BlackDiamond® X8 offer new flexibility in building enterprise and cloud spec data center networks with the X670V providing server connectivity at 1 and 10GbE while connecting into the BDX8 at either 10 or 40GbE under a single OS image in both devices.

Extreme Networks Summit® X670V Test Configuration*

	Hardware	Software Version	Port Density
Device under test	X670V www.extremenetworks.com	12.6	64
Test Equipment	Ixia XG12 High Performance Chassis	IxOS 6.10 EA SP2 IxNetwork 6.10 EA	
Ixia Line Cards	Xcellon Flex AP10G16S 16 port 10G module		
	Xcellon Flex Combo 10/40GE AP 16 port 10G / 4 port 40G		
	Xcellon Flex 4x40GEQSFP+ 4 port 40G		
	http://www.ixiacom.com/		
Cabling	10GbE Optical SFP+ connectors. Laser optimized duplex lc-lc 50 micron mm fiber, 850nm SPF+ transceivers www.leviton.com		
	Siemon QSFP+ Passive Copper Cable 40 GbE 3 meter copper QSFP-FA-010 www.siemon.com		
	Siemon Moray Low Power Active Optical Cable Assemblies Single Mode QSFP+ 40GbE optical cable QSFP30-03 7 meters www.siemon.com		

* Extreme X670V was configured in 48 10GbE ports and 4-40GbE ports or an equivalent 64-10GbE ports

Extreme Networks Summit® X670V Top-of-Rack Switch

For the fall Lippis/Ixia test we populated and tested the Extreme Networks Summit® X670V with 48-10GbE ports and 4-40GbE ports; its full capacity. Its average latency ranged from a low of 900 ns to a high of 1,415 ns for layer 2 traffic. Its average delay variation ranged between 7 and 10 ns, providing consistent latency across all packet sizes at full line rate.

For layer 3 traffic, the Extreme Networks X670V's measured average latency ranged from a low of 907 ns at 64Bytes to a high of 1,375 ns. Its average delay variation for layer 3 traffic ranged between 5 and 9 ns, providing consistent latency across all packet sizes at full line rate.

Extreme Networks Summit® X670V RFC 2544 Layer 2 Latency Test in Cut-Through Mode

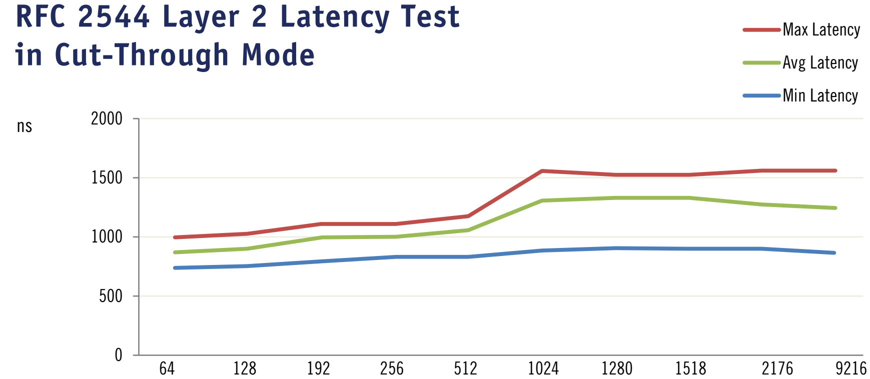

	64	128	192	256	512	1024	1280	1518	2176	9216
Max Latency	1020	1040	1200	1200	1360	1580	1560	1560	1580	1580
Avg Latency	900.192.	915.519	1056.519	1067.115	1208.731	1412.808	1415.962	1415.327	1411.923	1410.019
Min Latency	737	742	762	782	782	802	817	817	817	802
Avg. delay Variation	9.346	7.962	10.25	6.942	7.865	9.712	7.827	10.192	7.962	10.731

Extreme Networks Summit® X670V RFC 2544 Layer 3 Latency Test in Cut-Through Mode

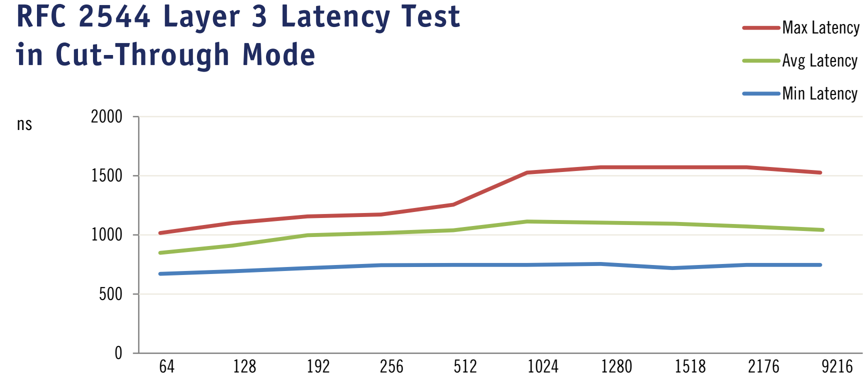

	64	128	192	256	512	1024	1280	1518	2176	9216
Max Latency	1040	1060	1160	1180	1320	1520	1540	1540	1540	1520
Avg Latency	907.923	943.135	1026.615	1029.731	1176.923	1375.019	1373.865	1374.077	1372.712	1370.327
Min Latency	762	747	790	770	810	810	847	790	810	810
Avg. delay Variation	8.135	5.423	8.942	5.558	7.019	6.885	4.558	9.346	5.058	9.712

Extreme Networks Summit® X670V Top-of-Rack Switch

The Extreme Networks X670V demonstrated 100% throughput as a percentage of line rate across all 48-10GbE and 4-40GbE ports. In other words, not a single packet was dropped while the Extreme Networks X670V was presented with enough traffic to populate its 48-10GbE and 4-40GbE ports at line rate simultaneously for both L2 and L3 traffic flows.

Two congestion test were conducted using the same methodology. A 10GbE and 40GbE congestion test stressed the X670V ToR switch's congestion management attributes for both 10GbE and 40GbE. The Extreme Networks X670V demonstrated 100% of aggregated forwarding rate as percentage of line rate during congestion conditions for both the 10GbE and 40GbE test. A single 10GbE port was flooded at 150% of line rate. In addition, a single 40GbE port was flooded at 150% of line rate.

The Extreme Networks X670V did not use HOL blocking which means that as the 10GbE and 40GbE ports on the X670V became congested, it did not impact the performance of other ports. Back pressure was detected. The X670V did send flow control frames to the Ixia test gear signaling it to slow down the rate of incoming traffic flow, which is common in ToR switches.

Extreme Networks Summit® X670V RFC 2544 L2 & L3 Throughput Test

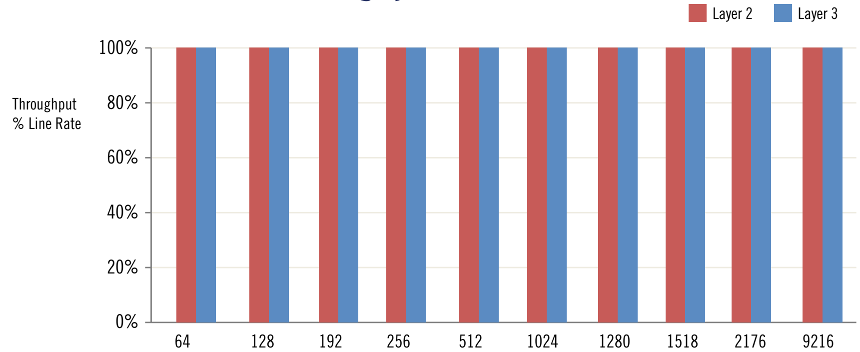

	64	128	192	256	512	1024	1280	1518	2176	9216
Layer 2	100	100	100	100	100	100	100	100	100	100
Layer 3	100	100	100	100	100	100	100	100	100	100

Extreme Networks Summit® X670V Congestion Test

150% of Line Rate into a single 10GbE

Frame (Bytes)	Agg Forwarding Rate (% Line Rate)	Head of Line Blocking	Back Pressure	Agg Flow Control Frames	
				LAYER 2	LAYER 3
64	100	no	yes	6976188	6641074
128	100	no	yes	4921806	5720526
192	100	no	yes	6453360	6751626
256	100	no	yes	4438204	6167986
512	100	no	yes	3033068	5425000
1024	100	no	yes	3125294	5668574
1280	100	no	yes	17307786	5540616
1518	100	no	yes	13933834	5320560
2176	100	no	yes	13393656	5339310
9216	100	no	yes	9994308	5042648

150% of Line Rate into a single 40GbE

Frame (Bytes)	Agg Forwarding Rate (% Line Rate)	Head of Line Blocking	Back Pressure	Agg Flow Control Frames LAYER 2	Agg Flow Control Frames LAYER 3
64	100	no	yes	10265682	10273162
128	100	no	yes	7951820	9517872
192	100	no	yes	9938718	10655066
256	100	no	yes	28243060	9508906
512	100	no	yes	26799154	9534112
1024	100	no	yes	31017876	7900572
1280	100	no	yes	25762198	9482604
1518	100	no	yes	28271714	9345090
2176	100	no	yes	27102102	9294636
9216	100	no	yes	22865370	9521042

Extreme Networks Summit® X670V Top-of-Rack Switch

The Extreme Networks X670V demonstrated 100% aggregated throughput for IP multicast traffic with latencies ranging from a 955 ns at 64Byte size packet to 1009 ns at 9216Byte size packets.

The Extreme Networks X670V performed well under cloud simulation conditions by delivering 100% aggregated throughput while processing a large combination of east-west and north-south traffic flows. Zero packet loss was observed as its latency stayed under 3.1 µs measured in cut through mode.

The Extreme Networks X670V represents a new breed of cloud network ToR switches with power efficiency being a core value. Its WattsATIS/port is 3.2 and TEER value is 295. Note that the total number of X670V 10GbE ports to calculate Watts/10GbE were 64, as 4-40Gbs ports or 16 10GbE ports were added to the 48-10GbE ports for Watts/ATIS/port calculation. The 40GbE ports are essentially four 10GbE ports from a power consumption point of view.

The Extreme Networks X670V power cost per 10GbE is calculated at $11.74 per 3-year period. The three-year cost to power the Extreme Networks X670V is estimated at $751.09 and represents less than 2.2% of its list price. Keeping with data center best practices, its cooling fans flow air front to back.

Extreme Networks Summit® X670V RFC 3918 IP Multicast Test Cut Through Mode

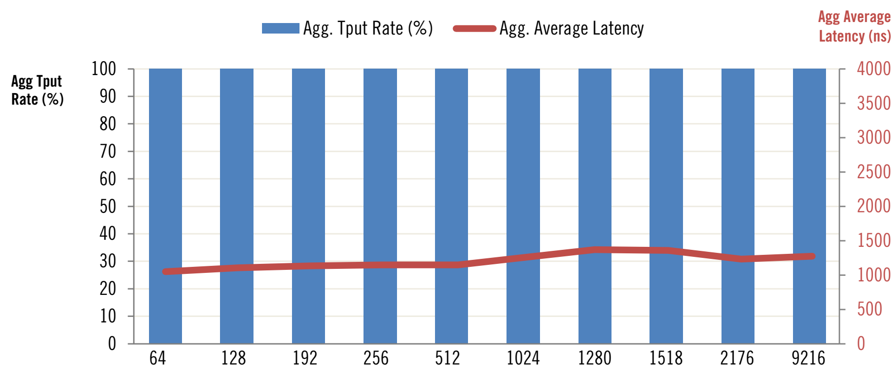

Agg Tput Rate (%)	100	100	100	100	100	100	100	100	100	100
Agg. Aver. Latency	955.824	968.431	987.51	985.98	998.157	1011	1012.529	1010.157	1009.824	1009.863

Extreme Networks Summit® X670V Cloud Simulation Test

Traffic Direction	Traffic Type	Cut-through Avg Latency (ns)
East-West	Database_to_Server	2568
East-West	Server_to_Database	1354
East-West	HTTP	2733
East-West	iSCSI-Server_to_Storage	1911
East-West	iSCSI-Storage_to_Server	3110
North-South	Client_to_Server	1206
North-South	Server_to_Client	1001

Extreme Networks Summit® X670V Power Consumption Test

$\text{Watts}_{\text{ATIS}}$/10GbE port	3.21
3-Year Cost/$\text{Watts}_{\text{ATIS}}$/10GbE	$11.74
Total power cost/3-Year	$751.09
3 yr energy cost as a % of list price	2.21%
TEER Value	295
Cooling	Front to Back, Reversible

Discussion:

The Extreme Networks X670V seeks to offer a range of server and ToR connectivity options to IT architects as they design their data centers and cloud network facilities. The X670V is competitive with other ToR switches from a price, performance, latency and power consumption point of view. The X670V when combined with the X8 offers a powerful two tier cloud network architecture as they bring together the performance advantages observed during the Lippis/Ixia test and software features found is ExtremeXOS® network operating system.

The Extreme Networks X670V was found to have low power consumption, front-to-back cooling, front-accessible components and a 1RU compact form factor. The Extreme Networks X670V is designed to meet the requirements for private and public cloud networks. Based upon these Lippis/Ixia test, it achieves its design goals.

The Lippis Report Test Methodology

To test products, each supplier brought its engineers to configure its equipment for test. An Ixia test engineer was available to assist each supplier through test methodologies and review test data. After testing was concluded, each supplier's engineer signed off on the resulting test data. We call the following set of testing conducted "The Lippis Test." The test methodologies included:

Throughput Performance: Throughput, packet loss and delay for layer-2 (L2) unicast, layer-3 (L3) unicast and layer-3 multicast traffic was measured for packet sizes of 64, 128, 256, 512, 1024, 1280, 1518, 2176, 9216 bytes. In addition, a special cloud computing simulation throughput test consisting of a mix of north-south plus east-west traffic was conducted. Ixia's IxNetwork RFC 2544 Throughput/Latency quick test was used to perform all but the multicast tests. Ixia's IxAutomate RFC 3918 Throughput No Drop Rate test was used for the multicast test.

Latency: Latency was measured for all the above packet sizes plus the special mix of north-south and east-west traffic blend. Two latency tests were conducted: 1) latency was measured as packets flow between two ports on different modules for modular switches, and 2) between far away ports (port pairing) for ToR switches to demonstrate latency consistency across the forwarding engine chip. Latency test port configuration was via port pairing across the entire device versus side-by-side. This meant that a switch with N ports, port 1 was paired with port (N/2)+1, port 2 with port (N/2)+2, etc. Ixia's IxNetwork RFC 2544 Throughput / Latency quick test was used for validation.

Jitter: Jitter statistics was measured during the above throughput and latency test using Ixia's IxNetwork RFC 2544 Throughput/Latency quick test.

Congestion Control Test: Ixia's IxNetwork RFC 2889 Congestion test was used to test both L2 and L3 packets. The objective of the Congestion Control Test is to determine how a Device Under Test (DUT) handles congestion. Does the device implement congestion control and does congestion on one port affect an uncongested port? This procedure determines if Head-of-Line (HOL) blocking and/or if back pressure are present. If there is frame loss at the uncongested port, HOL blocking is present. Therefore, the DUT cannot forward the amount of traffic to the congested port, and as a result, it is also losing frames destined to the uncongested port. If there is no frame loss on the congested port and the port receives more packets than the maximum offered load of 100%, then Back Pressure is present.

Video feature: **Click to view a discussion on the Lippis Report Test Methodology**

RFC 2544 Throughput/Latency Test

Test Objective: This test determines the processing overhead of the DUT required to forward frames and the maximum rate of receiving and forwarding frames without frame loss.

Test Methodology: The test starts by sending frames at a specified rate, usually the maximum theoretical rate of the port while frame loss is monitored. Frames are sent from and received at all ports on the DUT, and the transmission and reception rates are recorded. A binary, step or combo search algorithm is used to identify the maximum rate at which no frame loss is experienced.

To determine latency, frames are transmitted for a fixed duration. Frames are tagged once in each second and during half of the transmission duration, then tagged frames are transmitted. The receiving and transmitting timestamp on the tagged frames are compared. The difference between the two timestamps is the latency. The test uses a one-to-one traffic mapping. For store and forward DUT switches latency is defined in RFC 1242 as the time interval starting when the last bit of the input frame reaches the input port and ending when the first bit of the output frame is seen on the output port. Thus latency is not dependent on link speed only, but processing time too.

Results: This test captures the following data: total number of frames transmitted from all ports, total number of frames received on all ports, percentage of lost frames for each frame size plus latency, jitter, sequence errors and data integrity error.

The following graphic depicts the RFC 2554 throughput performance and latency test conducted at the iSimCity lab for each product.

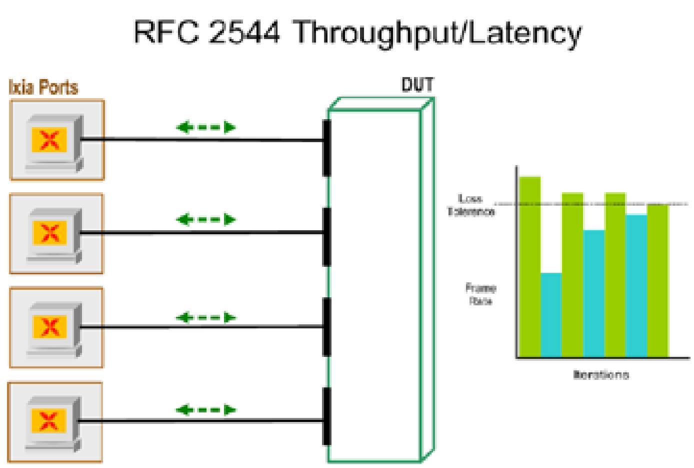

RFC 2889 Congestion Control Test

Test Objective: The objective of the Congestion Control Test is to determine how a DUT handles congestion. Does the device implement congestion control and does congestion on one port affect an uncongested port? This procedure determines if HOL blocking and/or if back pressure are present. If there is frame loss at the uncongested port, HOL blocking is present. If the DUT cannot forward the amount of traffic to the congested port, and as a result, it is also losing frames destined to the uncongested port. If there is no frame loss on the congested port and the port receives more packets than the maximum offered load of 100%, then Back Pressure is present.

Test Methodology: If the ports are set to half duplex, collisions should be detected on the transmitting interfaces. If the ports are set to full duplex and flow control is enabled, flow control frames should be detected. This test consists of a multiple of four ports with the same MOL (Maximum Offered Load). The custom port group mapping is formed of two ports, A and B, transmitting to a third port C (the congested interface), while port A also transmits to port D (uncongested interface).

Test Results: This test captures the following data: intended load, offered load, number of transmitted frames, number of received frames, frame loss, number of collisions and number of flow control frames obtained for each frame size of each trial are captured and calculated.

The following graphic depicts the RFC 2889 Congestion Control test as conducted at the iSimCity lab for each product.

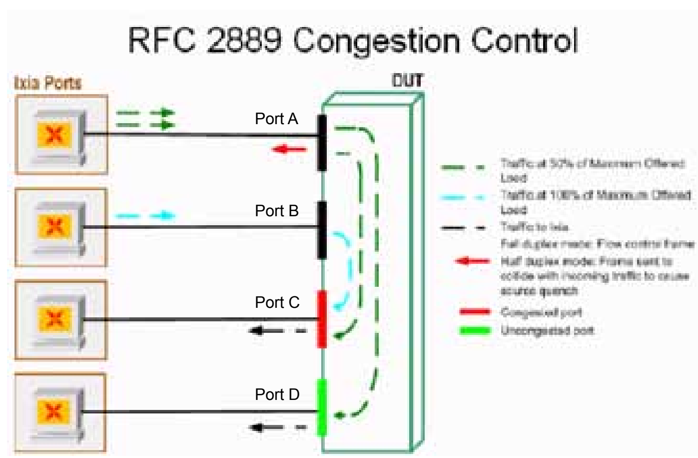

RFC 3918 IP Multicast Throughput No Drop Rate Test

Test Objective: This test determines the maximum throughput the DUT can support while receiving and transmitting multicast traffic. The input includes protocol parameters (IGMP, PIM), receiver parameters (group addressing), source parameters (emulated PIM routers), frame sizes, initial line rate and search type.

Test Methodology: This test calculates the maximum DUT throughput for IP Multicast traffic using either a binary or a linear search, and to collect Latency and Data

Integrity statistics. The test is patterned after the ATSS Throughput test; however this test uses multicast traffic. A one-to-many traffic mapping is used, with a minimum of two ports required.

If choosing OSPF or ISIS as IGP protocol routing, the transmit port first establishes an IGP routing protocol session and PIM session with the DUT. IGMP joins are then established for each group, on each receive port. Once protocol sessions are established, traffic begins to transmit into the DUT and a binary or linear search for maximum throughput begins.

If choosing "none" as IGP protocol routing, the transmit port does not emulate routers and does not export routes to virtual sources. The source addresses are the IP addresses configured on the Tx ports in data frame. Once the routes are configured, traffic begins to transmit into the DUT and a binary or linear search for maximum throughput begins.

Test Results: This test captures the following data: maximum throughput per port, frame loss per multicast group, minimum/maximum/average latency per multicast group and data errors per port.The following graphic depicts the RFC 3918 IP Multicast Throughput No Drop Rate test as conducted at the iSimCity lab for each product.

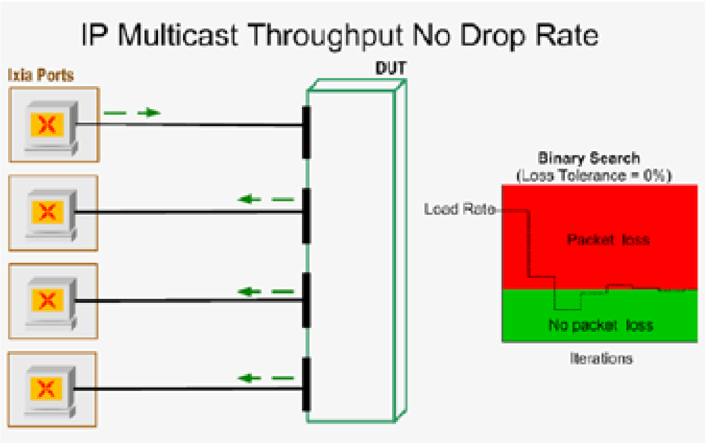

Power Consumption Test

Port Power Consumption: Ixia's IxGreen within the IxAutomate test suite was used to test power consumption at the port level under various loads or line rates.

Test Objective: This test determines the Energy Consumption Ratio (ECR), the ATIS (Alliance for Telecommunications Industry Solutions) Telecommunications Energy Efficiency Ratio (TEER) during a L2/L3 forwarding performance. TEER is a measure of network-element efficiency quantifying a network component's ratio of "work performed" to energy consumed.

Test Methodology: This test performs a calibration test to determine the no loss throughput of the DUT. Once the maximum throughput is determined, the test runs in automatic or manual mode to determine the L2/L3 forwarding performance while concurrently making power, current and voltage readings from the power device. Upon completion of the test, the data plane performance and Green (ECR and TEER) measurements are calculated. Engineers followed the methodology prescribed by two ATIS standards documents:

ATIS-0600015.03.2009: Energy Efficiency for Telecommunication Equipment: Methodology for Measuring and Reporting for Router and Ethernet Switch Products, and

ATIS-0600015.2009: Energy Efficiency for Telecommunication Equipment: Methodology for Measuring and Reporting - General Requirements

The power consumption of each product was measured at various load points: idle 0%, 30% and 100%. The final power consumption was reported as a weighted average calculated using the formula:

WATIS = 0.1*(Power draw at 0% load) + 0.8*(Power draw at 30% load) + 0.1*(Power draw at 100% load).

All measurements were taken over a period of 60 seconds at each load level, and repeated three times to ensure result repeatability. The final WATIS results were reported as a weighted average divided by the total number of ports per switch to derive at a WATTS per port measured per ATIS methodology and labeled here as $\text{WATTS}_{\text{ATIS.}}$

Test Results: The L2/L3 performance results include a measurement of WATIS and the DUT TEER value. Note that a larger TEER value is better as it represents more work done at less energy consumption. In the graphics throughout this report, we use $\text{WATTS}_{\text{ATIS}}$ to identify ATIS power consumption measurement on a per port basis.

With the $\text{WATTS}_{\text{ATIS}}$ we calculate a three-year energy cost based upon the following formula.

$\text{Cost/Watts}_{\text{ATIS}}\text{/3-Year} =$
$(\text{WATTS}_{\text{ATIS}}/1000)*(3*365*24)*(0.1046)*(1.33)$,
where $\text{WATTS}_{\text{ATIS}}$ = ATIS weighted average power in Watts

3*365*24 = 3 years @ 365 days/yr @ 24 hrs/day

0.1046 = U.S. average retail cost (in US$) of commercial grade power as of June 2010 as per Dept. of Energy Electric Power Monthly

(http://www.eia.doe.gov/cneaf/electricity/epm/table5_6_a.html)
1.33 = Factor to account for power costs plus cooling costs @ 33% of power costs.

The following graphic depicts the per port power consumption test as conducted at the iSimCity lab for each product.

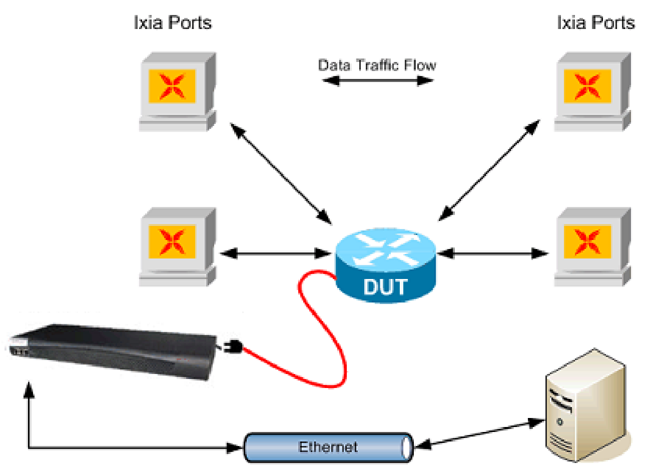

Public Cloud Simulation Test

Test Objective: This test determines the traffic delivery performance of the DUT in forwarding a variety of north-south and east-west traffic in cloud computing applications. The input parameters include traffic types, traffic rate, frame sizes, offered traffic behavior and traffic mesh.

Test Methodology: This test measures the throughput, latency, jitter and loss on a per application traffic type basis across M sets of 8-port topologies. M is an integer and is proportional to the number of ports the DUT is populated with. This test includes a mix of north-south traffic and east-west traffic, and each traffic type is configured for the following parameters: frame rate, frame size distribution, offered traffic load and traffic mesh. The following traffic types are used: web (HTTP), database-server, server-database, iSCSI storage-server, iSCSI server-storage, client-server plus server-client. The north-south client-server traffic simulates Internet browsing, the database traffic simulates server-server lookup and data retrieval, while the storage traffic simulates IP-based storage requests and retrieval. When all traffic is transmitted, the throughput, latency, jitter and loss performance are measured on a per traffic type basis.

Test Results: This test captures the following data: maximum throughput per traffic type, frame loss per traffic type, minimum/maximum/average latency per traffic type, minimum/maximum/average jitter per traffic type, data integrity errors per port and CRC errors per port. For this report we show average latency on a per traffic basis at zero frame loss.

The following graphic depicts the Cloud Simulation test as conducted at the iSimCity lab for each product.

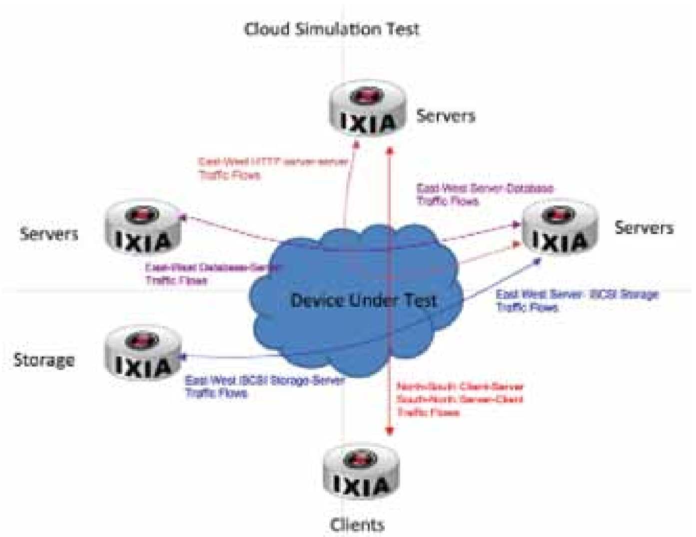

40GbE Testing For the test plan above, 24-40GbE ports were available for those DUT that support 40GbE uplinks or modules. During this Lippis/Ixia test 40GbE testing included latency, throughput congestion, IP Multicast, cloud simulation and power consumption test. ToR switches with 4-40GbE uplinks are supported in this test.

Terms of Use

This document is provided 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 vendors 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 Lippis Enterprises, Inc. (Lippis), 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. Lippis Enterprises, Inc. is not responsible for, and you agree to hold Lippis Enterprises, Inc. 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.

Lippis Enterprises, Inc. 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 www.lippisreport.com .

No part of any document may be reproduced, in whole or in part, without the specific written permission of Lippis Enterprises, Inc. 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.

About Nick Lippis

Nicholas J. Lippis III is a world-renowned authority on advanced IP networks, communications and their benefits to business objectives. He is the publisher of the Lippis Report, a resource for network and IT business decision makers to which over 35,000 executive IT business leaders subscribe. Its Lippis Report podcasts have been downloaded over 160,000 times; i-Tunes reports that listeners also download the Wall Street Journal's Money Matters, Business Week's Climbing the Ladder, The Economist and The Harvard Business Review's IdeaCast. Mr. Lippis is currently working with clients to design their private and public virtualized data center cloud computing network architectures to reap maximum business value and outcome.

He has advised numerous Global 2000 firms on network architecture, design, implementation, vendor selection and budgeting, with clients including Barclays Bank, Eastman Kodak Company, Federal Deposit Insurance Corporation (FDIC), Hughes Aerospace, Liberty Mutual, Schering-Plough, Camp Dresser McKee, the state of Alaska, Microsoft, Kaiser Permanente, Sprint, Worldcom, Cigitel, Cisco Systems, Hewlett Packet, IBM, Avaya and many others. He works exclusively with CIOs and their direct reports. Mr. Lippis possesses a unique perspective of market forces and trends occurring within the computer networking industry derived from his experience with both supply and demand side clients.

Mr. Lippis received the prestigious Boston University College of Engineering Alumni award for advancing the profession. He has been named one of the top 40 most powerful and influential people in the networking industry by Network World. TechTarget an industry on-line publication has named him a network design guru while Network Computing Magazine has called him a star IT guru.

Mr. Lippis founded Strategic Networks Consulting, Inc., a well-respected and influential computer networking industry-consulting concern, which was purchased by Softbank/Ziff-Davis in 1996. He is a frequent keynote speaker at industry events and is widely quoted in the business and industry press. He serves on the Dean of Boston University's College of Engineering Board of Advisors as well as many start-up venture firm's advisory boards. He delivered the commencement speech to Boston University College of Engineering graduates in 2007. Mr. Lippis received his Bachelor of Science in Electrical Engineering and his Master of Science in Systems Engineering from Boston University. His Masters' thesis work included selected technical courses and advisors from Massachusetts Institute of Technology on optical communications and computing.