Lippis Report

Cisco Nexus® 9508 Power Efficiency

A Report on the:

Cisco Nexus® 9508 Switch

October, 2013

Acknowledgements

We thank the following people for their help and support in making possible this firstevaluation of the fastest 40GbE data center Ethernet switch:

We thank the following people for their help and support in making possible this first evaluation of the fastest 40GbE data center Ethernet switch:

Bradley Wong, Distinguished Engineer at Cisco, Lilian Quan, Technical Marketing Engineer at Cisco and Arun Selvarajan Technical Marketing Engineer for their help and engineering support.

Michael Githens, Lab Program Manager at Ixia, for his technical competence, extra effort and dedication to fairness as he executed test week and worked with participating vendors to answer their many questions.

Jim Smith, VP of Marketing at Ixia, for his support and contributions to creating a successful industry event.

Bill Nicholson for his graphic artist skills that make this report look amazing.

Jeannette Tibbetts for her editing that makes this report read as smooth as a test report can. Steven Cagnetta for his legal review of all agreements.

Cisco Nexus® 9508 Power Efficiency

Cisco launched its Nexus® 9500 series of data center modular switches, which Cisco promises offers not only the highest port density of 10/40/100GbE, and the most programmable and fastest packet forwarding modular data center switch in the industry, but the most power efficient too. To verify and measure the power efficiency claim, Cisco engaged the Lippis/Ixia team to test the new Nexus® 9508 modular data center switch for power draw at scale.

One of the many unique attributes of the Cisco Nexus® 9508 series is its new approach to chassis design. The Nexus® 9508 is an eight-slot chassis, meaning it supports eight line card payload slots – an important distinction as these slots are used solely for line cards versus a mix of line cards and supervisor modules. Two supervisor modules are located underneath the line card modules offering full 1-for-1 redundancy. Under the supervisor

modules are eight power supply slots capable of supplying some 3k Watts of AC power each to the Cisco Nexus® 9508; however only two are needed to power a fully populated chassis. An additional one or two power supplies can be used for redundancy in a 2-plus-1 or 2-plus-2 grid configuration. Load sharing across power supplies increases reliability.

The remaining four power supply slots increase mean time between failure, or MTBF, and provide additional power headroom for future upgrades to support 100GbE at scale, pluggable optics and multiple future generations of ASICs that would enable the Cisco Nexus® 9508 to scale to some 10Tbps fabric modules. Pluggable optics provides connectivity flexibility and mitigates against fixed optics deployments due to limited power budget. If history is any guide, then this chassis will be in operation for

well over a decade supporting multiple generations of line cards, fabric modules, supervisors and ASICs.

One of the key design attributes is that the Nexus® 9508 has no midplane for line card-to-fabric module connectivity enabling unobstructed front-to-back airflow, thus contributing to its power and cooling efficiency.

Fan trays speed is automatically regulated by system software—increasing fan speed if additional cooling is required and/or slowing fans down if additional cooling is not required, all in an effort to conserve power consumption. There are three fan trays, two system controllers and six fabric modules accessible from the back of the chassis. The bottom line is that the Cisco Nexus® 9508 chassis was designed for reliability, future scalability and power efficiency in keeping with today's market requirement of green data center computing.

     Evaluation conducted on Ixia test equipment   www.lippisreport.com

As mentioned, the Cisco Nexus® 9508 does not have a midplane for connecting line cards to fabric modules, allowing for completely unobstructed front-to-back airflow. The line cards contain only two-to-four ASICS with no buffer bloat and the mix of 28nm Cisco and 40nm Broadcom custom and merchant ASICs respectively. In terms of power, the Cisco Nexus® 9508's power supplies are platinum rated and are said to exhibit 90% to 94% power efficiency across all workloads.

All of these chassis design points driving more power efficiency leads to a design that has no midplane, no airflow disruption, auto-regulating fans, efficient power supplies, smaller number of ASICs with integrated buffering, load sharing between power supplies and pluggable optics, which we measure in terms of Watts/10GbE. Note that historically, data center switches have not been the most power efficient class of network devices, considering a 5-to-10W per 10GbE port draw as typical for

a data center aggregation switch with 352 10GbE ports, which adds up to a total power draw of some 2 to 4kW. An entire server rack draws between 8kW to 10kW. Without new power-efficient designs, power consumption would only increase as 10, 40 and 100GbE port density requirements increase and consumes more power. Therefore, in order to cater for the bump in speed and performance, there needs to be a new approach in designing data center switches that are much more power efficient.

Cisco Nexus® 9508 Power Test Configuration

No data center switch has been tested for power draw at full 288 40GbE ports, thus Lippis/Ixia engineers designed a new approach to test at this scale. To test the Nexus® 9508, we connected 288 40GbE ports to three Ixia XG12 High Performance Chassis running its IXOS 6.50 and IxNetwork 7.10. To deliver

288 40GbE ports to the Nexus® 9508, an Ixia CXP port capable of supplying 120Gbps of network traffic was split into three 40GbE QSFP+ ports. This was accomplished via the Ixia Xcellon-Multis XM40GE12 QSFP+FAN 40GE Load Modules, which support four CXP ports. Eight Ixia Xcellon-Multis load modules populated each XG12 chassis, delivering 96 40GbE streams of line rate traffic. Three XG12 populated with 96 40GbE each delivered 288 40GbE of line rate traffic flow into the Nexus® 9508 at varying packet sizes.

The Nexus® 9508's 288 40GbE ports were connected via Ixia Multis CXP-to-3-40GE QSFP Active Optical Cable (AOC). The optical wavelength was 850NM, and 3-meter optical cable connected test gear and Nexus 9508. The Ixia modules connect to the Nexus® 9508 via its 9636PQ 40GbE modules that support 12-40GbE QSFP+ plus 48-10GbE SFP+ ports. Lastly, the Nexus® 9508 ran a pre-release version of the NXOS 6.1(2)I1(1) operating system.

Cisco Nexus® 9508 Test Configuration			
	Hardware	**Software Version**	**Port Density**
Device under test	Cisco Nexus® 9508 http://www.cisco.com/en/US/products/	NXOS 6.1(2)I1(1)	288-40GbE
Test Equipment	(3) Ixia XG12 High Performance Chassis	IxOS 6.40 EA IxNetwork 7.10 EASP1	
	(32) Xcellon-Multis XM40GE12 QSFP+FAN 40GE Load Module		
	http://www.ixiacom.com/		
Cabling	Ixia's Multis CXP-to-3-40GE QSFP Active Optical Cable (AOC)		
	http://www.ixiacom.com/pdfs/datasheets/xcellon-multis-100-40-10-loadmodule.pdf		

Cisco Nexus® 9508 Power Measurement Findings

With the Cisco Nexus® 9508 connected to Ixia test gear, the team completed the standard Lippis/Ixia iMix tests of 0, 30 and 100% line rate traffic load. These were transmitted to the Nexus® 9508; see test methodology below. A Voltech PM300 three-phase power analyzer was used to measure power consumption at each traffic load level. With power levels in hand, engineers then calculated the power consumption of the Cisco Nexus® 9508 on a 10 GbE and 40GbE per port basis. To compare different electronic equipment, we then calculate the Cisco Nexus® 9508's TEER (Telecommunications Energy Efficiency Ratio) value. TEER is a measure of network-element efficiency quantifying a network component's ratio of "work performed" to energy consumed. Note that a larger TEER value is better as it represents more work done at less energy consumption.

We found that the Cisco Nexus® 9508 is the most power efficient modular data center switch in the industry with a power draw per 10GbE port of only 3.85 Watts and 15.4 Watts per 40GbE port, inclusive of optics and live traffic. For reference, many Christmas tree light bulbs draw some 5-Watts. We calculate that it would cost $5,402.84 annually to power a fully populated 288 40GbE port Cisco Nexus® 9508 data center modular switch.

Evaluation conducted on Ixia test equipment   www.lippisreport.com

The three-year cost to power the Cisco Nexus® 9508 is estimated at $16,208 when fully populated with 1,152 10GbE ports. The three-year energy cost is some 3.52% of list price. The Cisco Nexus® 9000's TEER value is 246 – the highest TEER value we have measured in these Lippis/Ixia tests to date. Remember higher TEER values are better than lower ones. Previous data center modular switch TEER values ranged from a high of 117 to a low of 44 making the Nexus® 9508 the most power-efficient modular switch by significant margin.

The Cisco Nexus® 9508 demonstrates the epitome of mechanical chassis engineering. The Cisco engineers did not just deliver incremental improvement to the status quo in data center switching, but a leap forward in new/enhanced features that set a new high bar in power efficiency. The Cisco Nexus® 9508 is the most power efficient chassis tested during the Lippis/Ixia indus-

try tests. Previous Lippis tested core switches draw power from a high of 22 Watts to a low of 9 Watts per 10GbE port, but at port densities in the range of 128-to-352 10GbE ports rather than the Nexus® 9508's 1,152 10GbE. The Cisco Nexus® 9508 draws less than half the power of the lowest core switches we have tested with over three times the port density and traffic load!

This Lippis test of the Nexus® 9508 verifies that Cisco engineers' efficient modular chassis design of no midplane, no airflow disruption, auto-regulating fans, efficient power supplies, small number of ASICs with integrated buffering, load sharing between power supplies and pluggable optics do equate to a new industry standard in low power consumption for modular data center switches.

Power Consumption Test

Test Objective: This test determines the Energy Con¬sumption Ratio (ECR), the ATIS (Alliance for Telecom¬munications Industry Solutions) 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 auto¬matic or manual mode to determine the L2/L3 forwarding performance while concurrently making power, current and voltage readings from the Voltech PM300 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 Telecom¬munication 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

Cisco Nexus® 9508 Power Consumption Test	
Watts$_{ATIS}$/10GbE Port	3.85
3-Year Cost/ Watts$_{ATIS}$/10Gbe	$14.07
Watts/40GbE port	15.4
Total Power cost/3-Year	$16,208
Teer Value	246
Cooling	Front to Back

Evaluation conducted on Ixia test equipment   www.lippisreport.com

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 simply as Watts, Watts/10GbE or Watts/40GbE.

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 Watts to identify ATIS power consumption measurement on a per port basis.

With the Watts, we calculate a three-year energy cost based upon the following formula:
Cost/Watts/3-Year = (Watts /1000)*(3*365*24)*(0.1046)*(1.33), where Watts = 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. We use the June 2010 Dept. of Energy Electric Power Monthly rate as to maintain consistency over the past few years of testing. For more information on Dept. of Energy Electric Power Monthly rates, please see (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 consump¬tion test as conducted at the iSimCity lab for each product.
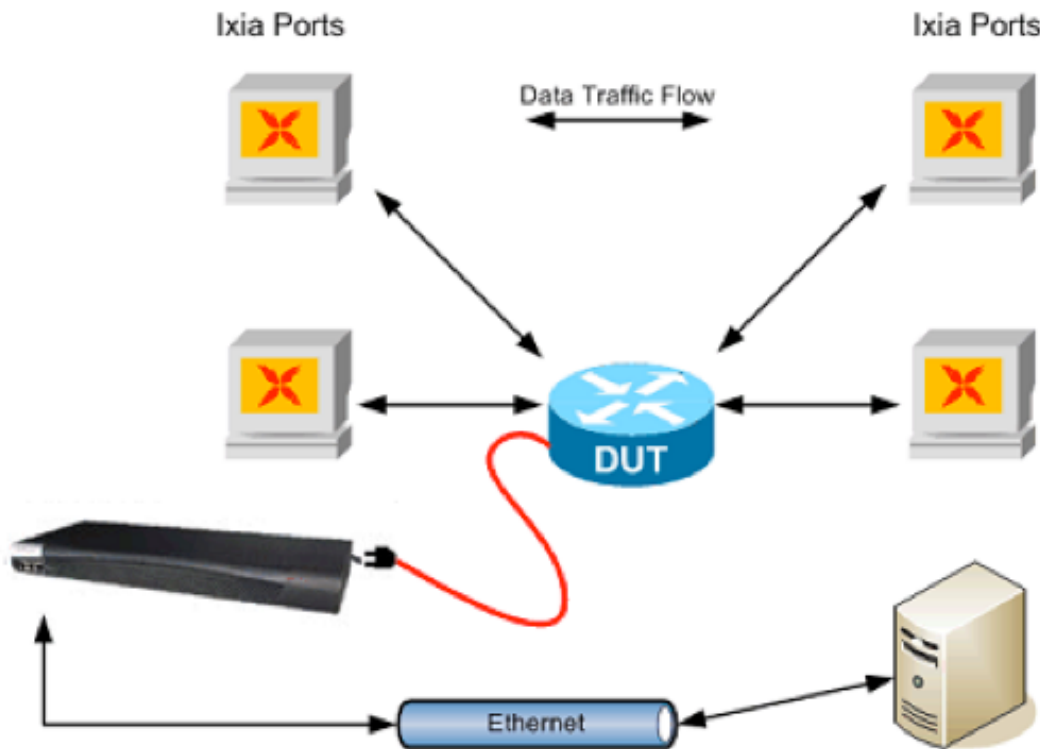
Terms of Use

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 is not responsible for, and you agree to hold Lippis 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 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. 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 200,000 times; ITunes 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. He is also the co-founder and conference chair of the Open Networking User Group, which sponsors a bi-annual meeting of over 200 IT business leaders of large enterprises. Mr. Lippis is currently working with clients to design their private and public virtualized data center cloud computing network architectures with open networking technologies 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, 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 firms' 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.