

Mellanox SX1016 & SX1036 10/40GbE Switches

Performance and Power Consumption Evaluation

EXECUTIVE SUMMARY

As data center demands continually increase, network architects seek out the best performance combined with operational efficiency and scalability. Mellanox is focused on meeting these needs by developing L2/L3 Ethernet switches that offer high throughput, extremely low latency and ultra low power consumption.


Mellanox commissioned Tolly to evaluate SwitchX[®]-based SX1016 and SX1036 Ethernet switches benchmarking throughput, latency, multicast performance and power consumption. The SX1036 delivers wire-speed throughput running 36 ports of 40GbE, as well as when running 48 ports of 10GbE along with 12 ports of 40GbE with cut through latency less than 223ns for L2 40GbE and 269ns for L2 10GbE. See "The Bottom Line" and Figure 1.

Mellanox notes that the above configuration represents the new Mellanox switch, SX1024, 2012 Best Of Interop Finalist, featuring 48 SFP+ ports and 12 QSFP ports. Similarly, the SX1016 delivers wire-speed running 64 ports of 10GbE. See Figures 2 and 4.

THE BOTTOM LINE

The Mellanox SX1036 Switch delivers:

- 1 36 ports of 40GbE, 2.88Tbps, at full wire speed on a single chip
- 2 40GbE L2 latency less than 223ns for all frame sizes
- 3 40GbE L3 latency less than 333ns for all packet sizes
- 4 10GbE L2 latency less than 269ns for all frame sizes
- 5 10GbE L3 latency less than 426ns for all packet sizes
- 6 Ultra low power consumption. Less than 0.6W per 10Gbps of throughput





Test Results

Layer 2

SX1036

RFC 2544 tests were run using frame sizes ranging from 64-bytes to "jumbo" frames of 9216-bytes in an all-to-all port configuration. All tests demonstrated wire-speed throughput with no loss.

Tests were run first using 40GbE and again running 10GbE. (When 10GbE results were measured, an additional 12 ports of 40 GbE were running wire speed.) Average cut-through latency was below 230ns across all frame sizes at 40GbE and

below 270ns at 10GbE rates. See Figure 1.

SX1016

The same RFC 2544 tests were run on the SX1016 which also delivered wire-speed throughput with no loss. Average latency was below 315ns across all frame sizes. See Figure 2.

Layer 3

SX1036

Tests were run using packet sizes ranging from 64-bytes to "jumbo" packets of 9216-bytes in an one-to-one port configuration. All tests demonstrated wire-speed throughput with no loss.

Mellanox Technologies, LTD

SX1016 & SX1036

10GbE & 40GbE Performance & Power Consumption



Tested April 2012

Tests were run first using 40GbE and again running 10GbE. Average latency was below 335ns across all packet sizes for 40GbE and below 430ns in the 10GbE tests. See Figure 3.

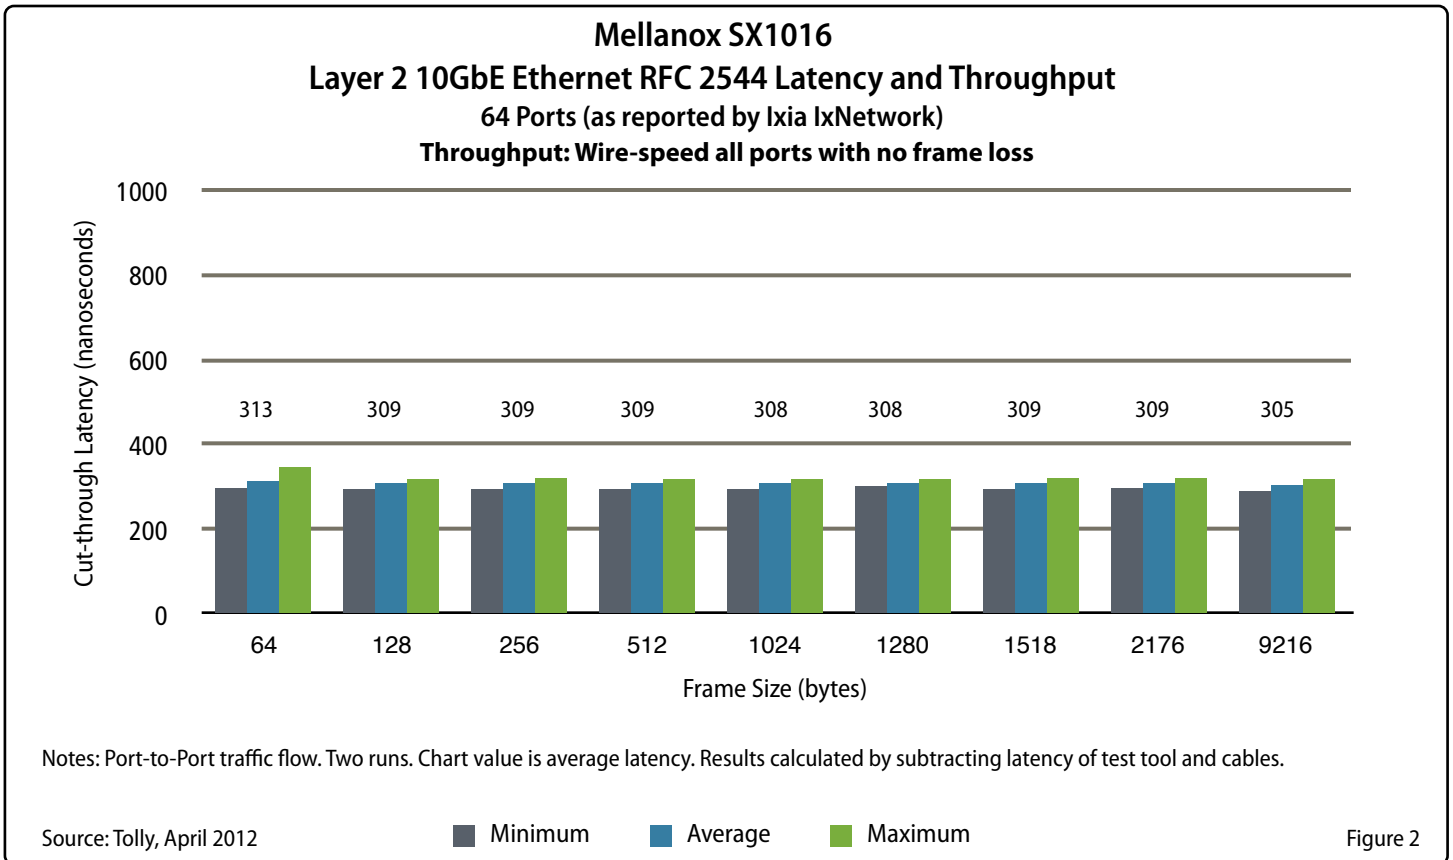



Figure 2



Mellanox SX1036 Layer 3 40GbE/10GbE Ethernet RFC 2544 Latency and Throughput

(as reported by Ixia IxNetwork)
Throughput: Wire-speed with no packet loss



Notes: Two ports tested. Two runs. Chart value is average latency. Results calculated by subtracting latency of test tool and cables.

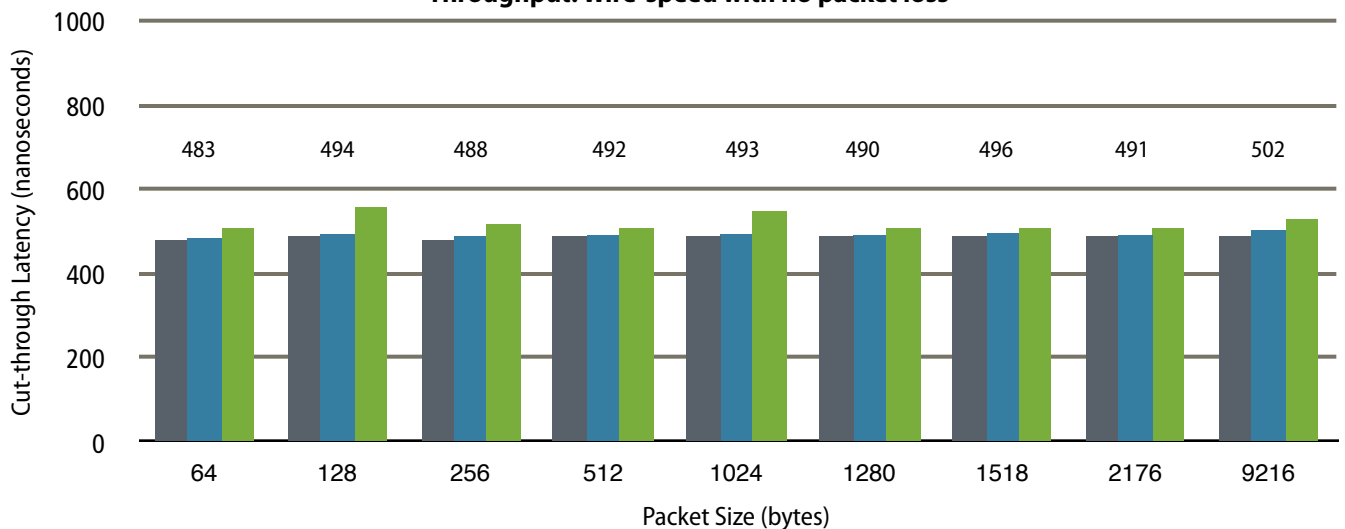
Source: Tolly, April 2012

■ Minimum ■ Average ■ Maximum

Figure 3

Mellanox SX1016 Layer 3 10GbE Ethernet RFC 2544 Latency and Throughput

(as reported by Ixia IxNetwork)
Throughput: Wire-speed with no packet loss



Notes: Two ports tested. Two runs. Chart value is average latency. Results calculated by subtracting latency of test tool and cables.

Source: Tolly, April 2012

■ Minimum ■ Average ■ Maximum

Figure 4




Figure 5

SX1016

The same RFC 2544 tests were run on the SX1016 which also delivered wire-speed throughput with no loss. Average latency was below 505ns across all packet sizes. See Figure 4.

IP Multicast

SX1036

RFC 3918 multicast throughput and latency tests were run using packet sizes ranging from 128-bytes to “jumbo” frames of 9216-bytes in a one-to-all port configuration. All tests demonstrated wire-speed throughput with no loss.

Tests were run first using 40GbE and again running 10GbE. Average cut-through latency was below 215ns across all packet sizes. See Figure 5.

SX1016

The same RFC 3918 multicast throughput and latency tests were run on the SX1016, using packet sizes ranging from 64-bytes to “jumbo” frames of 9216-bytes, which also delivered wire-speed throughput with no loss. Average latency was below 325ns across all frame sizes. See Figure 6.

Congestion Control

SX1036

L2 and L3 congestion control tests (RFC 2889) were run across the range of frame sizes used for the L2 performance tests, 64-byte to 9216-byte “jumbo” frames. Results (not illustrated) showed that there was no head-of-line blocking and that back pressure was generated.

SX1016

The same tests were run on the SX1016 with the same results.

Power Consumption

SX1036


Power measurements of the SX1036 were taken according to that ATIS specification. That is, measuring power in three system states: 0% load, 10% load and 100% load. The weighted values (10%, 10% and 80% respectively) are used to determine the ATIS power value.

Running 36 ports of 40GbE over copper cable, the SX1036 delivered an ATIS value of 78.30 which translates into 0.55W per 10Gbps of throughput. See Table 1.

SX1016

The same test was run on the SX1016 running 64 ports of 10GbE over copper cable. The SX1016 delivered an ATIS value of 61.64 which translates into 0.96W per 10Gbps of throughput. See Table 1.

Mellanox SX1016
IP Multicast 10GbE Ethernet RFC 3918 Latency and Throughput
1 Port to 63 Ports Configuration (as reported by Ixia IxNetwork)
Throughput: Wire-speed all ports with no packet loss



Notes: Two runs. Chart value is average latency. Results calculated by subtracting latency cables.

Source: Tolly, April 2012

Figure 6

Mellanox SX1036 & SX1016
Power Measurements and Watt-per-10-Gigabit Throughput
No Frame Loss

Device	Cable Type	Ports	Power Draw @ 100% (W)	ATIS Weighted Power (W)	Watts-per-10-Gigabit Throughput (ATIS-based)
SX1036	Copper	36 x 40GbE	83.04	78.3	0.55
SX1016	Copper	64 x 10GbE	63.56	61.64	0.96

Notes: Tests run with all ports. Measurements taken at 0%, 10% and 100% load and weighted per ATIS specification.

Source: Tolly, April 2012

Table 1

Systems Under Test: Configuration

Device	Description	Ports	Versions
SX1036	40GbE and 10GbE L2/L3 Switch	36 x 40GbE or 48 x 10GbE & 12 x 40GbE	(Software) SX_3.2.0085 (Hardware) SX_9.1.3372
SX1016	10GbE L2/L3 Switch	64 x 10 GbE	(Software) SX_3.2.0085 (Hardware) SX_9.1.3372

Source: Tolly, April 2012

Table 2

Test Setup & Methodology

Performance Testing

Device defaults were used for both systems under test with the exception of setting the MTU size to 9K for "jumbo" frames. See Table 2 and Figure 7.

Standard RFC performance tests were run with bidirectional port-to-port traffic flow (as noted in the

main text and figures) using the Ixia IxNetwork version 6.20.601.18 EA. All latency measurements were taken running at 100% load.

For all performance testing except IP multicast, cut-through latency was calculated by measuring the latency of 1m cable for SX1036 36 x 40GbE, 3m for SX1036 48 x 10GbE/12 x 40GbE and 2m for the SX1016 10GbE and the latency of the test gear and subtracting these values from the values reported by the test tool.


For the IP multicast test, only cable latency was subtracted from the reported value.

Power Measurement (ATIS)

Power measurements were recorded and calculated using the ATIS specification. This provides for a value that combines measurements from three states: idle, low load and full load - as noted earlier in this report.

For more information on ATIS, see <http://www.atis.org>.

Mellanox SX1016 & SX1036 10/40GbE Switches



Source: Mellanox, LTD, April 2012

Figure 7




About Tolly...

The Tolly Group companies have been delivering world-class IT services for more than 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 by email at sales@tolly.com, or by telephone at +1 561.391.5610.

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

Test Equipment Summary

Vendor	Product	Web
ixia	ixNetwork 6.20.601.18 EA ixOS 6.20.800.8 EA	 http://www.ixiacom.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.