

Data Center 40GE Switch Study

**Cisco Nexus 9508
DR 140126L**



Contents

1	Executive Summary	3
2	Introduction	4
3	Test Bed Setup and How We Did It	5
4	Throughput and Latency Performance Test (RFC 2544)	8
4.1	RFC 2544 Throughput Test System and DUT Configuration	8
4.2	RFC 2544-based 40GE Layer 3 Throughput	9
4.3	RFC 2544 40GE Layer 3 Latency – 50% Load	10
4.4	RFC 2544 40GE Layer-3 Latency – 100% Load	11
4.5	RFC 3393-based 40GE Layer 3 Delay Variance - 50% Load.....	12
4.6	RFC 3393-based 40GE Layer-3 Delay Variance - 100% Load	13
5	Fully Meshed Throughput and Latency Test (RFC 2889)	14
5.1	RFC 2889-based Fully Meshed Test System and DUT Configuration.....	14
5.2	RFC 2889-based 40GE Layer-3 Full Mesh Throughput.....	15
5.3	RFC 2889-based Full Mesh L3 Latency – 50% Load	16
5.4	RFC 2889-based Full Mesh L3 Latency – 100% Load	17
6	IP Multicast Throughput and Latency (RFC 3918)	18
6.1	RFC 3918-based IP Multicast Test System and DUT Configuration	18
6.2	RFC 3918-based 40GE IP Multicast Throughput.....	19
6.3	RFC 3918-based 40GE IP Multicast Latency – 100% Load.....	20
7	Mean Time between Failure (MTBF) Analysis	21
7.1	Component Count Used for MTBF Calculations	23
8	Competitive Analysis	24
8.1	RFC 2544-based 40GE Layer 3 Throughput Test – 100% Load	24
8.2	RFC 2544-based 40GE Layer 3 Latency Test – 50% Load.....	25
8.3	RFC 2544-based 40GE Layer 3 Latency Test – 100% Load.....	26
8.4	RFC 3393-based 40GE Layer 3 Delay Variance - 50% Load.....	27
8.5	RFC 3393-based 40GE Layer 3 Delay Variance - 100% Load.....	28
8.6	RFC 2889-based 40GE Layer 3 Full Mesh Throughput	29
8.7	RFC 2889-based Full Mesh L3 Latency – 100% Load	30
8.8	RFC 3918-based 40GE IP Multicast Throughput	31
8.9	RFC 3918-based 40GE IP-Multicast Latency – 100% Load.....	32
8.10	Data Center Switch Power Consumption	33
9	About the Miercom 40GE Switch Industry Study	34
10	About Miercom	34
11	Use of This Report	34

1 Executive Summary

Miercom is pleased to publish these results, of recent testing we conducted of one of the highest-density and highest-throughput data-center switches we have ever seen – the powerful 8-slot Cisco Nexus 9508.

Conducted in January 2014, Miercom engineers were invited to give the impressive new switch a work-out, and we did. It was configured to scale – that is, packed with eight of the vendor's latest and highest density line cards – the *36p 40G QSFP+ ACI spine line card* -- each offering 36 x 40GE ports. The system we tested, Nexus 9508 with eight line cards, supported 288 x 40GE (QSFP+ fiber) ports – all in one 13 RU high chassis. Other line cards (not tested) enable the same 8-slot system to provide 1,152 x 10GE ports.

The system incorporates a novel chassis design. Rather than classical multi-slot switches, where a backplane or mid-plane typically connects line cards and fabric modules, Nexus 9500 is the first chassis in the industry to employ a mid-plane free design. Its fabric modules are directly attached to all the line cards through connecting pins, eliminating the airflow obstruction caused by a chassis mid-plane. This midplane-free design greatly increases cooling efficiency and eases upgrade to higher speeds in the future. The Nexus 9500 switch supports up to six fabric modules, each yielding 5.12 terabit/second (Tbps) of bidirectional connectivity throughput.

In addition to validating the vendor's published throughput figures, our testing examined aspects of the switch as far as replicating packets for multicast traffic handling, and the delay, or latency, that the switch applies to packets traversing the system. In each test scenario we were frankly impressed.

Key findings for the Cisco Systems Nexus 9508:

- The switch delivered full line-rate throughput on every 40GE port for all packet sizes, with zero packet loss, based on full-mesh, 288-port x 40GE, high-density throughput testing. Using IP packets and hundreds of flows with different IP source and destination addresses, each 36p 40G line card can transmit 1.44 Tbps of traffic.
- Traffic traversing the switch incurs consistent low latency, even under heavy load, regardless of packet size, based on the fully populated configuration tested.
- The switch applies very low jitter – Delay Variance – regardless of frame size, again based on our fully populated test configuration.

All of the performance testing was conducted in accordance with widely accepted IETF standards, including Requests for Comments (RFCs) 2544, 2889, 3399 and 3918.

Miercom independently substantiates the performance of the Cisco Systems Nexus 9508 and awards the ***Miercom Performance Verified Certification*** in recognition of the product's proven performance in our ongoing Data Center Class 40GE Switch Testing Study.

2 Introduction

Like the other models in the family, the Cisco Nexus 9508 is modular and employs the same innovative "mid-plane-free" architecture. A chassis mid-plane or backplane – essentially a high-speed bus – has long been the traditional means in modular switches for connecting line cards into and across the switching fabric. In the Nexus 9500 Series, line cards and fabric modules directly attach to each other via connecting pins and a precise alignment mechanism. Each of up to six fabric modules connects directly to all line cards.

This is one of the highest density and highest throughput 8-port switches we have tested and confirmed "wire speed" performance on all ports concurrently. In other words, the internal architecture is more than capable of switching the 288 ports of 40GE traffic delivered to the system in our tests conducted. That is a bit more than 23 Terabits per second of bidirectional traffic.

The test traffic was not all random bits. It was legitimate IP packets, and the Nexus 9508 had to examine each, check its forwarding table, and route each one appropriately. The testing is discussed more in the following sections.

We created this load with a battery of Spirent test systems, which collectively delivered 40GE of traffic on each of the Nexus 9508's 288 ports, and then checked the outputs on all of the ports to make sure it was all there and all handled correctly. The results proved that all of the IP packets were processed and forwarded by the Nexus 9508 properly.

Since the test traffic was all IP packets, all of the regular IP traffic rules had to be obeyed – interframe gaps, packet-handling overhead and so on. This means that traffic applied never actually filled the full 40GE clock rate of the optical link. This is true of any link carrying IP traffic for that matter. Payload gets closer to 100 percent of link speed as larger frame sizes are used (1,518-byte, and "Jumbo" Ethernet frames, 9,216 bytes). Conversely, frame *rate* diminishes as the frame size grows. The entire spectrum of frame sizes was applied in testing the Nexus 9508, and in *all* cases the maximum load was accepted and forwarded – with no frame loss.

A number of vendors are preparing or already shipping switches designed to function at the heart of today's data centers. Miercom is slated to test several of these in the months ahead. But the current leader, based on all the test metrics we applied here – latency, throughput – is the Cisco Nexus 9508. And the bar for the others has been set very high.

3 Test Bed Setup and How We Did It

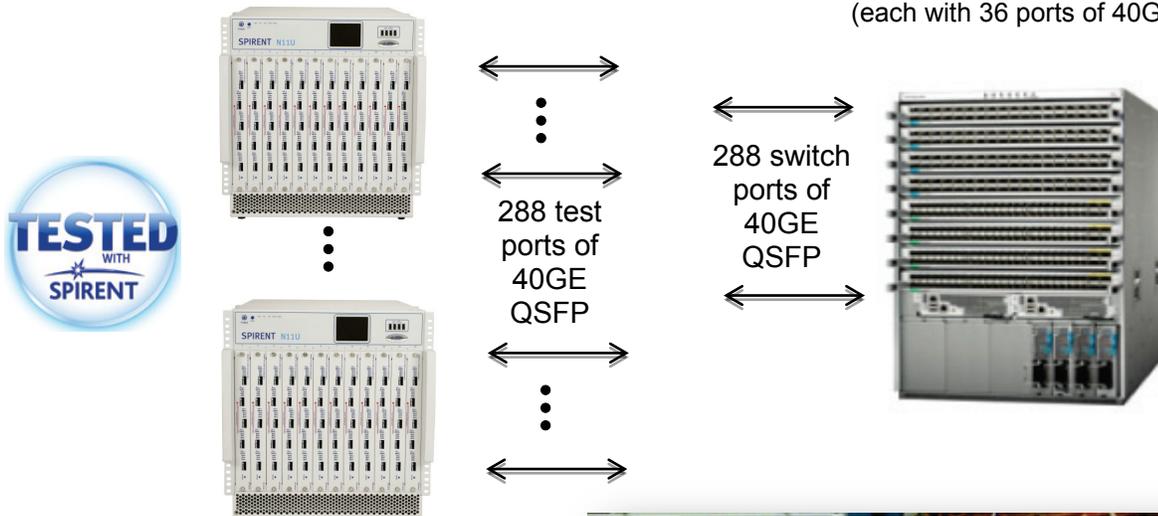
All testing of the Cisco Nexus 9508 employed the same test system – multiple Spirent N11U multi-slot chassis test systems, which are centrally controlled (see details of the Spirent test system below). Sufficient N11U systems were aggregated to yield 288 bidirectional 40GE test ports, one for each corresponding port on the Cisco Nexus 9508.

Test System:

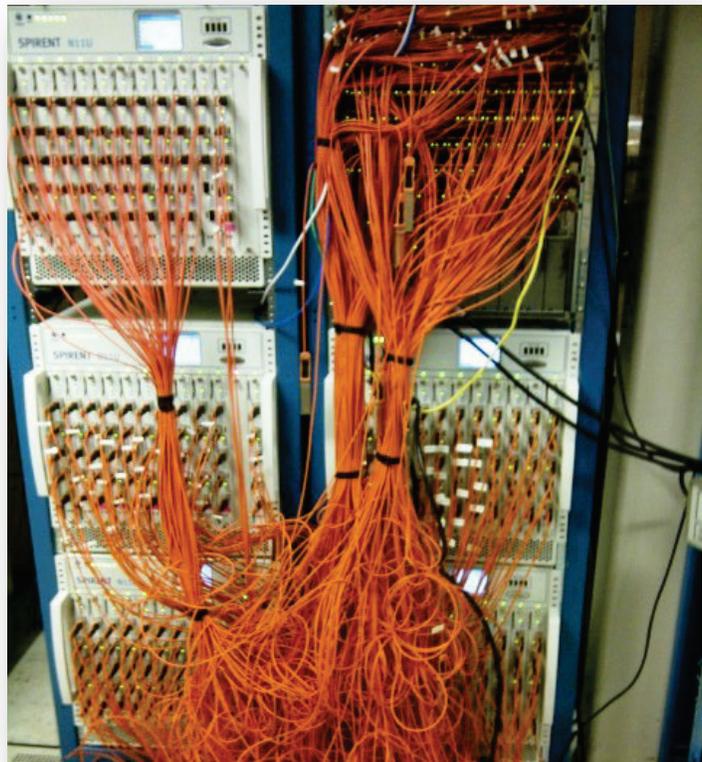
Multiple Spirent N11U multi-slot chassis systems

Device Under Test (DUT):

Cisco Nexus 9508
Single chassis; 8 line cards
(each with 36 ports of 40GE)



*Actual Test Site:
Shown right is the switch under test,
surrounded by five Spirent test systems.*



Once the array of Spirent N11U tests systems were connected to the Nexus 9508, one Spirent test port each connected to a 40GE port on the switch, and then connectivity confirmed on each port, we did not have to re-cable for the duration of the testing. The multiple chassis of the Spirent system are centrally controlled.

The battery of tests we applied are each discussed individually in this report. They examine different aspects of the device under test, or DUT, in this case the Nexus 9508. The procedures for testing Layer-2 and Layer-3 switches and routers have been more-or-less standardized in recent years, and this testing employed a number of those procedures. Of course the sheer number and speed of switch ports, and the mind-boggling volume of traffic applied to fill them, are orders of magnitude more today for data-center switches than just a few years ago. Even so, the same procedures still apply.

Four of the standards we used in this testing, and which the Spirent test system incorporates, are publicly available as Internet Requests for Comments, or RFCs. The ones applicable here include RFCs 2544, 2889, 3393 and 3918 – for throughput and latency measurements of IPv4 (Layer-3) bidirectional traffic.

RFC 2544, issued in 1999, describes how to conduct basic benchmark tests for latency and throughput measurement. Bidirectional Layer-3 (IP) traffic is applied on port pairs on the device under test (DUT) so that test traffic is processed across the switch fabric.

RFC 2889, issued in 2000, is a reference for conducting more stressful full-mesh tests for latency and throughput measurement. The Spirent test system provides a fully meshed bidirectional traffic flow for these measurements and fully stresses the switch fabric.

RFC 3393, a 2002 document that defines the metric for variation in delay, or jitter, of IP packets passing through a system.

RFC 3918, a 2004 specification, addresses throughput and latency measurement for IP multicast traffic. Based on RFC 3918, the Spirent system supports a combination of traffic profiles with an adjustable number of transmit and receive ports for multicast traffic flows.

Tests were conducted for unicast and multicast traffic throughput, latency, and delay variation. The maximum throughput achievable through the switch was verified, as was proper multicast packet replication.

Latency measurements for all tests were conducted using FIFO (First In, First Out) mode as opposed to LIFO (Last In, First Out) mode with the test equipment. This is the appropriate mode for testing “Cut-Through” low latency switches such as the Nexus 9608 included in this report. The mode difference records actual latency for the transit time of the frame through the switch without subtracting (erroneously) extra time for the frame ingress to the receiving port. We also use FIFO mode for testing “Store and Forward” technology switches. It is equally applicable for testing switches employing store and forward operation in that the test will reveal the actual latency experience as a customer or connected entity would experience directly connected to the switch. This practice may show latency results slightly higher than other labs report for our latency measurements, especially for those tests involving larger frame sizes (approximately 1.2 μ sec for 9,216

Byte frames). Miercom tests all switches the same way, equally, to ensure fair and accurate comparison data.

The multiple Spirent N11U chassis were driven by Spirent's sophisticated TestCenter application, which offers a vast library of test methodologies. This was the primary traffic-generation system that delivered test network traffic through the Nexus 9508 switch.

Spirent TestCenter is an end-to-end testing solution that delivers high performance with deterministic results, so that network equipment manufacturers, Service Providers and enterprises can use it to test, measure and assure their networks and deploy services with confidence.

Spirent TestCenter provides measurement solutions for next generation networks – from traditional performance testing to the rigorous analysis of Data Centers, cloud computing, Virtualization, Mobile Backhaul and High Speed Ethernet.

Network equipment manufacturers, Service Providers, product designers as well as testing architects/specialists trust Spirent TestCenter for its unmatched performance, flexibility, realism and optimization.

The tests in this report are intended to be reproducible for users who want to recreate them, with the appropriate test and measurement equipment. Those interested in repeating these tests may contact Miercom at reviews@miercom.com for more details on the configurations applied in this testing. A Miercom professional services sales representative can provide details for assistance.

4 Throughput and Latency Performance Test (RFC 2544)

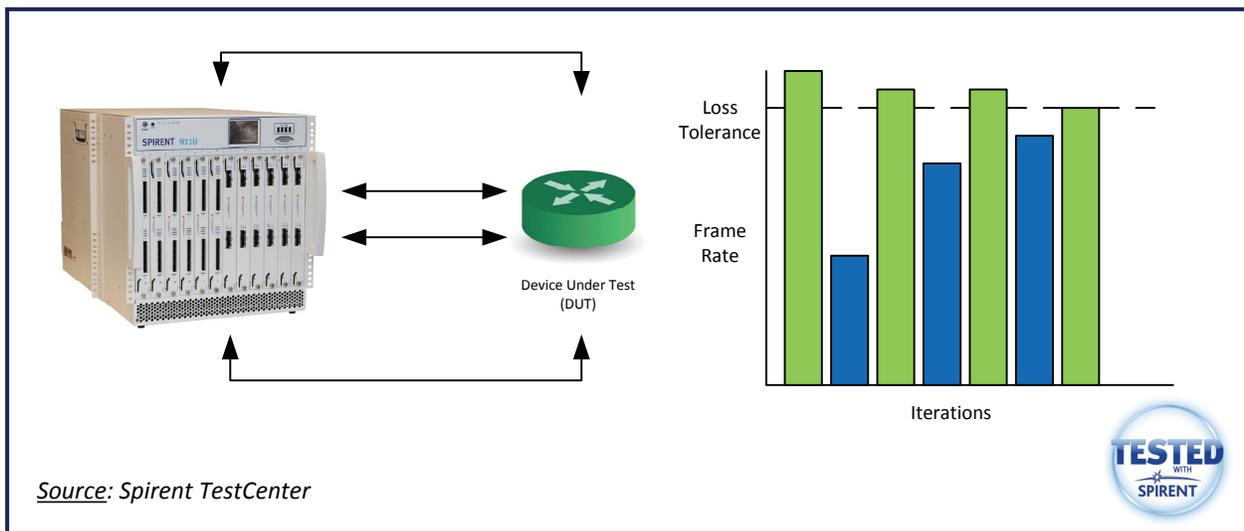
This throughput test determines the maximum rate at which the Nexus 9508 switch receives and forwards traffic without frame loss. Frames are sent at a specified rate, and then the rate is continually stepped up, using a binary search algorithm, to determine the maximum rate at which the switch does not lose frames. Frames can be MAC only, IPv4, IPv6 (with or without Extension Headers) or an IPv4/IPv6 mixture. IPv4 frames, of a complete spectrum of frame sizes, were used for testing the Nexus 9508.

Once the maximum traffic rate is established for a particular frame size, latency through the switch is then calculated – by subtracting the transmit time stamp from the receive time stamp. Based on traffic tests that usually run for a minute or two, the minimum, maximum and average latencies are reported.

The test system's load generator was configured to forward and receive traffic to and from each directly connected port on the switch. As a rule, frames are initially sent at the maximum theoretical rate based on the speed of the port. This test is configured with a one-to-one traffic mapping. The results below show the maximum throughput of the switch without frame loss. When a switch accepts and successfully processes and forwards all traffic at the maximum theoretical rate based on the speed of the port, the switch is said to perform at "wire speed" or "full line rate" for the particular packet size.

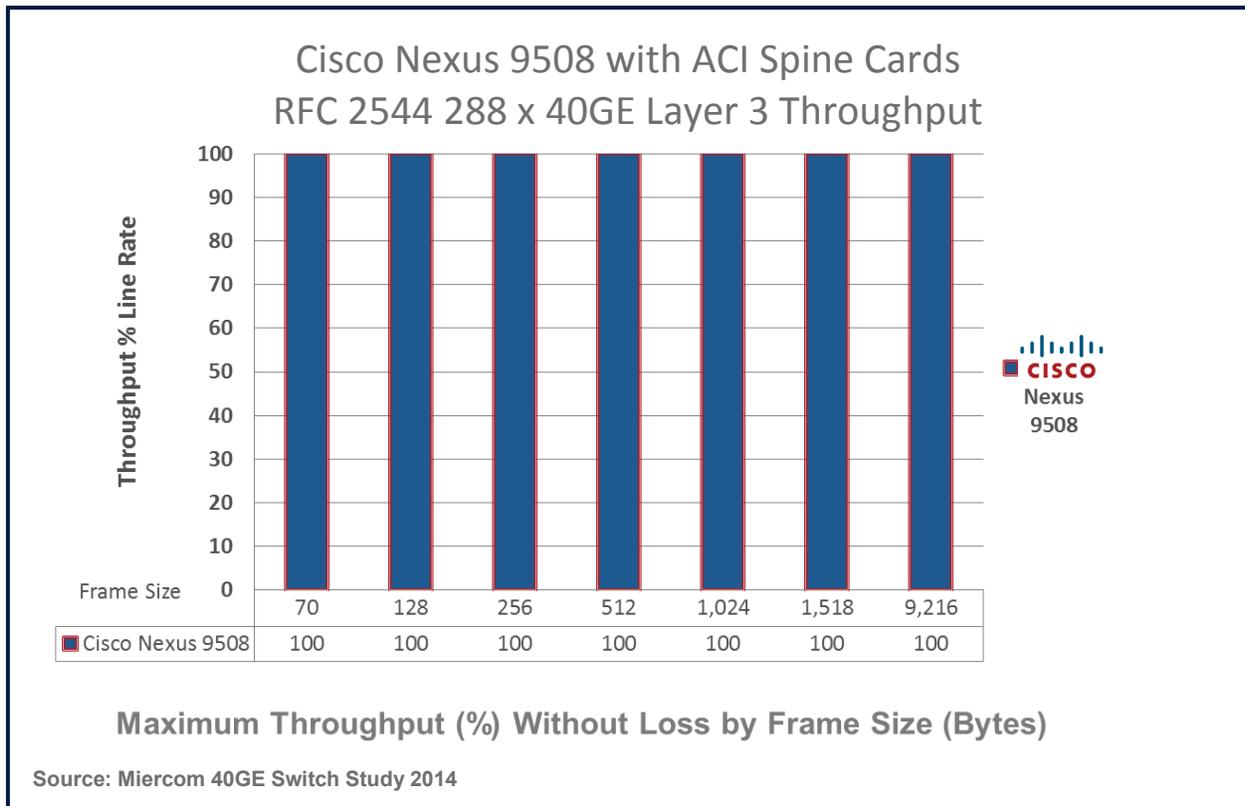
The Nexus 9508 was configured for Layer-3 switching (IP routing). Port-pair combinations were assigned within the test system so that bi-directional traffic was transmitted between line-card ports across the fabric modules, in accordance with RFC 2544. All 288 of the Nexus 9508's 40GE ports were connected to the Spirent load-generation system for these tests. The test system established that traffic delivered to and received from each port on the Nexus 9508 could be sent at "wire speed," at all tested packet sizes, without any data loss.

4.1 RFC 2544 Throughput Test System and DUT Configuration



4.2 RFC 2544-based 40GE Layer 3 Throughput

"The Cisco Nexus 9508 with eight 36-port 40GE ACI spine line cards delivered line-rate performance, for all packet sizes during RFC 2544-based Layer-3 Throughput Tests."

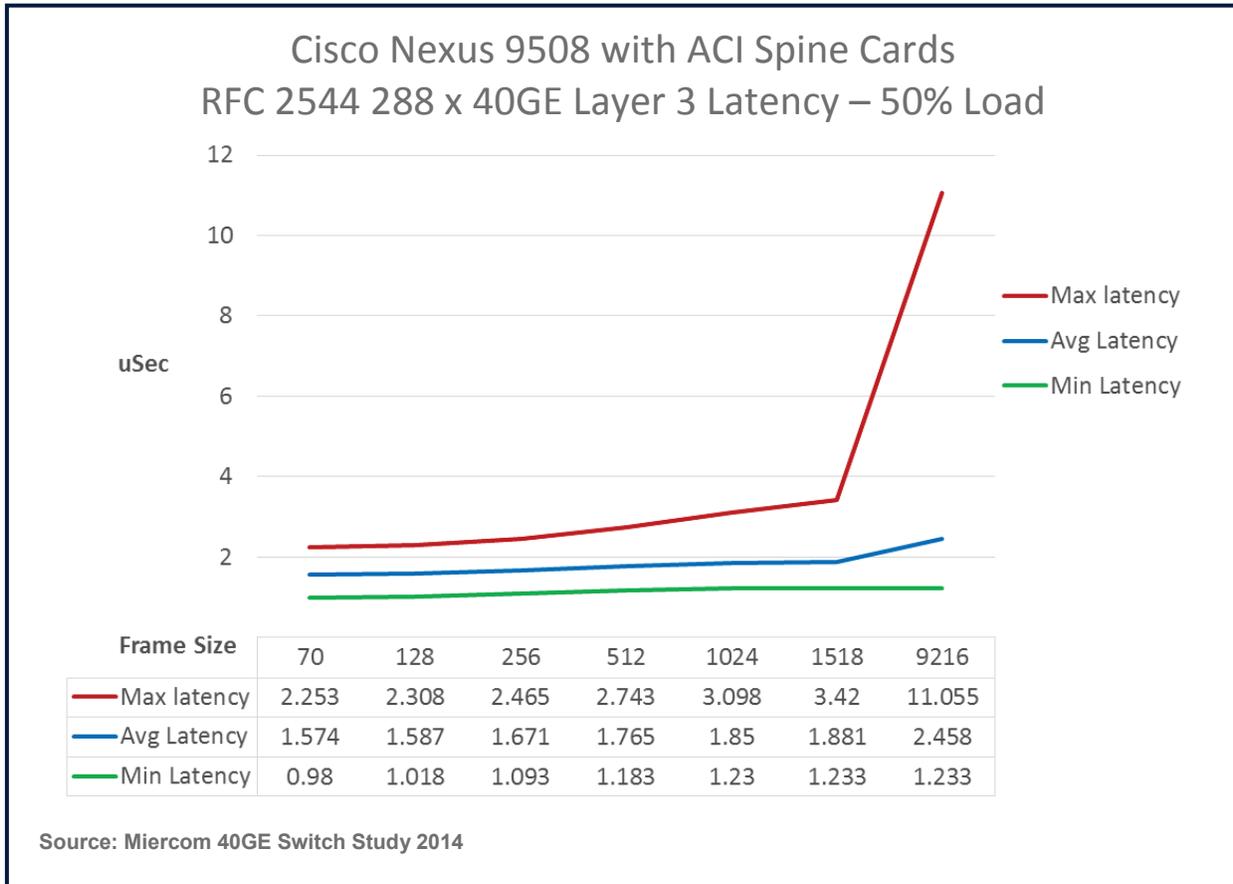


The Cisco Nexus 9508 with eight 36p 40G QSFP+ ACI spine line cards exhibits full line-rate forwarding performance with all 288 of its 40GE ports loaded. Test results shown are for 40GE port-pair testing, cross-fabric configuration. The switch was configured with 288 ports under maximum sustainable load without loss for Layer 3 (IP) unicast traffic for the particular frame size. Testing was conducted in accordance with RFC 2544.

Observations - Throughput was tested for different packet sizes in the range of 70 Bytes to 9,216 Bytes. The Cisco Nexus 9508 could handle full-line-rate traffic on all ports, at all frame sizes, without incurring any loss.

4.3 RFC 2544 40GE Layer 3 Latency – 50% Load

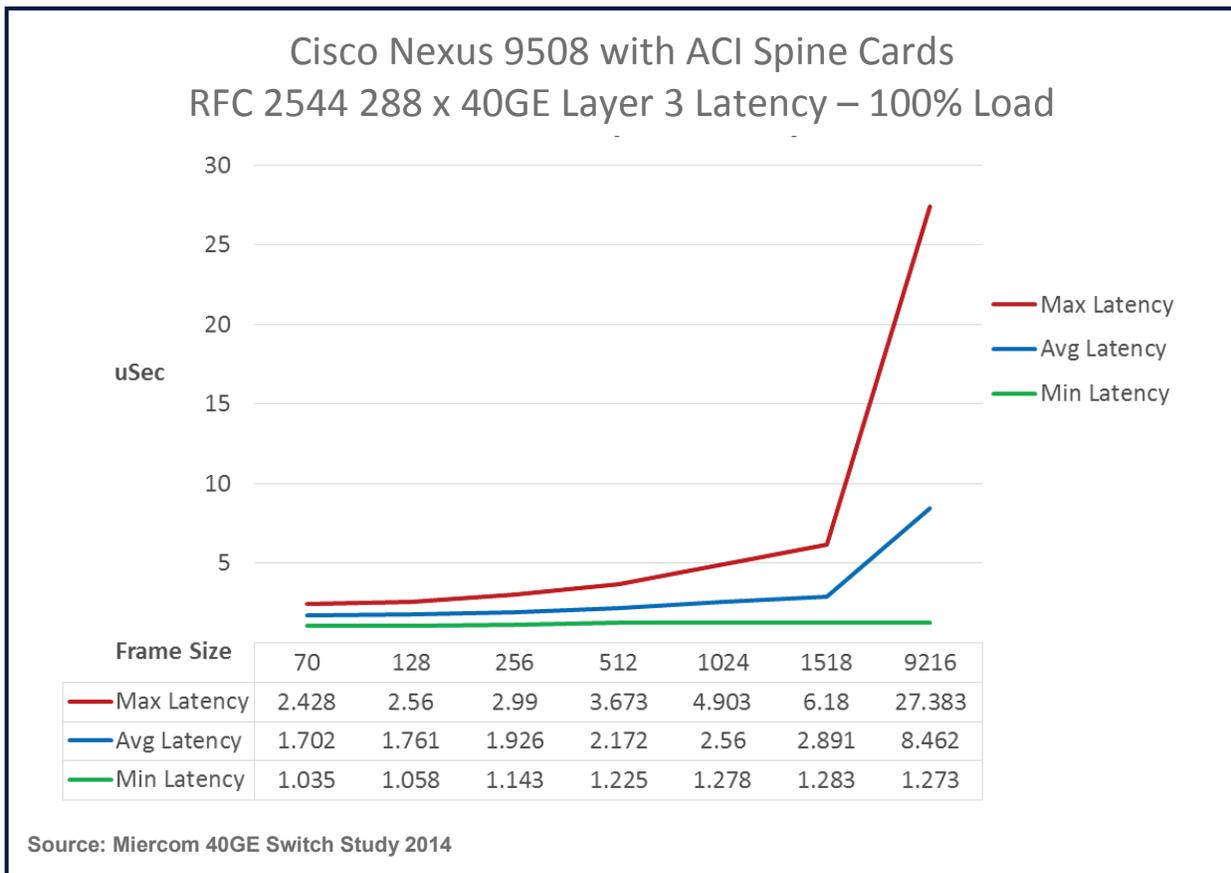
"The Cisco Nexus 9508 with eight 36-port 40GE ACI spine line cards exhibits consistent low latency for all packet sizes tested during RFC 2544-based Latency Tests with 50% Load"



The Cisco Nexus 9508 with eight 36p 40G QSFP+ ACI spine line cards exhibited the latency test results shown above during 40GE Data Center switch testing. The switch was subjected to a 50 percent traffic load of Layer 3 (IP) unicast traffic on all of its 288 ports, for the specified frame size, and exhibited low and consistent latency. From the smallest packet size of 70-Bytes to the largest 9216-Bytes, the average latency ranges from 1.6 to 2.5 microseconds (uSec). Tests were conducted in accordance with RFC 2544.

4.4 RFC 2544 40GE Layer-3 Latency – 100% Load

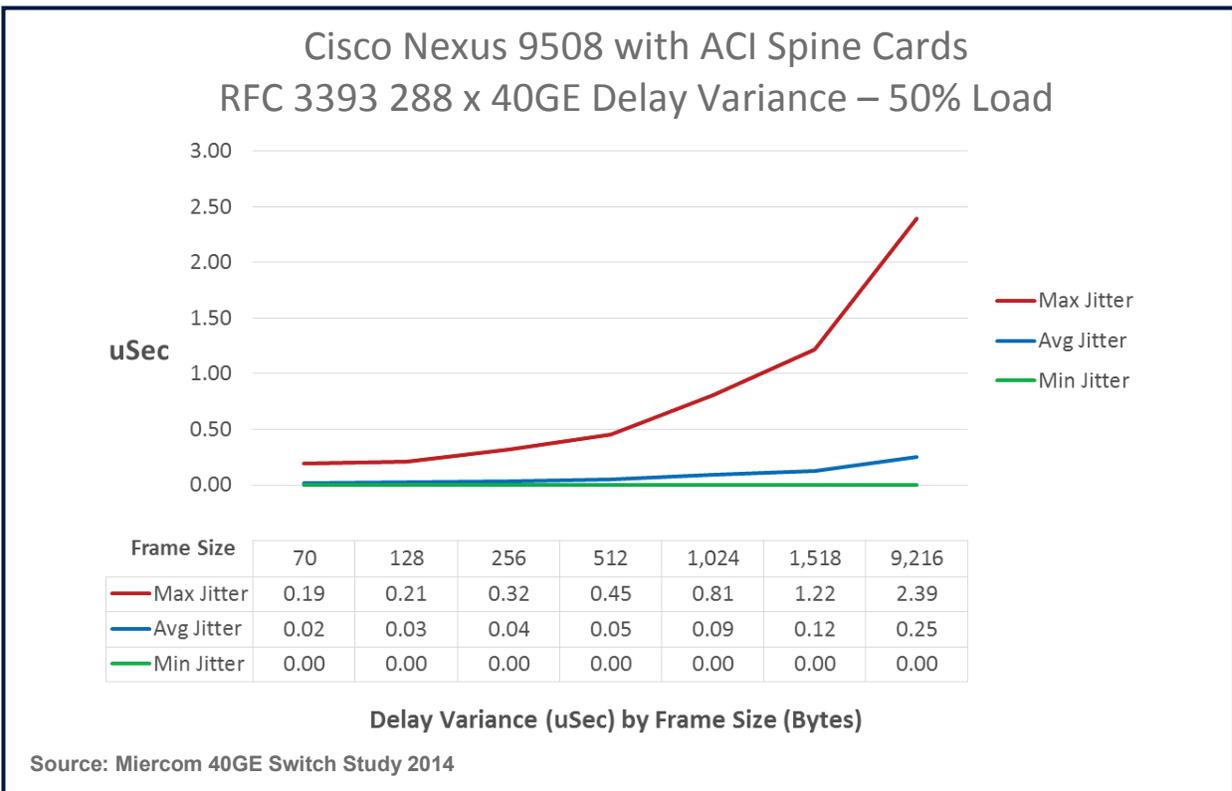
"The Cisco Nexus 9508 with eight ACI spine line cards exhibits consistent low latency for all packet sizes tested during RFC 2544-based Latency Tests, even with the switch fully loaded 100% on all ports."



The Cisco Nexus 9508 with eight 36p 40G QSFP+ ACI spine line cards exhibited the latency results shown above during 40GE Data Center switch testing. The switch was subjected to a 100 percent traffic load of Layer 3 (IP) unicast traffic on all of its 288 ports, for the specified frame size, and exhibited low and consistent latency. From the smallest packet size of 70-Bytes to the largest 9216-Bytes, the average latency ranges from 1.7 to 8.4 microseconds (uSec). Tests were conducted in accordance with RFC 2544.

4.5 RFC 3393-based 40GE Layer 3 Delay Variance - 50% Load

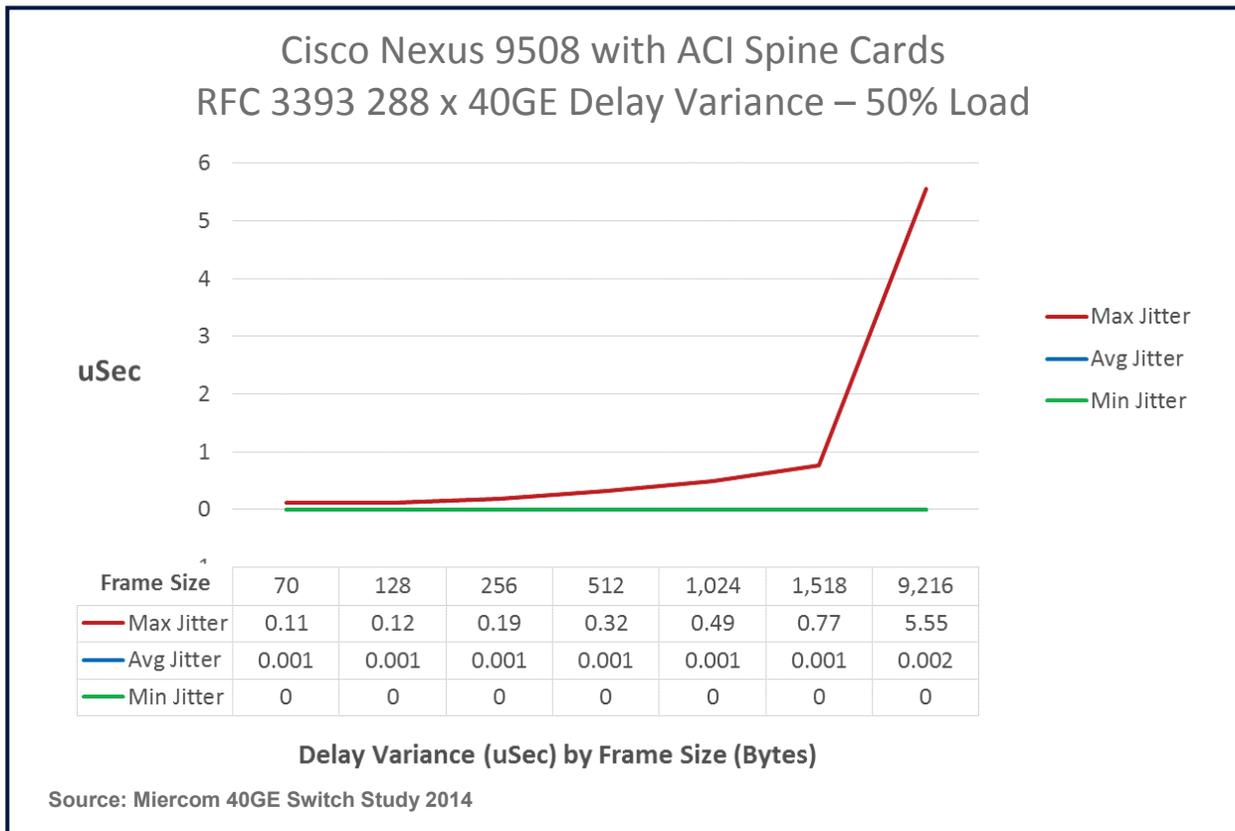
"The Cisco Nexus 9508 with eight ACI spine line cards exhibits very little variance and consistent latency for all packet sizes tested during RFC 3393-based Latency Tests with 50% load."



The Cisco Nexus 9508 exhibited very little variance in latency, also called "jitter" – less than one quarter of a microsecond on average for all packet sizes up to 9,216-byte. The Cisco switch configured with 288 ports was subjected to a 50 percent traffic load. Layer 3 IP unicast traffic was used for the specified frame size. Tests were conducted in accordance with RFC 3393.

4.6 RFC 3393-based 40GE Layer-3 Delay Variance - 100% Load

"The Cisco Nexus 9508 with eight ACI spine line cards exhibits very little variance and very consistent latency for all packet sizes tested during RFC 3393-based Latency Tests, even at 100% load."



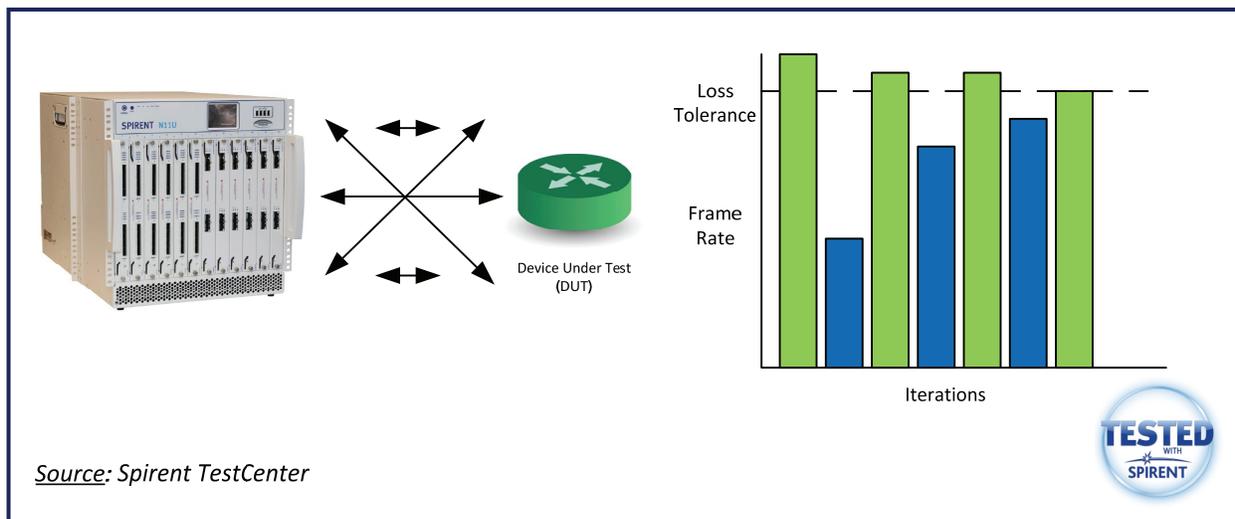
The Cisco Nexus 9508 exhibited very little variance in latency, also called "jitter" –10 nanoseconds or less, on average, for all packet sizes up to 9,216-byte. The Cisco switch configured with 288 ports was subjected to a 100 percent traffic load. Layer 3 IP unicast traffic was used for the specified frame size. Tests were conducted in accordance with RFC 3393.

5 Fully Meshed Throughput and Latency Test (RFC 2889)

Test Description - The Fully Meshed throughput performance test, as described in RFC 2889, determines the total number of frames that the device under test (DUT) can handle when receiving frames on all ports. The test results show the total number of frames transmitted from, and the total number of frames received on, all ports. In addition, the percentage loss of frames for each frame size is also determined.

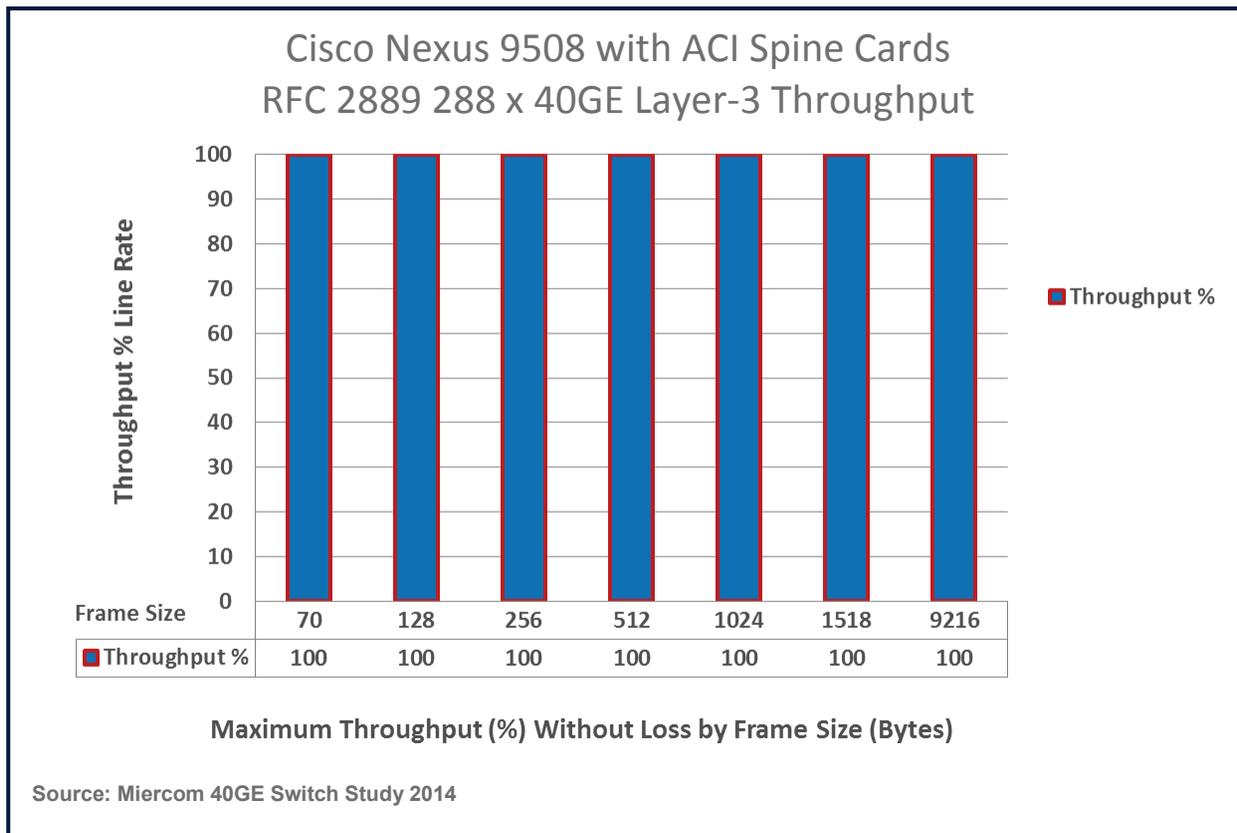
Procedure and Configuration - In accordance with RFC 2889 and best known networking practices, all ports of the Spirent test system are connected, and traffic flows of fixed packet sizes are sent in a mesh-distribution fashion. The device under test is configured for Layer-3 switching (IP routing). The test inherently stresses the switch by sending a "mesh traffic" load distribution, with traffic traversing both the local line card and other line cards, therefore forcing traffic across all the fabric modules. The total number of frames obtained for each frame size for the fully populated switch is recorded. A bidirectional traffic load is used for this test (each port is sending and receiving traffic simultaneously).

5.1 RFC 2889-based Fully Meshed Test System and DUT Configuration



5.2 RFC 2889-based 40GE Layer-3 Full Mesh Throughput

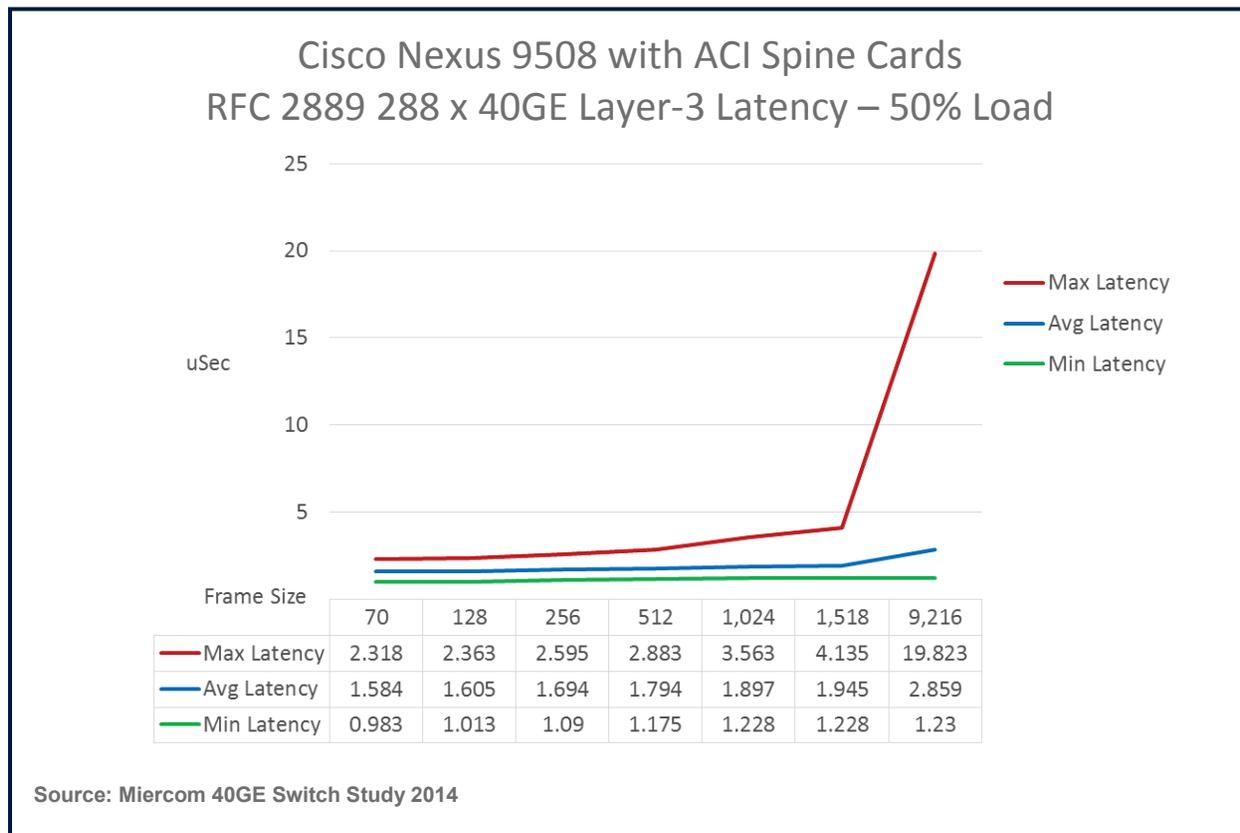
"The Cisco Nexus 9516 with eight ACI spine line cards proved full line-rate mesh throughput performance for all packet sizes tested during RFC 2889-based Throughput Tests without a single packet dropped."



The Cisco Nexus 9508 exhibits full-line-rate traffic handling across 288 x 40GE ports loaded. These results are for testing of 40GE ports in a cross-fabric configuration, with full mesh traffic loads. Tests subjected the switch to the maximum sustainable load that it could maintain without loss of Layer 3 unicast traffic for the specified frame size. Tests were conducted in accordance with RFC 2889.

5.3 RFC 2889-based Full Mesh L3 Latency – 50% Load

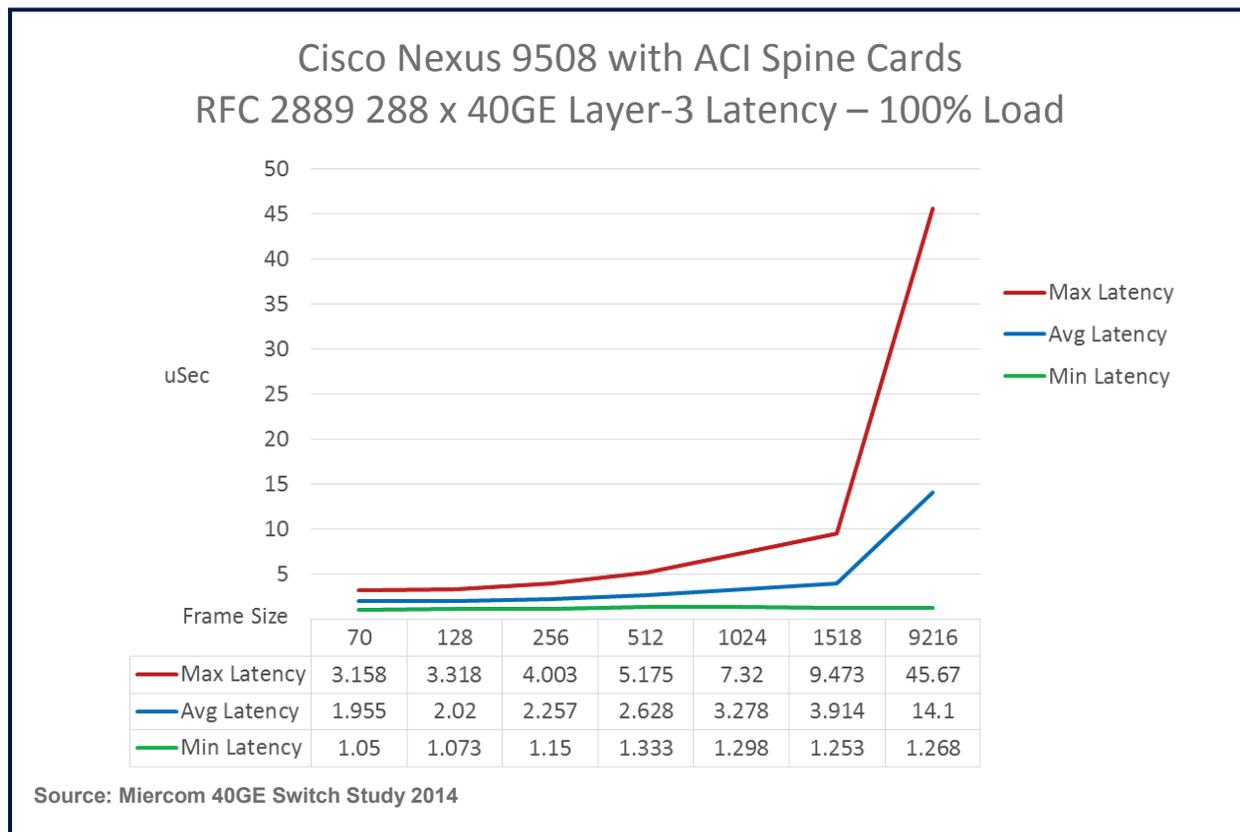
"The Cisco Nexus 9508 with eight ACI spine line cards exhibits consistent low latency for all packet sizes tested during RFC 2889-based Full Mesh Latency Tests with 50% load."



The Cisco Nexus 9508 exhibits consistent low latency for the full range of packet sizes tested across all 288 of its 40GE ports with 50 percent load on all ports. From the smallest packet size of 70-Bytes to the largest 9,216-Bytes, the average latency ranges from 1.6 to 2.9 microseconds. Results are shown for 288 x 40GE full mesh testing at 50 percent traffic load. Tests were conducted in accordance with RFC 2889.

5.4 RFC 2889-based Full Mesh L3 Latency – 100% Load

"The Cisco Nexus 9508 with eight ACI spine line cards exhibits consistent low latency for all packet sizes tested during RFC 2889 Full Mesh Latency Tests, even with the switch loaded 100 percent on all ports."



The Cisco Nexus 9508 exhibits consistent low latency for the full range of packet sizes tested across all 288 of its 40GE ports with 100 percent load on all ports. From the smallest packet size of 70-Bytes to the largest 9216-Bytes, the average latency ranges from 2.0 to 14.1 microseconds. Results are shown for 288 x 40GE full-mesh testing at 100 percent traffic load. Tests were conducted in accordance with RFC 2889.

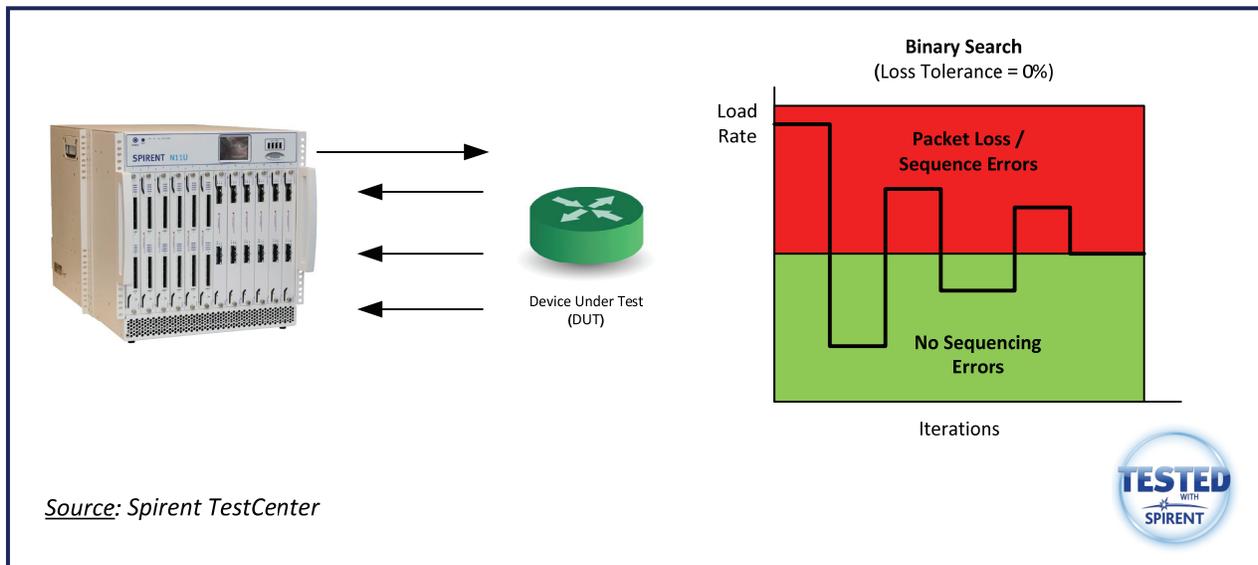
6 IP Multicast Throughput and Latency (RFC 3918)

Description - RFC 3918 describes Multicast Forwarding and Latency test, for measuring the forwarding throughput and the minimum, maximum and average latency of multicast frames. These are frames generated and received by the test and measurement equipment. The multicast frames are sent to "clients" on multiple subnets (ports), which are also configured via the test and measurement system.

The test reveals how much processing overhead the device under test (DUT) requires, on average, to forward multicast frames. In a typical scenario the tester defines the multicast protocol (our testing used IGMPv1 2 3; PIM-SM; SSM), and the number of multicast groups to be sent. Traffic streams are automatically built by the Spirent test system. A combination of throughput, latency, group capacity, frame loss, join delay, and leave delay can be calculated from the results.

This test is used to measure the IP multicast forwarding throughput of the DUT and then calculate the latency of traffic. This is done by examining the "sent" timestamp that the test system places within test frames. When frames are received from the ports, the test system compares this timestamp with the current time and calculates the difference, which is the latency. The test system records the average, minimum and maximum latencies for each multicast group, in microseconds (μ s).

6.1 RFC 3918-based IP Multicast Test System and DUT Configuration

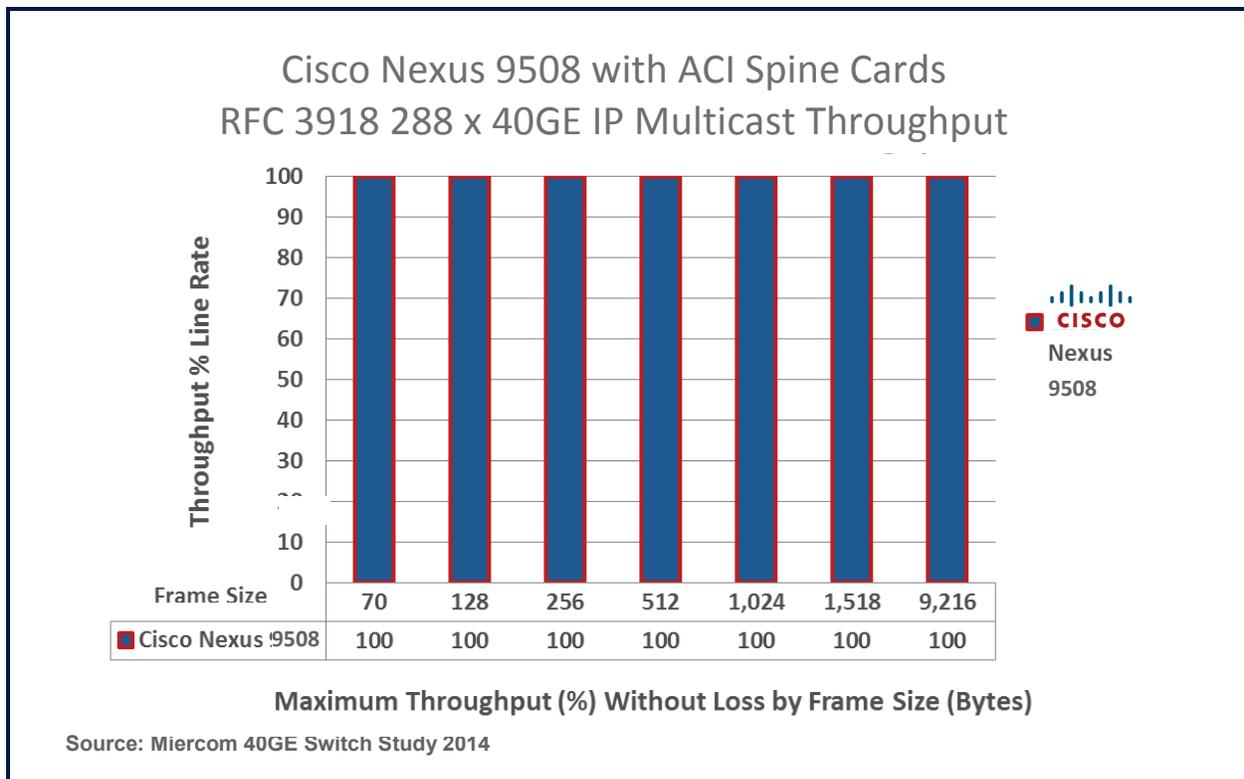


6.2 RFC 3918-based 40GE IP Multicast Throughput

IP multicast is a process and protocol for sending the same IP packet stream to a group of interested receivers. It is an efficient way to achieve one-to-many or many-to-many communication between affiliated groups of IP devices. IP multicast is heavily used by financial trading applications, where performance is critical to those applications.

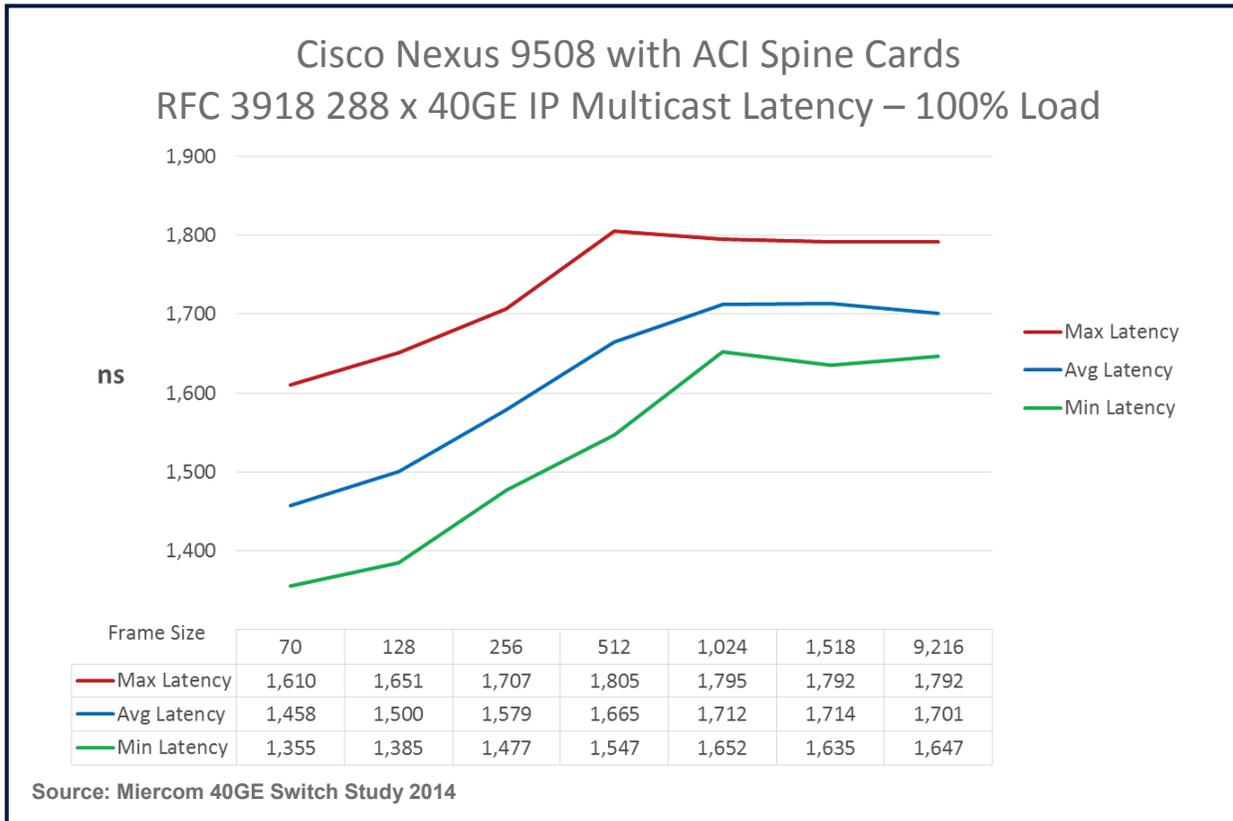
To understand the IP-multicast-handling capability of the Nexus 9508 we performed RFC 3918-based IP multicast throughput tests, where IP multicast traffic was sent from one input port to all the rest of the ports on the switch. So there were 287 receivers on the Nexus 9508 (traffic sent in on one port and delivered via IP multicast to all the other ports). The graph below shows the throughput testing results.

"The Cisco Nexus 9508 with eight ACI spine line cards delivered line-rate performance for all packet sizes during RFC 3918-based IP Multicast Tests."



The Cisco Nexus 9508 with eight 36p 40G QSFP+ ACI spine line cards exhibits full line-rate multicast handling across all 288 of its 40GE ports, from a fully loaded input traffic stream. Tests were conducted in accordance with RFC 3918.

6.3 RFC 3918-based 40GE IP Multicast Latency – 100% Load



The Cisco Nexus 9508 exhibits low latency for IP Multicast traffic across all 288 of its 40GE ports, under 100 percent traffic load. From the smallest packet size of 70-Bytes to the largest 9216-Bytes, the average latency ranges from 1.5 to 1.7 microseconds (uSec). The test results shown are for 288 x 40GE IP Multicast testing. The Cisco switch was configured with 288 ports under 100 percent load without loss for Layer 3 IP Multicast traffic for the specified frame size. Tests were conducted in accordance with RFC 3918.

7 Mean Time between Failure (MTBF) Analysis

Miercom engineers audited the analysis for the MTBF calculations of the line cards for the Cisco Nexus 9508 Switch relative to competitive offerings. Overall, the life expectancy of the line cards before a failure, based on subcomponent analysis failure statistics, projects a calculated expected lifetime for a fully populated Nexus 9508 to be 55,996 hours. This calculation is based on the analysis of the subcomponent count shown in the table on the next page.

The analysis covered 3 different systems:

- 1) 8 slot chassis based on a Broadcom Dune reference design. Line cards have 36p 40GE QSFP ports, but only 30 of those ports are line rate for all packet sizes.
- 2) 8 slot chassis using Broadcom Trident 2 ASICs on the fabric modules.
 - a) Line cards have 36p 40G QSFP ports using 3 Broadcom Trident 2 ASICs.
 - b) Line cards have 36p 40G QSFP ports using 2 Cisco ASICs.

Note: Calculated MTBF numbers based on a detailed subcomponent analysis cannot be compared with “observed” MTBF numbers based on failure rates seen in an installed base.

System MTBF: The MTBF analysis in this report assumes a fully loaded 8-slot chassis, yielding 288 x 40GE ports.

System 1: Broadcom Dune based solution - 19,981 hours

System 2a: Broadcom T2 based solution - 45,162 hours

System 2b: Cisco ASIC line card and Broadcom fabric modules - 55,996 hours

System Design

All designs are modular chassis (8 line-card slots) with 6 fabric modules.

Component count per line card

- | | |
|--|---------|
| 1) Broadcom Dune line card: 36 ports (30 ports line rate) of 40GE: | 6 ASICs |
| 2a) Broadcom T2 line card: 36 ports (36 ports line rate) of 40GE: | 3 ASICs |
| 2b) Cisco Custom ASIC: 36 ports (36 ports line rate) of 40GE: | 2 ASICs |

The large number of components and IC-Memory adversely impacts the MTBF for systems using Type 1 line cards.

Also, ASICs leveraging the latest technology allow for higher density of 40GE ports, and in turn for designs with less components, yielding a higher system MTBF.

Current Broadcom ASICs are based on 40-nm ASIC technology. Cisco ASICs used in the Nexus 9500 Series are based on 28-nm ASIC technology. This allows for more built-in functionality in the ASIC while requiring less power and less external components.

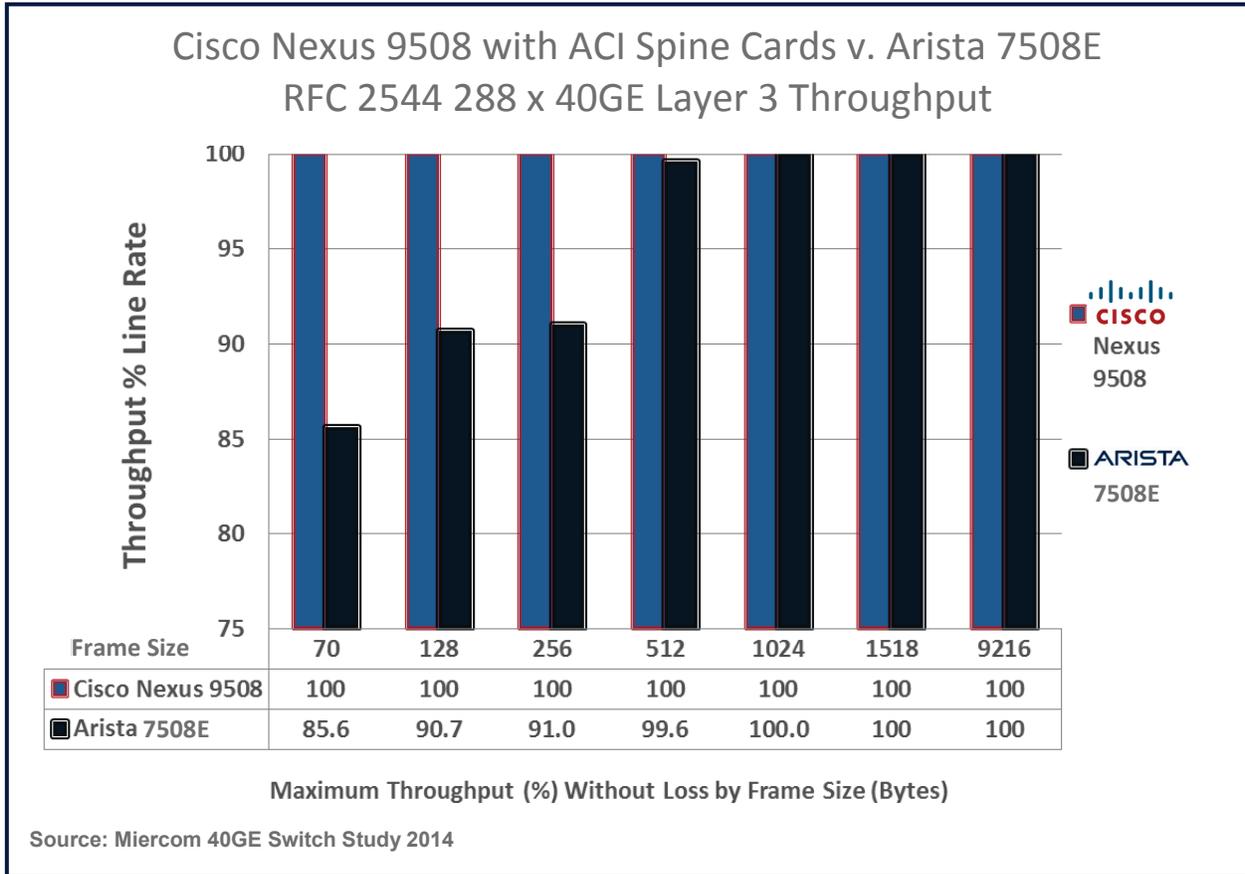
7.1 Component Count Used for MTBF Calculations

Number of Components per Line Card (Few components is better)			
Components	Dune	Trident II	Cisco Custom ASIC
Capacitors	7,990	4,083	2,640
Connectors	18	22	24
Diodes	85	96	97
IC - MPU	3	1	1
IC - Linear	84	62	54
IC-Logic	48	12	18
IC-Mem	103	15	11
Inductros	454	332	157
Misc	39	29	29
Optics	2	2	2
Resistors	4,218	491	524
Transistors	25	19	17
TOTAL	13,067	5,164	3,574

The 36p 40G QSFP+ ACI spine line card designed for the Cisco Nexus 9516 and 9508 has fewer components that could fail. Less subcomponents in design lead to higher product reliability. The components listed here show the component count for line cards using the Cisco custom ASIC compared to competitive offerings. The sum of the individual component life expectations were added to calculate the total projected MTBF in hours.

8 Competitive Analysis

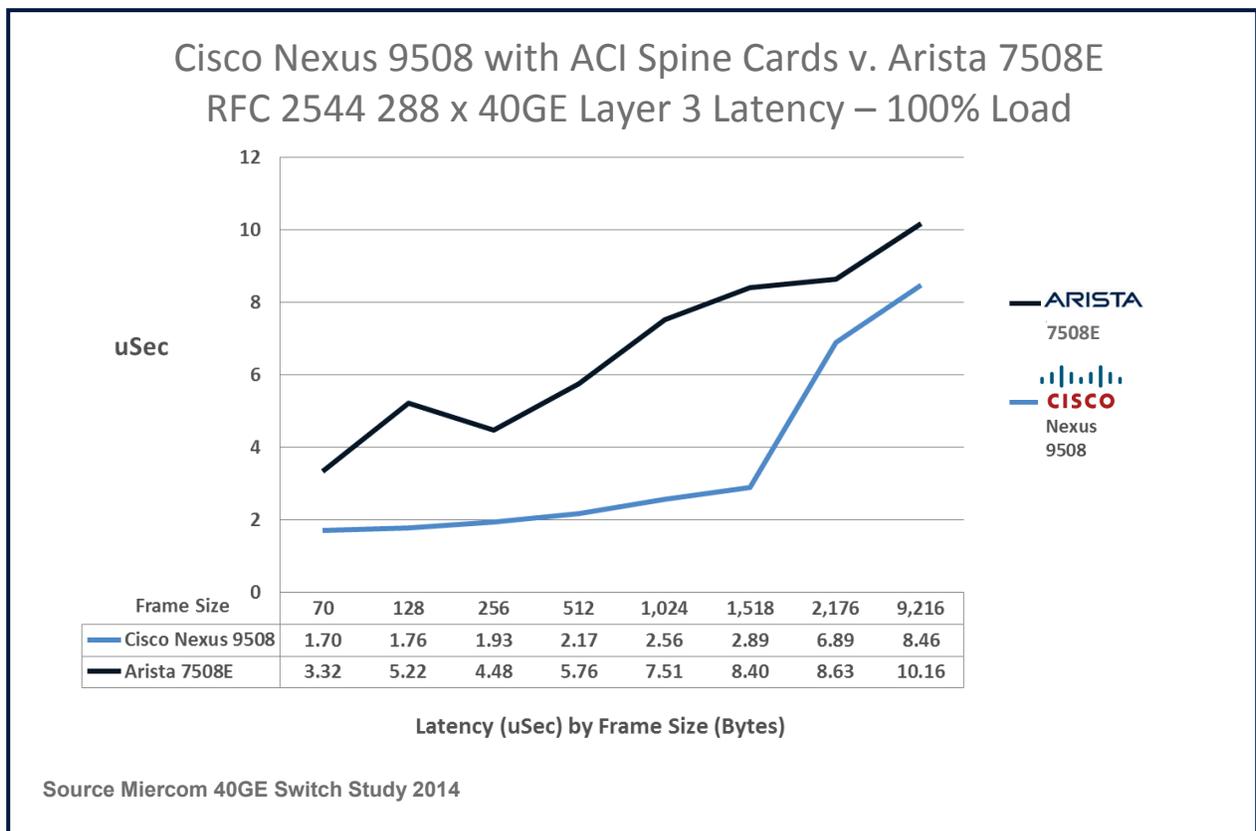
8.1 RFC 2544-based 40GE Layer 3 Throughput Test – 100% Load



Cisco Nexus 9508 with eight 36p 40G QSFP+ ACI spine line cards exhibits full line-rate traffic handling ability across all 288 x 40GE ports loaded. The Arista 7508E switch drops traffic for smaller packet sizes with 288 x 40GE ports loaded. Test results for 40GE cross fabric configuration, port pair testing are shown above. Arista switch was configured with 288 ports and the Cisco switch with 288 ports under maximum sustainable load without loss for Layer 3 unicast traffic at the specified frame size. Tests were conducted in accordance with RFC 2544.

8.3 RFC 2544-based 40GE Layer 3 Latency Test – 100% Load

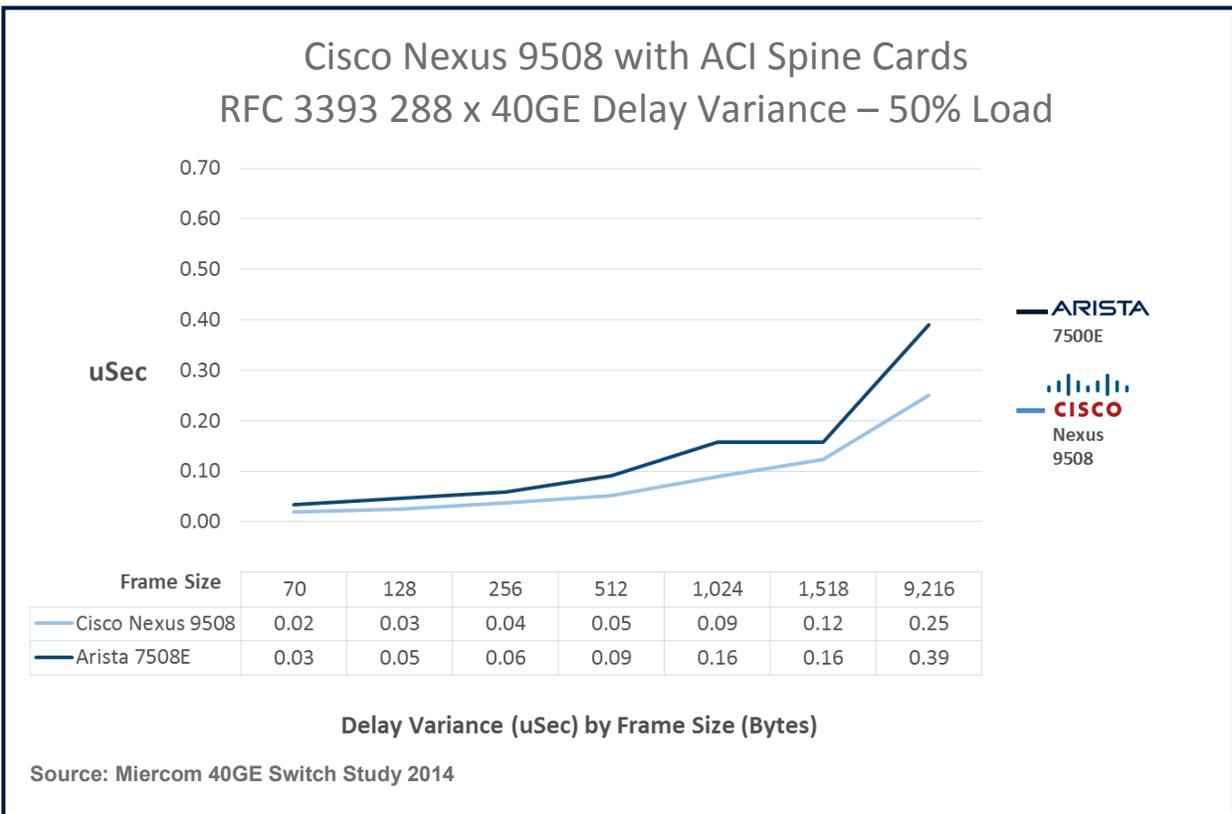
"Even under the most extreme conditions of traffic load, the Cisco Nexus 9508 with eight ACI spine line cards exhibits consistent low latency for all frame sizes tested. Arista 7508E exhibits higher latency and significant frame loss at smallest frame sizes."



The Cisco Nexus 9508 with eight 36p 40G QSFP+ ACI spine line cards and Arista 7508E latency test results at 99% traffic utilization for 40GE Data Center Switch Testing are shown above. The Cisco Nexus 9508 exhibited better (lower) and more consistent latency compared to the Arista 7508E for all packet sizes tested. Tests were conducted in accordance with RFC 2544. Latency measurements were conducted using FIFO (First In, First Out) which is the appropriate mode for testing low latency switches in "Cut-Through" mode of operation.

8.4 RFC 3393-based 40GE Layer 3 Delay Variance - 50% Load

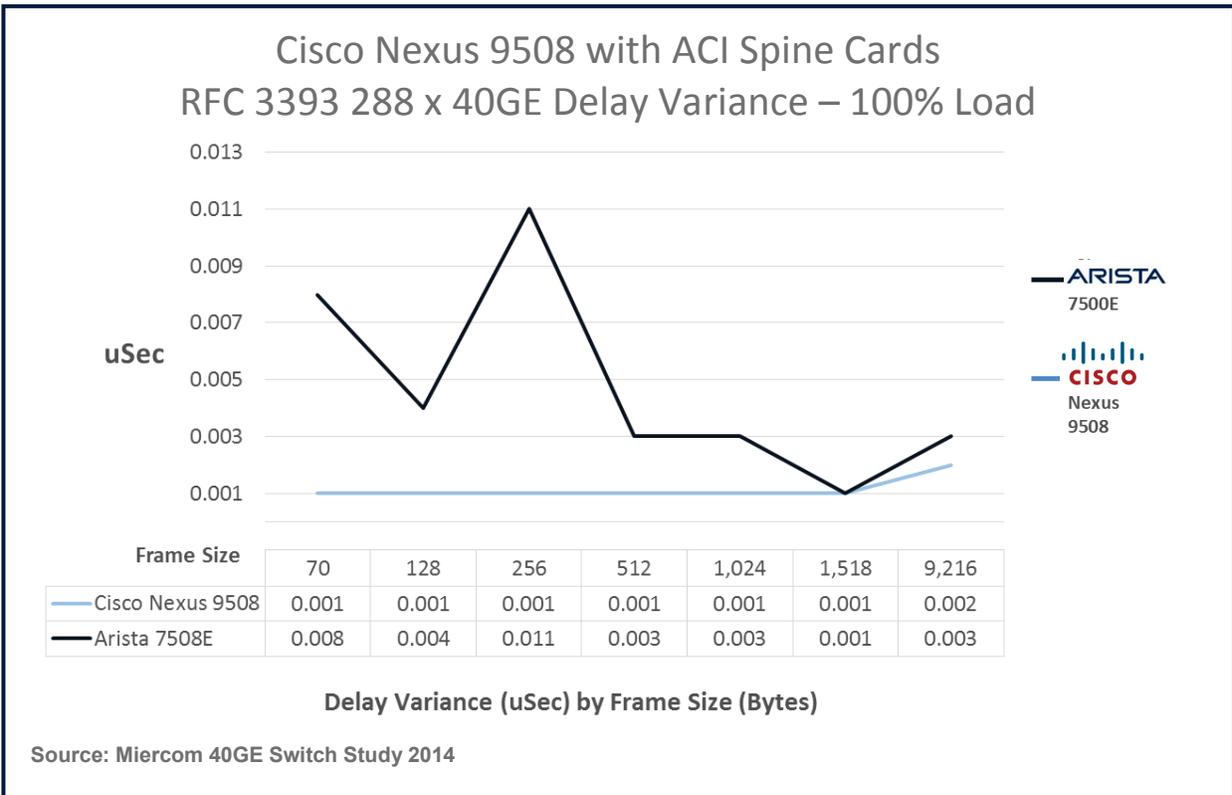
"The Cisco Nexus 9508 with eight ACI spine line cards exhibits very little variance and consistent latency for all packet sizes tested during RFC 3393-based Latency Tests with 50% load."



The Cisco Nexus 9508 exhibited very little variance in latency, also called "jitter" – less than 1 microsecond on average for all packet sizes up to 9,216-Bytes. The Cisco switch configured with 288 ports was subjected to a 50 percent traffic load. Layer 3 IP Unicast traffic was used for the specified frame size. Tests were conducted in accordance with RFC 3393.

8.5 RFC 3393-based 40GE Layer 3 Delay Variance - 100% Load

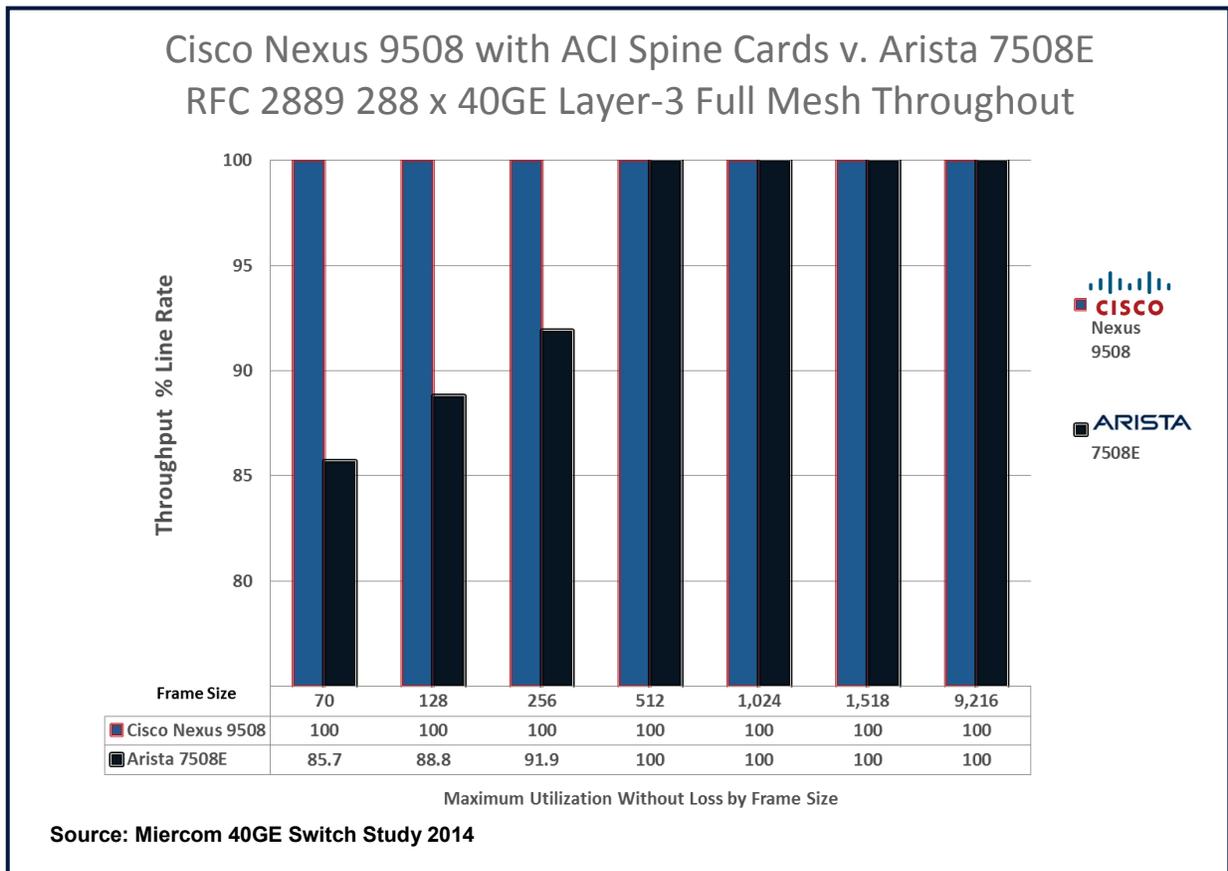
"The Cisco Nexus 9508 with eight ACI spine line cards exhibits very little variance and very consistent latency for all packet sizes tested during RFC 3393-based Latency Tests, even at 100% load."



The Cisco Nexus 9508 exhibited very little variance in latency, also called "jitter" – 10 nanoseconds or less, on average, for all packet sizes up to 9,216-Bytes. The Cisco switch configured with 288 ports was subjected to a 100 percent traffic load. Layer 3 IP Unicast traffic was used for the specified frame size. Tests were conducted in accordance with RFC 3393.

8.6 RFC 2889-based 40GE Layer 3 Full Mesh Throughput

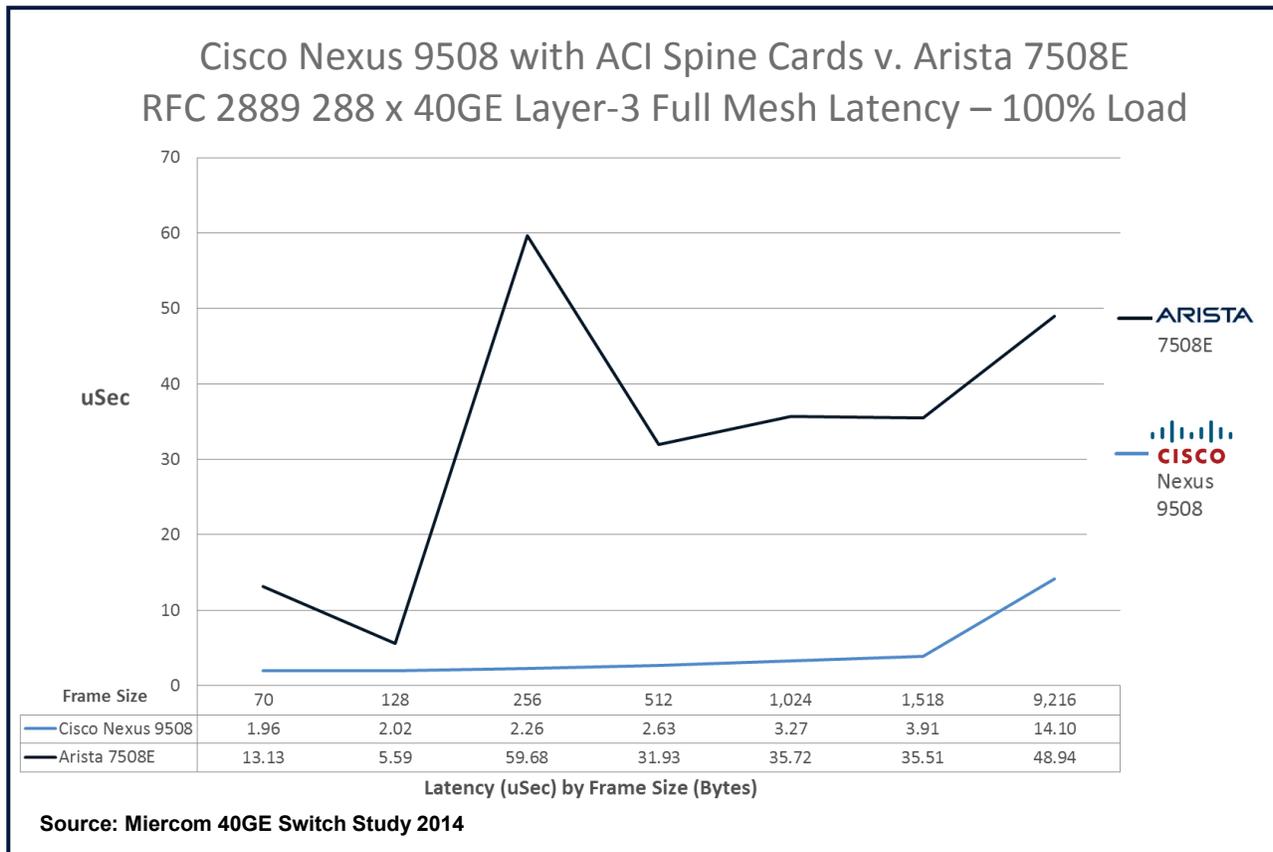
"The Cisco Nexus 9516 with eight ACI spine line cards proved full line-rate mesh throughput performance for all packet sizes tested during RFC 2889-based Throughput Tests without a single packet dropped."



The Cisco Nexus 9508 with eight 36p 40G QSFP+ ACI spine line cards and Arista 7508E full mesh throughput test results are shown above. The Cisco Nexus 9508 exhibits full line-rate traffic handling across 288 x 40GE ports loaded. The Arista product demonstrated lower throughput without loss with its 288 ports tested for smaller frame sizes. These results are for testing of 40GE ports in a cross-fabric configuration, with full mesh traffic loads. Tests subjected the switch to the maximum sustainable load that it could maintain without loss of Layer 3 Unicast traffic for the specified frame size. Tests were conducted in accordance with RFC 2889 using a Layer 3 traffic configuration.

8.7 RFC 2889-based Full Mesh L3 Latency – 100% Load

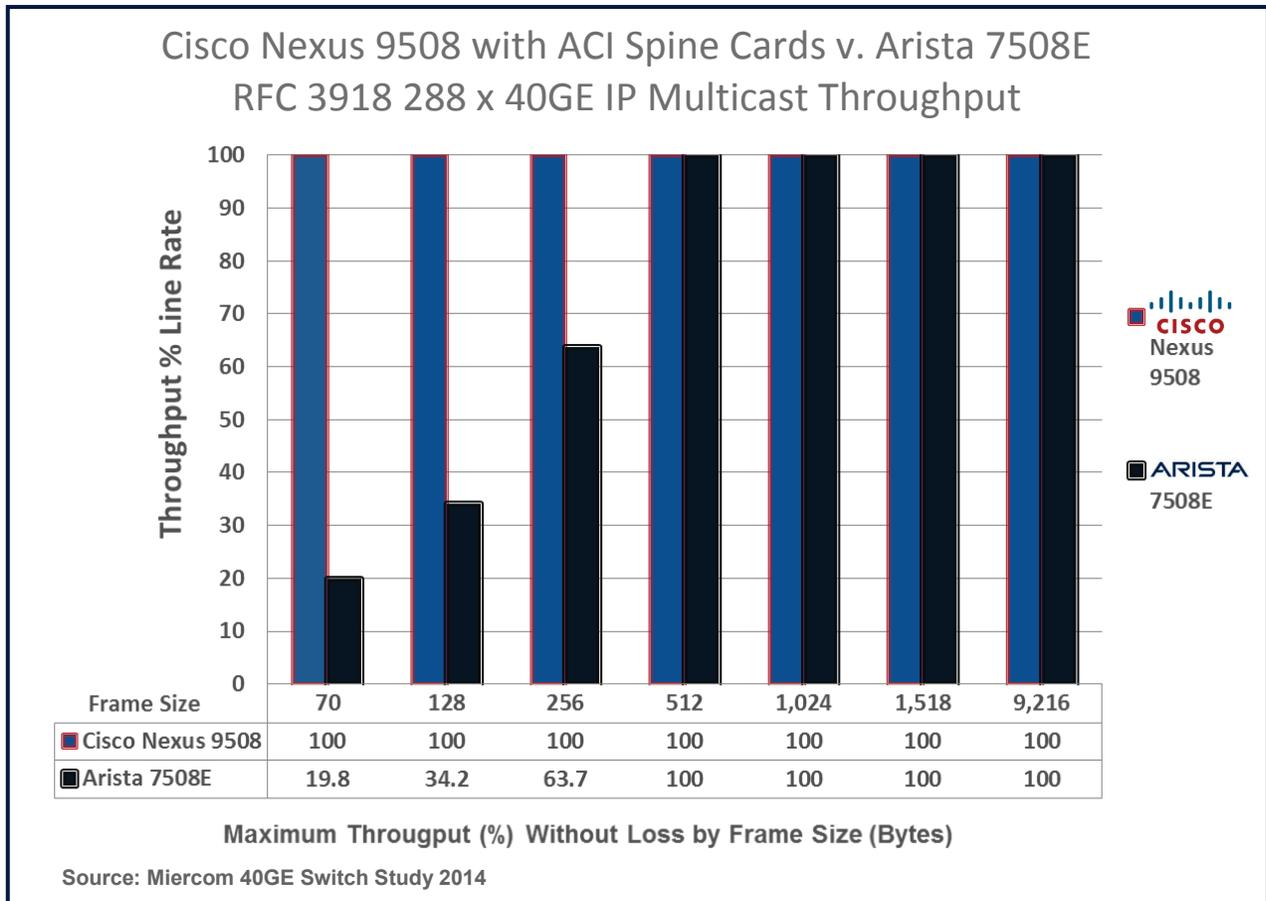
"The Cisco Nexus 9508 with eight ACI spine line cards exhibits consistent low latency for all packet sizes tested during RFC 2889 Full Mesh Latency Tests, even with the switch loaded 100 percent on all ports."



Full mesh latency test results are shown above for the Cisco Nexus 9508 with eight 36p 40G QSFP+ ACI spine line cards and the Arista 7508E. From the smallest packet size of 70-Bytes to the largest 9216-Bytes, the average latency for the Cisco Nexus 9508 is consistently lower than that for the Arista 7508E. Results are shown for the Cisco Nexus 9508 tested with 288 x 40GE ports full mesh testing at 100 percent traffic load. The Arista 7508 was tested with 288 x 40GE ports full mesh testing at 100 percent traffic load. Tests were conducted in accordance with RFC 2889 using a Layer 3 traffic configuration. Latency measurements were conducted using FIFO (First In, First Out) which is the appropriate mode for testing low latency switches in "Cut-Through" mode of operation.

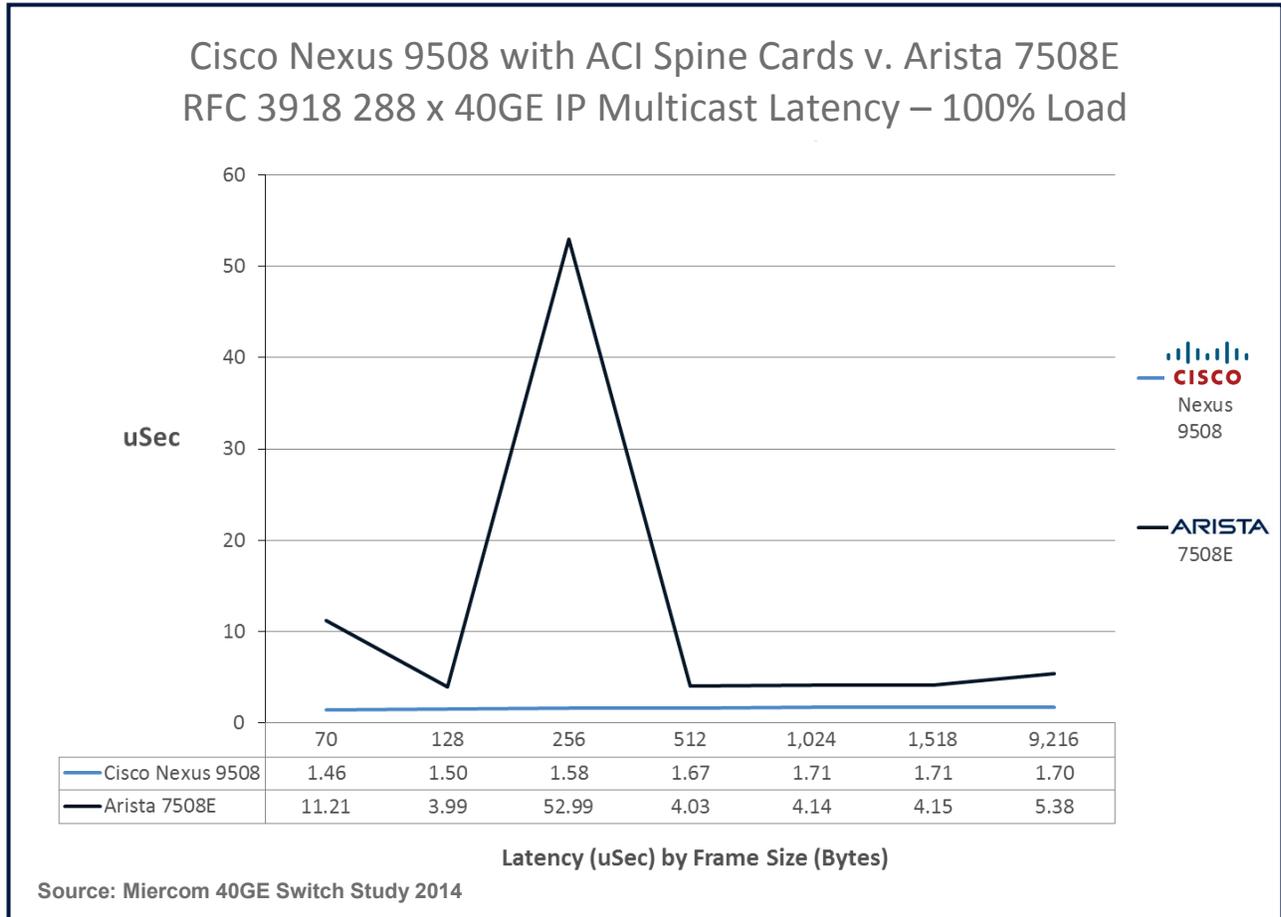
8.8 RFC 3918-based 40GE IP Multicast Throughput

"The Cisco Nexus 9508 with eight ACI spine line cards delivered line-rate performance for all packet sizes during RFC 3918-based IP Multicast Tests."



The Cisco Nexus 9508 with eight 36p 40G QSFP+ ACI spine line cards exhibits full line-rate multicast handling across all 288 of its 40GE ports, from a fully loaded input traffic stream. The Arista 7508E demonstrated less than line rate performance for its 288 ports tested with IP multicast traffic with frames smaller than 512 bytes. Tests were conducted in accordance with RFC 3918 in an accumulated test mode, one multicast group with a single source transmitting to multiple destinations.

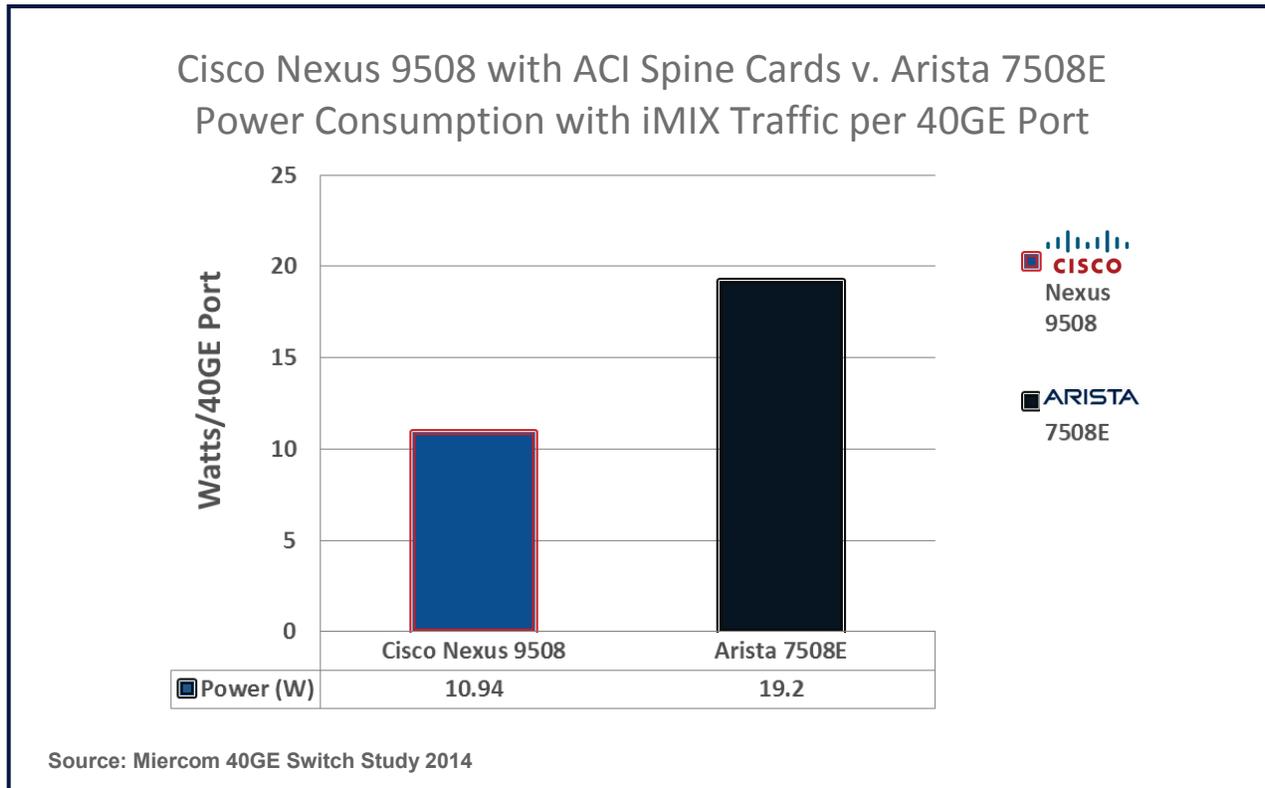
8.9 RFC 3918-based 40GE IP-Multicast Latency – 100% Load



The Cisco Nexus 9508 with eight 36p 40G QSFP+ ACI spine line cards exhibits low latency for IP Multicast traffic across all 288 of its 40GE ports, under 100 percent traffic load. From the smallest packet size of 70-Bytes to the largest 9216-Bytes, the average latency for the Cisco Nexus 9508 is consistently lower than that for the Arista 7508E. Results are shown for the Cisco Nexus 9508 tested with 288 x 40GE ports at 100 percent traffic load without loss for IP Multicast traffic at the specified frame size. The Arista 7508E was tested with 288 x 40GE ports full mesh testing at 100 percent traffic load. Tests were conducted in accordance with RFC 3918. Latency measurements were conducted using FIFO (First In, First Out) which is the appropriate mode for testing low latency switches in “Cut-Through” mode of operation.

8.10 Data Center Switch Power Consumption

Power consumption was measured for the Cisco Nexus 9508 and then calculated on a per-port basis for the switch, configured in a full 288 x 40GE port configuration. Testing was conducted at full line-rate, with bidirectional traffic using a Layer 3 iMIX traffic load.



The Cisco Nexus 9508 with spine line cards exhibits an average of 10.94 watt/40GE port power consumption. This is 43% less power per 40GE port than Arista 7508E with its 36-port 40GE line card. Calculations and power measurements are based on the switch configured with 288 x 40GE ports and full bidirectional traffic loads of Layer 3 iMIX traffic applied to all ports.

9 About the Miercom 40GE Switch Industry Study

This report was sponsored by Cisco Systems, Inc. The data was obtained completely and independently as part of the Miercom's 40GE Data Center Switch Industry Study. The study is an ongoing endeavor in which all vendors have equal opportunity to participate and contribute to the most comprehensive and challenging test program in the industry.

All vendors with products featured in this report were afforded the opportunity before, during, and after testing was complete to comment on the testing results and demonstrate their product's performance. Any vendor with a product tested by Miercom in one of our published studies that disagrees with our findings is extended an opportunity to retest and provide a demonstration of their product's performance (at no charge for the testing.)

10 About Miercom

Miercom has published hundreds of network-product-comparison analyses in leading trade periodicals including *Network World*, *Business Communications Review - NoJitter*, *Communications News*, *xchange*, *Internet Telephony* and other leading publications. Miercom's reputation as the leading, independent product test center is undisputed.

Miercom's private test services include competitive product analyses, as well as individual product evaluations. Miercom features comprehensive certification and test programs including: Certified Interoperable, Certified Reliable, Certified Secure and Certified Green. Products may also be evaluated under the Performance Verified program, the industry's most thorough and trusted assessment for product usability and performance.

11 Use of This Report

Every effort was made to ensure the accuracy of the data contained in this report but errors and/or oversights can occur. The information documented in this report 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 were reasonably verified by Miercom but beyond our control to verify to 100 percent certainty.

This document is provided "as is," by Miercom and 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 in this report.

No part of any document may be reproduced, in whole or in part, without the specific written permission of Miercom or Cisco Systems, 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.