

# Cisco Nexus 6004 Switch 40 Gigabit Ethernet Performance Validation

White Paper

June 2013

## What You Will Learn

Cisco, in partnership with Spirent Communications, conducted performance benchmark tests for the Cisco Nexus® 6004 Switch. The tests evaluated switch performance with various traffic profiles and Cisco® NX-OS Software features up to and including the most stressful testing conditions. This document presents the test results.

## Overview

The Cisco Nexus 6004 Switch is a high-performance, high-density, low-latency 10 and 40 Gigabit Ethernet and Fibre Channel over Ethernet (FCoE) switch. This compact four-rack-unit (4RU) 10- and 40-Gbps switch provides line-rate Layer 2 and 3 switching. It runs the industry-leading Cisco NX-OS Software operating system, providing customers with features and capabilities that are widely deployed worldwide.

The Cisco Nexus 6004 offers 48 x 40-Gbps ports on the base of the chassis and 4 line-card expansion module (LEM) slots; 12 x 40-Gbps port LEMs can be plugged into each slot, for a system total of 96 x 40-Gbps ports in an Enhanced Quad Small Form-Factor Pluggable (QSFP+) form factor.

Figure 1 shows the Cisco Nexus 6004.

**Figure 1:** Cisco Nexus 6004 with Four LEMs



The increasing need for high-bandwidth data transfer rates in the data center challenges organizations to increase the available bandwidth to both the servers and the backplane. One approach has been to use link aggregation and PortChannels consisting of multiple 10-Gbps links. However, this approach may not achieve the desired throughput because of hashing mechanism imbalances and the need to support higher bandwidth flows. The solution is faster links using native 40 Gigabit Ethernet.

Native implementation of 40 Gigabit Ethernet is different from link aggregation, even though both look alike physically. Because of hashing issues and because only one link can be used per flow, full bandwidth utilization is difficult to achieve with link aggregation. With native 40 Gigabit Ethernet, because the traffic is split across multiple lanes at a very low level, full bandwidth is utilized.

The goal of this document is to show the native 40 Gigabit Ethernet performance of the Cisco Nexus 6004 in different traffic scenarios.

To fully validate the Cisco Nexus 6004, rigorous testing was performed. IETF RFC 2544<sup>1</sup>, 2889<sup>2</sup>, and 3918<sup>3</sup> are widely accepted in the industry as standard benchmarking methodologies to evaluate switch performance by measuring throughput, packet forwarding rate, and forwarding latency at various conditions. Spirent TestCenter 4.16 has built-in RFC 2544, 2889, and 3918 test suites, which were used in the tests.

The testing was performed for the full 40-Gbps configuration (96 x 40-Gbps ports).

<sup>1</sup> IETF Network Working Group: Benchmarking Methodology for Network Interconnect Devices

<sup>2</sup> IETF Network Working Group: Benchmarking Methodology for LAN Switching Devices

<sup>3</sup> IETF Network Working Group: Methodology for IP Multicast Benchmarking

Highlights of the Cisco Nexus 6004 performance test results include:

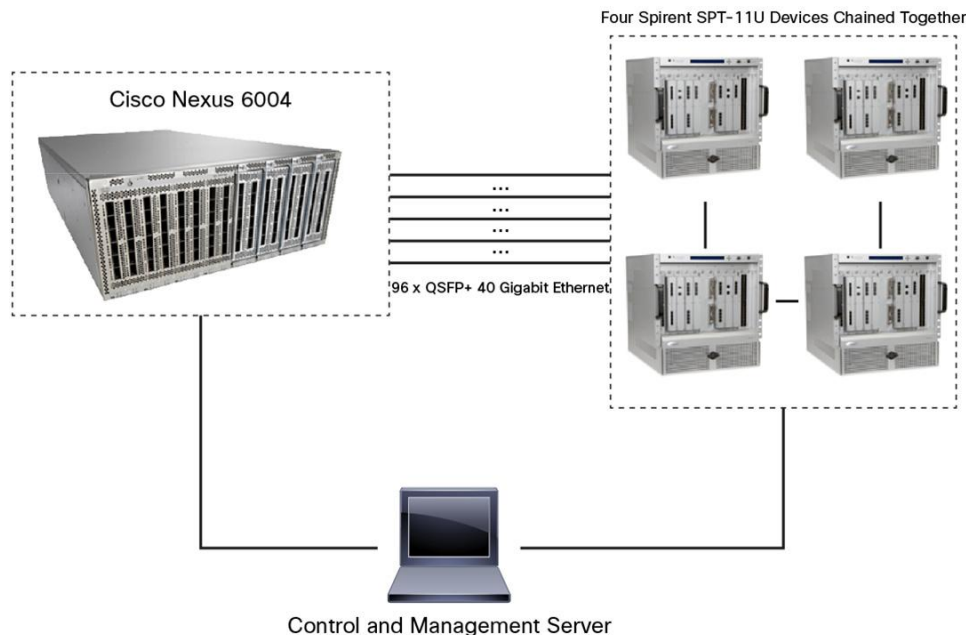
- 100 percent throughput without any packet loss
- Average latency of 1 microsecond
- Identical latency characteristics for Layer 2 and 3 unicast, broadcast, and multicast traffic
- Line-rate Switched-Port Analyzer (SPAN) replication capability

Spirent configured and validated the Spirent TestCenter configurations. Cisco provided the switch configuration through the standard Cisco NX-OS command-line interface (CLI) during the tests. All test results are based on these validated configurations.

## Test Bed

Figure 2 provides an overview of the test bed.

**Figure 2:** 96 x 40-Gbps Test Bed Configuration



## How Testing Was Performed

The test topology and test bed include a Cisco Nexus 6004 Switch, Spirent SPT-11U devices with 96 x 40 Gigabit Ethernet ports, and a control and management server on an out-of-band network. Table 1 shows the software matrix for the test configuration.

**Table 1:** Test Configuration Software Matrix

Device	Software
1 Cisco Nexus 6004 Switch with 96 x 40 Gigabit Ethernet interfaces	Cisco NX-OS Software 6.0(2)N1(1)
4 Spirent SPT-11U chassis	Spirent TestCenter 4.16

Table 2 shows the hardware matrix for the test configuration.

**Table 2:** Test Configuration: 40-Gbps Hardware Matrix

Device	Software	Notes
<b>24 Spirent TestCenter modules, each with 4 x 40 Gigabit Ethernet ports</b>	Spirent TestCenter 4.16	Total of 96 x 40 Gigabit Ethernet ports
<b>96 x 5m 40-Gbps fiber cables</b>	-	

The Spirent TestCenter RFC benchmark test suites were used to test the performance of the Cisco Nexus 6004.

First, the Cisco Nexus 6004 and Spirent TestCenter were configured to run RFC 2544 test suites to measure switch performance, including throughput and latency, for unicast traffic in a port-pair topology for both Layer 2 and Layer 3 traffic.

Then the multicast performance was measured with RFC 3918 for both Layer 2 and Layer 3 traffic. Again, the throughput and latency were measured in each test.

Finally, the SPAN performance was measured, as well as SPAN and multicast simultaneous replication.

The throughput and latency tests were conducted to learn the following information:

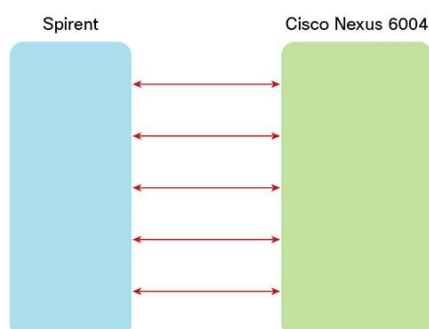
- Throughput is defined as the maximum forwarding rate that the switch can achieve without losing a single packet. The throughput test measures the number of packets that the switch can move at any given time internally: typically, 1 second. For example, 100 percent line rate for a 10 Gigabit Ethernet interface can move 14 million packets per second (mpps).
- Switch forwarding latency is the amount of time it takes for the switch to move a packet from its incoming port to its outgoing port. Essentially, the latency test measures how quickly a switch can forward a packet. For cut-through switches such as the Cisco Nexus 6004, this test uses the First In, First Out (FIFO) method, subtracting the time that the first bit exits the switch from the time that the first bit entered the switch.

Packet-size, traffic-load, and traffic-pattern metrics were tested in all test cases to help ensure adequate test coverage.

- Packet size (bytes): 64, 96, 128, 256, 512, 1024, 1518, 2048, 4096, and 9100 bytes
- Traffic load (percent of interface line rate): 100 percent
- 40-Gbps traffic pattern (96 ports): Port pair for unicast, and 1-to-95 fanout for multicast

In the port-pair configuration for 40-Gbps mode, all 96 ports were configured as 48 pairs; each pair consists of two ports to send and receive bidirectional traffic. Figure 3 provides an overview of the port-pair topology.

**Figure 3:** Port-Pair Topology



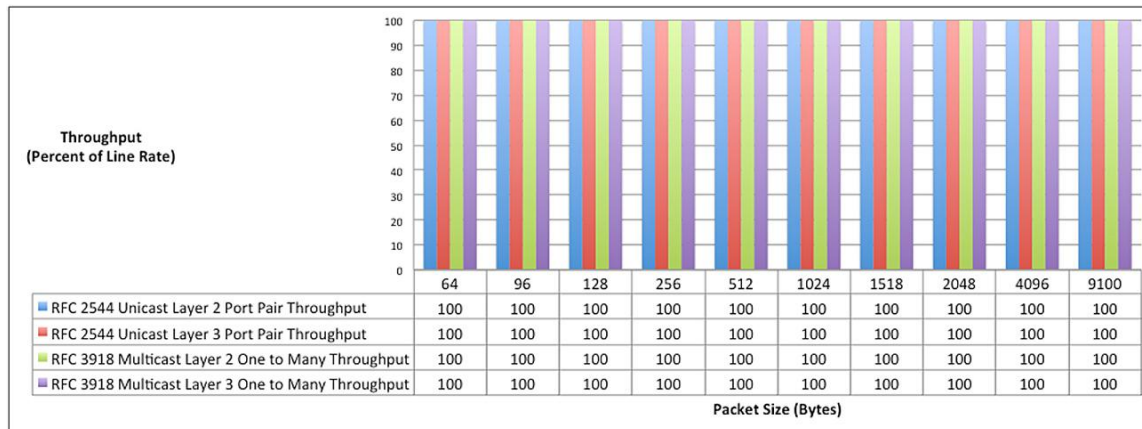
All tests in this report are intended to be reproducible by customers who want to re-create them in their labs with the same topology, hardware, software, and test application configurations. Current or prospective customers interested in repeating these results can contact their local Spirent and Cisco representatives to obtain additional details.

## Test Results

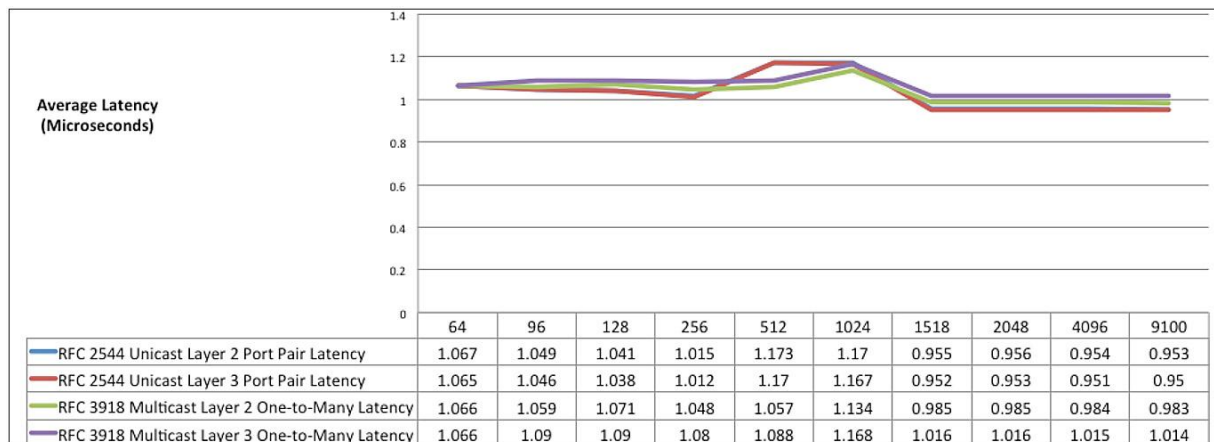
### 96 x 40-Gbps Overview

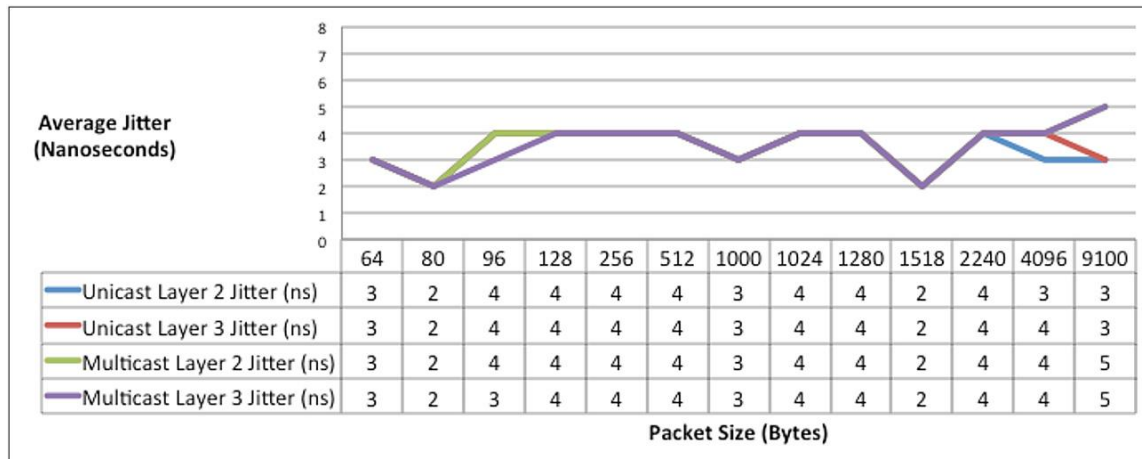
Figures 4 and 5 provide an overview of the 40-Gbps test results.

**Figure 4:** Throughput Overview for 96 x 40-Gbps Ports



**Figure 5:** Unicast and Multicast 100 Percent Load Latency Overview for 96 x 40-Gbps Ports





## RFC 2544: Unicast Port Pair Throughput and Latency for 96 x 40-Gbps Ports

### Overview

RFC 2544 provides an industry-standard benchmark testing methodology to measure unicast port-pair packet throughput and latency. Spirent TestCenter has a full-featured RFC 2544 test package to benchmark those measurements for both Layer 2 and Layer 3.

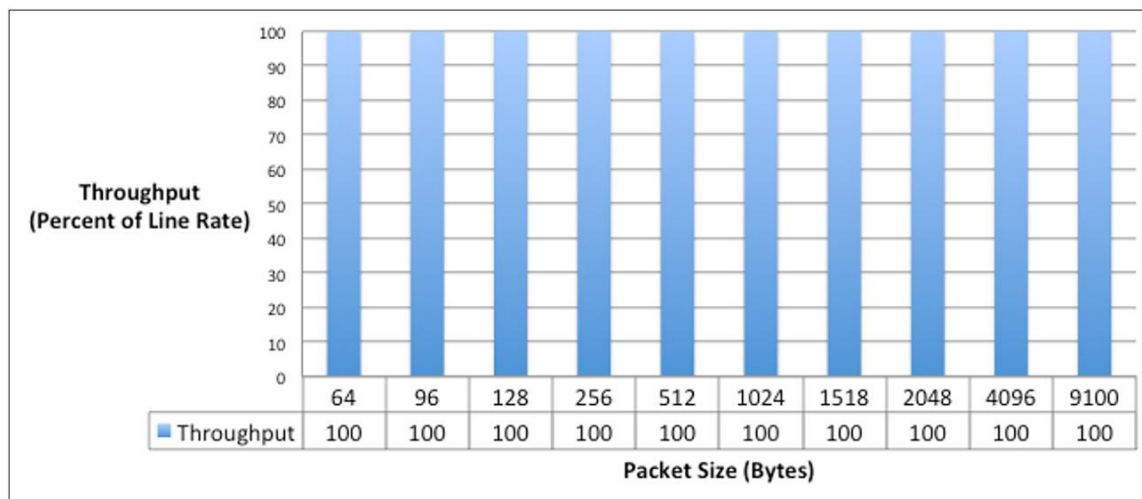
### Configuration

The test was configured using Spirent TestCenter RFC benchmarking test suites with the topology and test metrics described earlier in this document. First, Layer 2 switching performance was measured for the port-pair configuration. Then the configuration was changed to Layer 3, and the same measurements (throughput and latency) were taken. Spirent engineers confirmed that the Spirent TestCenter equipment was properly configured.

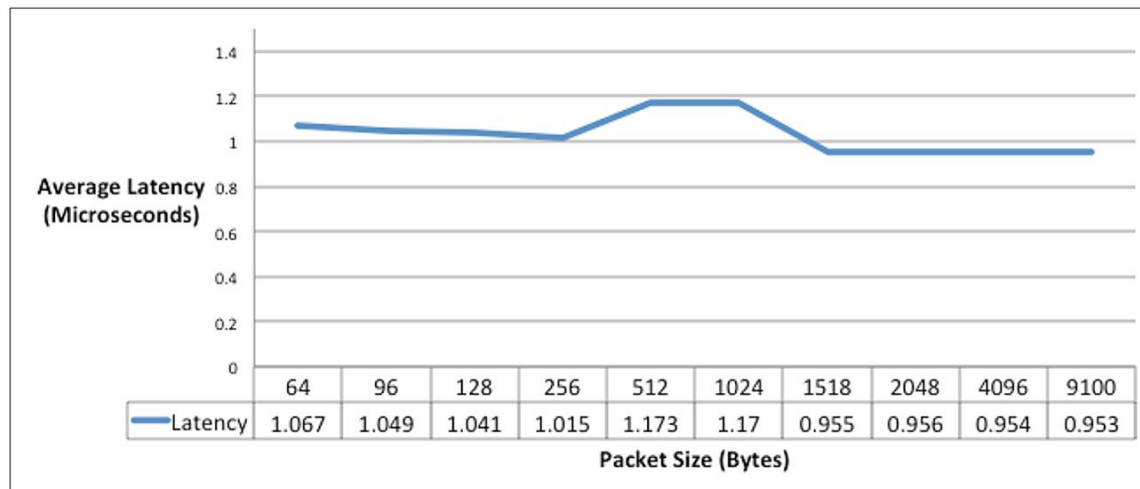
### Test Results

Figures 6 through 9 show the test results.

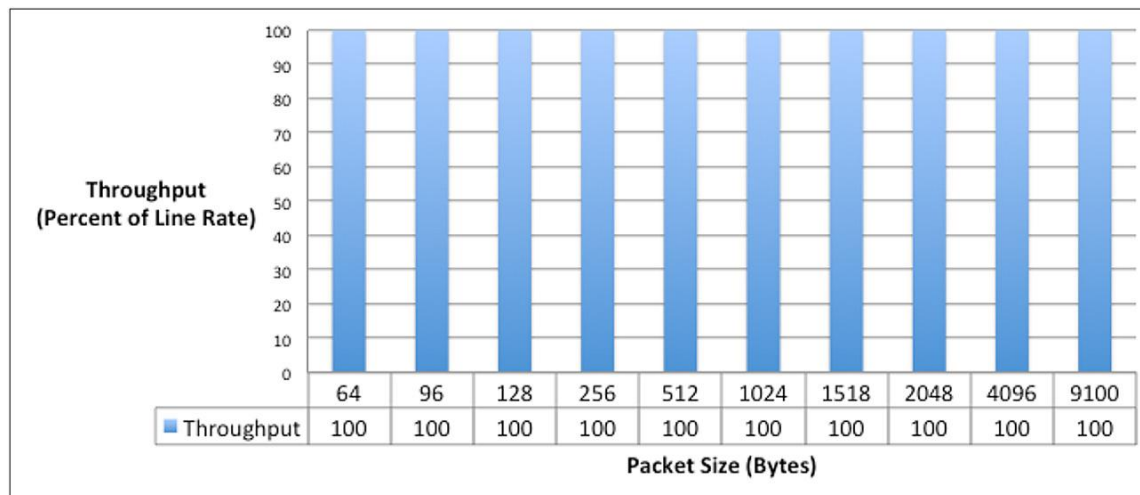
**Figure 6:** RFC 2544 Unicast Layer 2 Port-Pair Throughput for 96 x 40-Gbps Ports



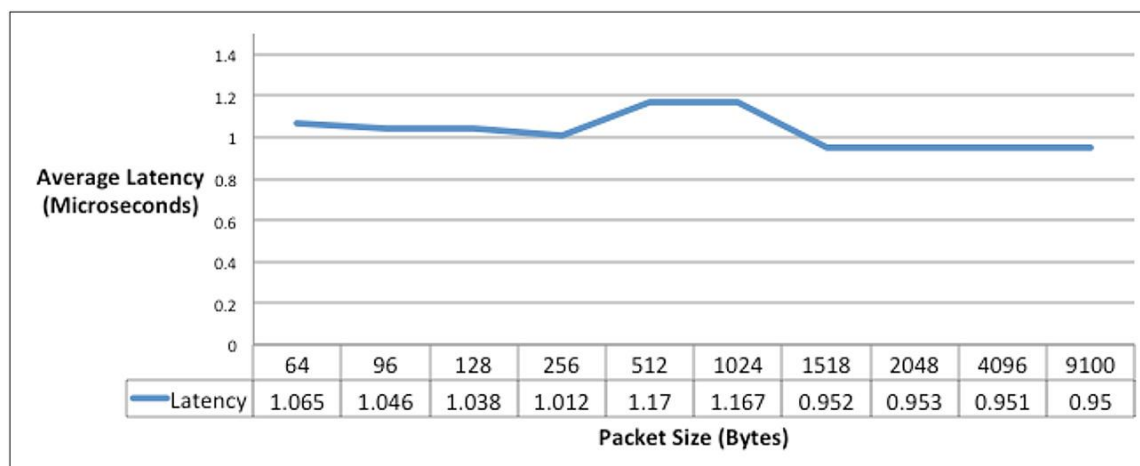
**Figure 7:** RFC 2544 Unicast Layer 2 Port-Pair 100 Percent Load Latency for 96 x 40-Gbps Ports



**Figure 8:** RFC 2544 Unicast Layer 3 Port-Pair Throughput for 96 x 40-Gbps Ports



**Figure 9:** RFC 2544 Unicast Layer 3 Port-Pair 100 Percent Load Latency for 96 x 40-Gbps Ports



## RFC 3918: Multicast Throughput and Latency for 96 x 40-Gbps Ports

### Overview

RFC 3918 provides an industry-standard benchmark testing methodology for measuring multicast throughput and latency. The test can be performed for both Layer 2 and Layer 3 packet forwarding. Spirent TestCenter has a full-featured RFC 3918 test package capable of running all the tests defined in RFC 3918. Switch throughput and forwarding latency are measured for each test.

### Configuration

The test was configured using Spirent TestCenter RFC benchmarking test suites with the topology and test metrics described earlier in this document. First, the switch performance was tested against Layer 2 multicast traffic in a 1-to-95 fanout configuration in which receivers on all 95 interfaces join the multicast groups. Then the configuration was changed to Layer 3, and the same measurements (throughput and latency) were taken. Spirent engineers confirmed that the Spirent TestCenter equipment was properly configured.

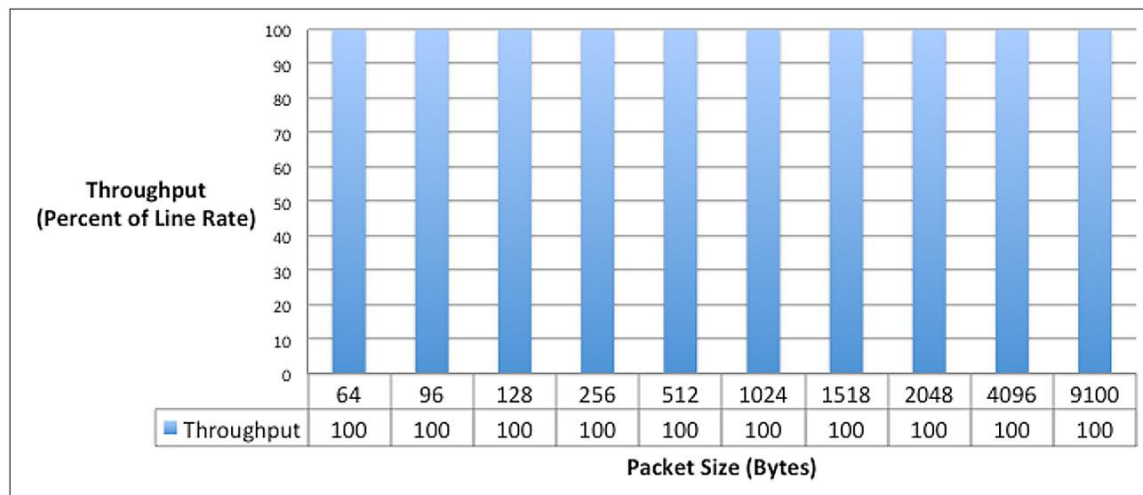
The multicast performance is verified in accumulated mode. In this mode, each receive port joins every group, with the result that the same multicast traffic goes out multiple ports.

The tests were performed with 96 multicast groups.

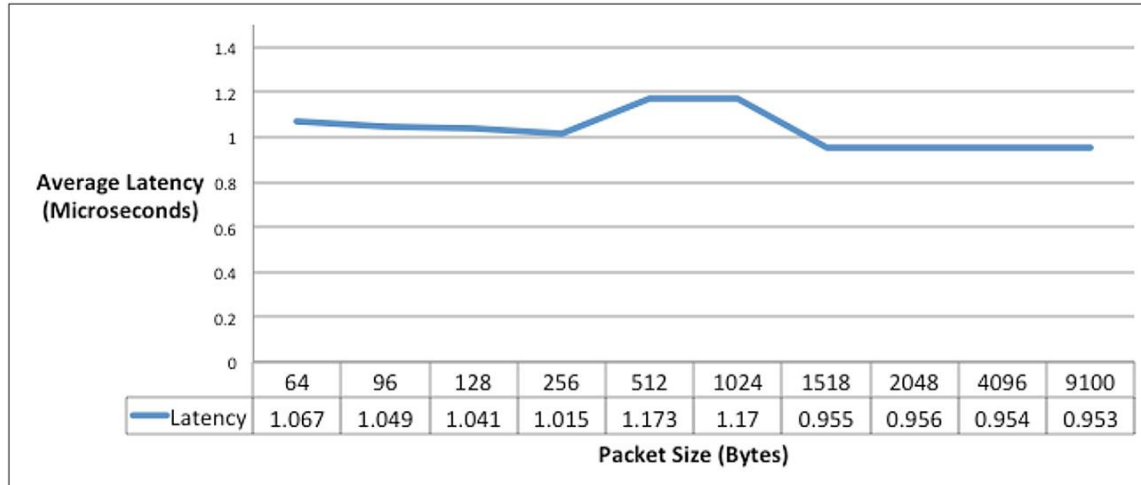
### Test Results

Figures 10 through 13 show the test results.

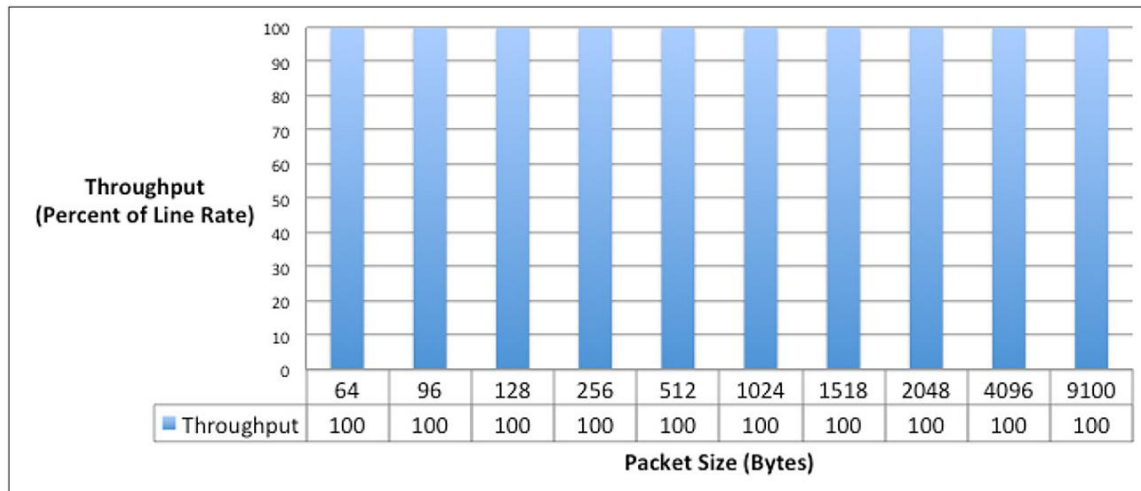
**Figure 10:** RFC 3918 Multicast Layer 2 One-to-Many Throughput for 96 x 40-Gbps Ports



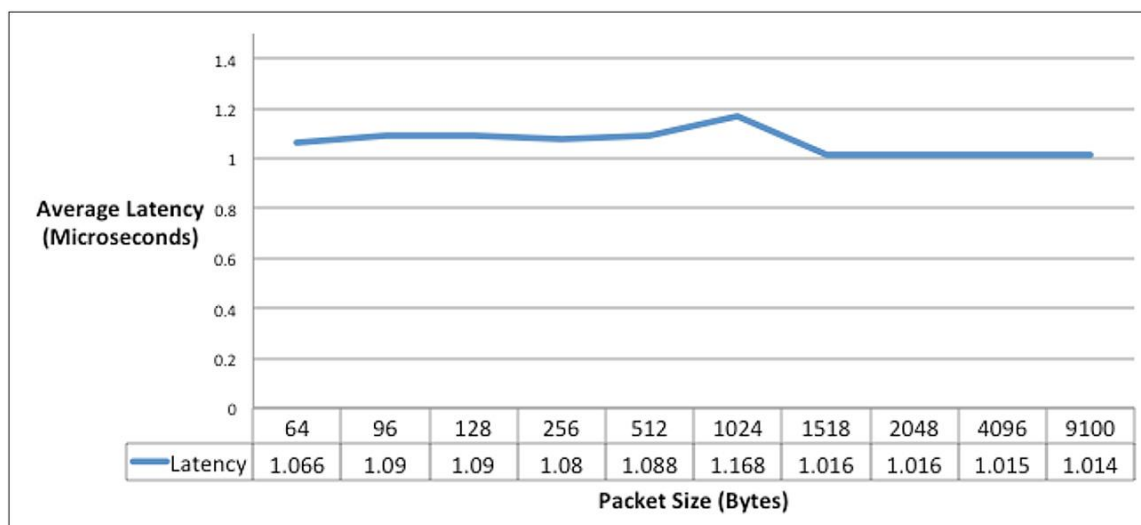
**Figure 11:** RFC 3918 Multicast Layer 2 One-to-Many 100 Percent Load Latency for 96 x 40-Gbps Ports



**Figure 12:** RFC 3918 Multicast Layer 3 One-to-Many Throughput for 96 x 40-Gbps Ports



**Figure 13:** RFC 3918 Multicast Layer 3 One-to-Many 100 Percent Load Latency for 96 x 40-Gbps Ports



## SPAN 40 Gigabit Ethernet Performance

### Overview

The Cisco Nexus 6004 scales up to 16 line-rate SPAN sessions, for both Layer 2 and Layer 3, with Cisco NX-OS 6.0(2)N1(1). This section details a test topology and results that show 40 Gigabit Ethernet line-rate SPAN capability for 16 simultaneous SPAN sessions.

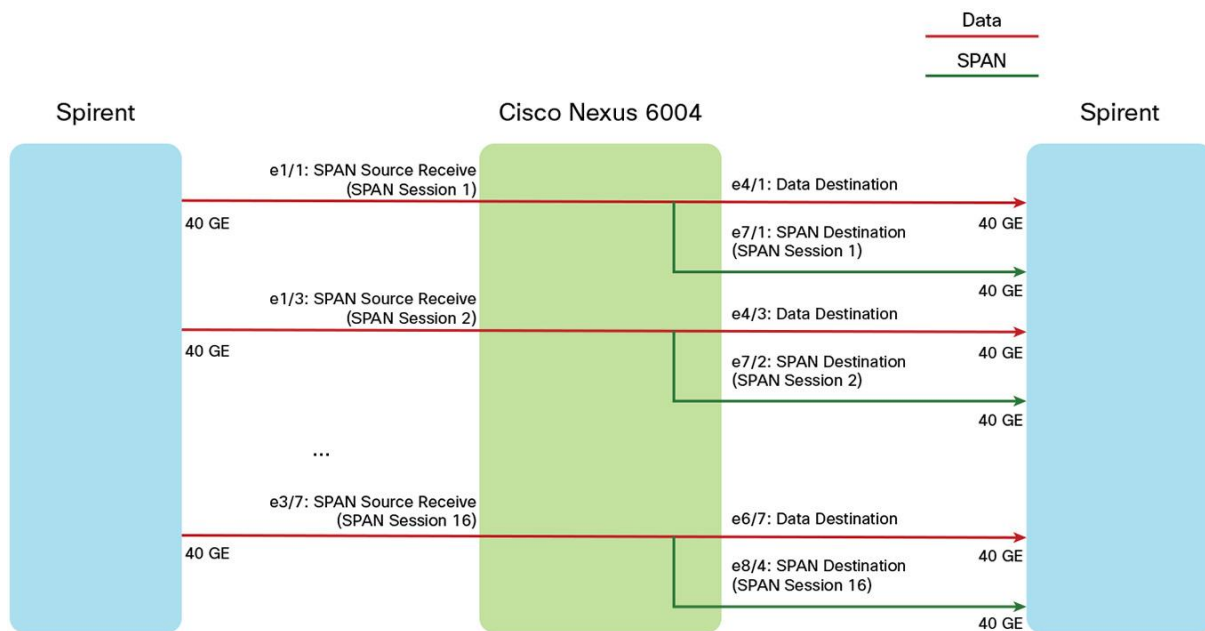
### Configuration

In this test scenario, the data traffic consists in 16 source ports simultaneously sending 40 Gigabit Ethernet Layer 3 unicast traffic to 16 destination ports. The packet size is 1024 bytes.

Sixteen Local SPAN sessions are configured. Each session has one of the data source ports configured as the SPAN source for receiving traffic. Each session has one destination port configured; this port is different from the original data destination port.

Figure 14 shows the topology.

**Figure 14:** SPAN Topology

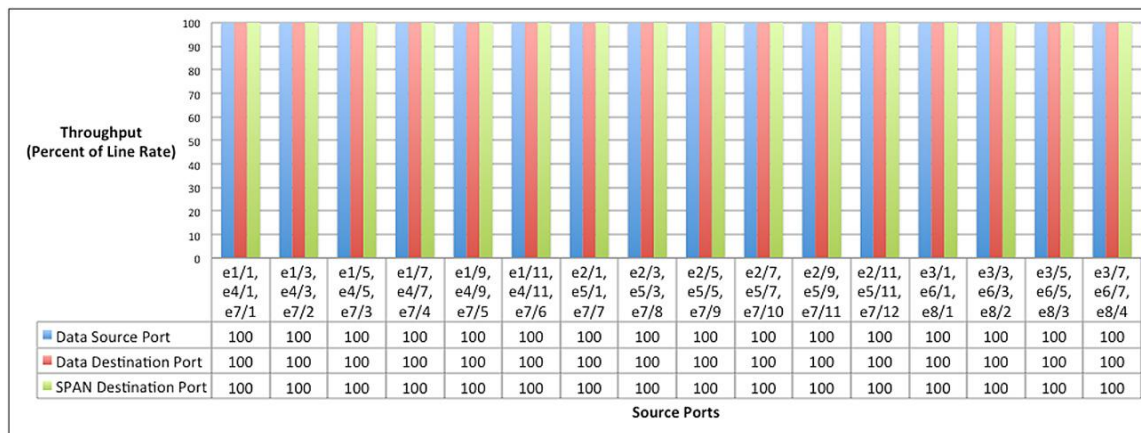


### Test Results

The results show that all SPAN destination ports receive 40 Gigabit Ethernet line-rate traffic with no packet loss. The data flow is not affected, and the data destination ports receive 40 Gigabit Ethernet line-rate traffic with no packet loss.

Figure 15 shows the test results.

**Figure 15: SPAN Test Results**



## SPAN and Multicast 40 Gigabit Ethernet Performance

### Overview

The Cisco Nexus 6004 can perform both line-rate SPAN and multicast replication simultaneously without any effect on data traffic performance. This section discusses a test scenario and results that show the SPAN and multicast capability of the Cisco Nexus 6004.

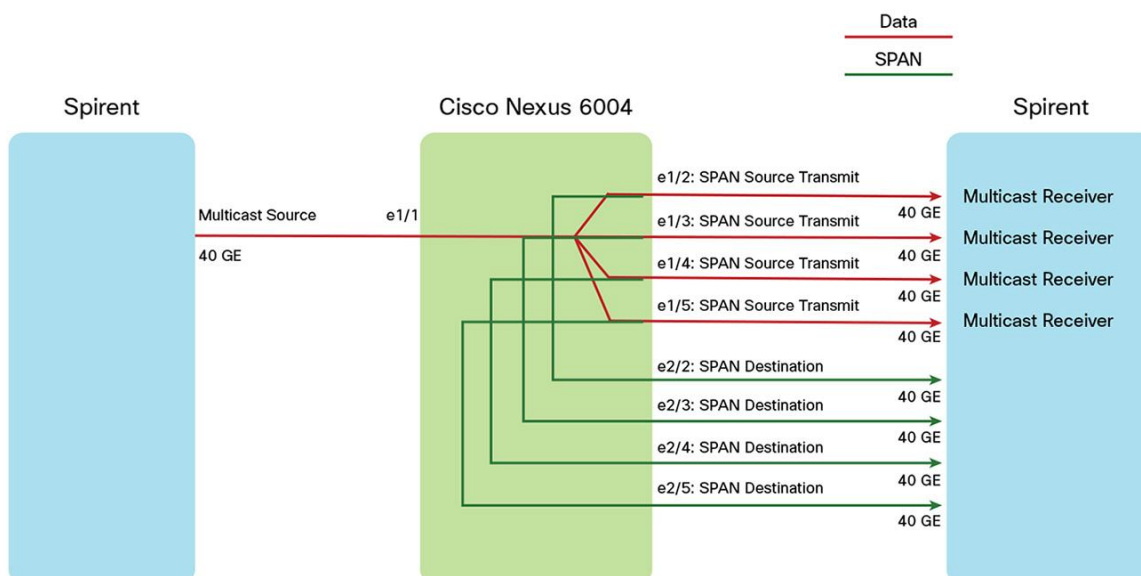
### Configuration

In this test scenario, one port of the Spirent traffic generator is configured as a multicast source, sending 40 Gigabit Ethernet line-rate Layer 3 multicast traffic. Spirent TestCenter is configured in accumulated mode with five multicast groups. The packet size is 1024 bytes. There are four multicast receivers.

Four local SPAN sessions are defined. Each session has one of the multicast receiver ports configured as a source for transmission. Each session has one destination port configured; this port is different from the original data source port.

Figure 16 shows the topology.

**Figure 16: SPAN and Multicast Topology**

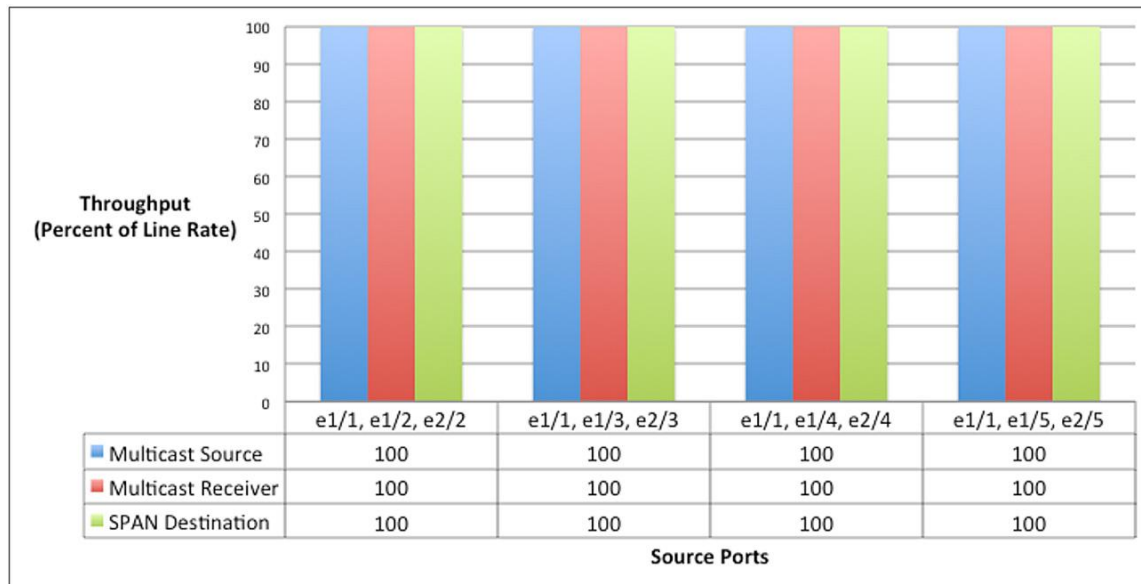


## Test Results

The results show that all multicast receivers and all SPAN destinations receive 40 Gigabit Ethernet line-rate traffic with no packet loss.

Figure 17 shows the results.

**Figure 17: SPAN and Multicast Test Results**



## RFC 2889: Broadcast Rate with Latency Measurement for 96 x 40-Gbps Ports

### Overview

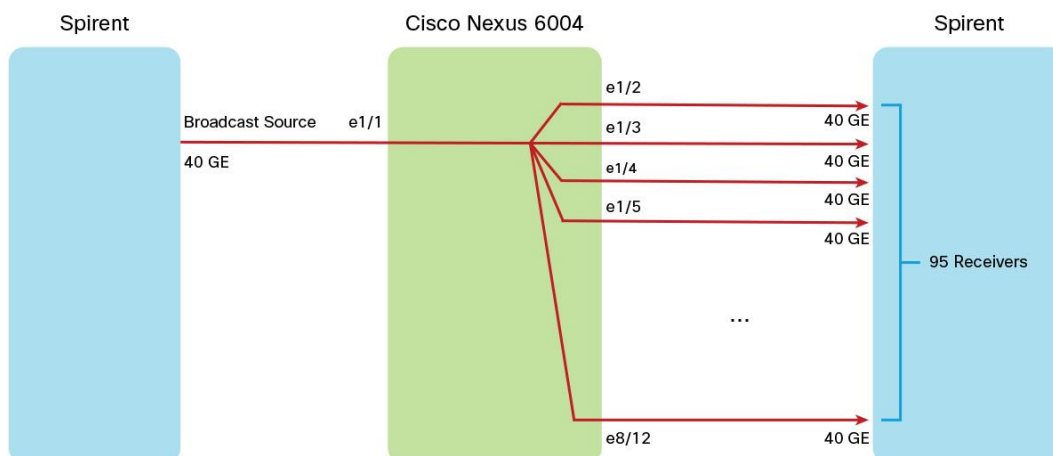
The objective of the Broadcast Frame Forwarding and Latency Test is to determine the throughput and latency of the switch when forwarding broadcast traffic.

### Configuration

The test scenario consists of one 40 Gigabit Ethernet source port sending Layer 2 broadcast traffic to all other 95 40 Gigabit Ethernet ports.

Figure 18 shows the topology.

**Figure 18: RFC 2889 Broadcast Measurement Topology**

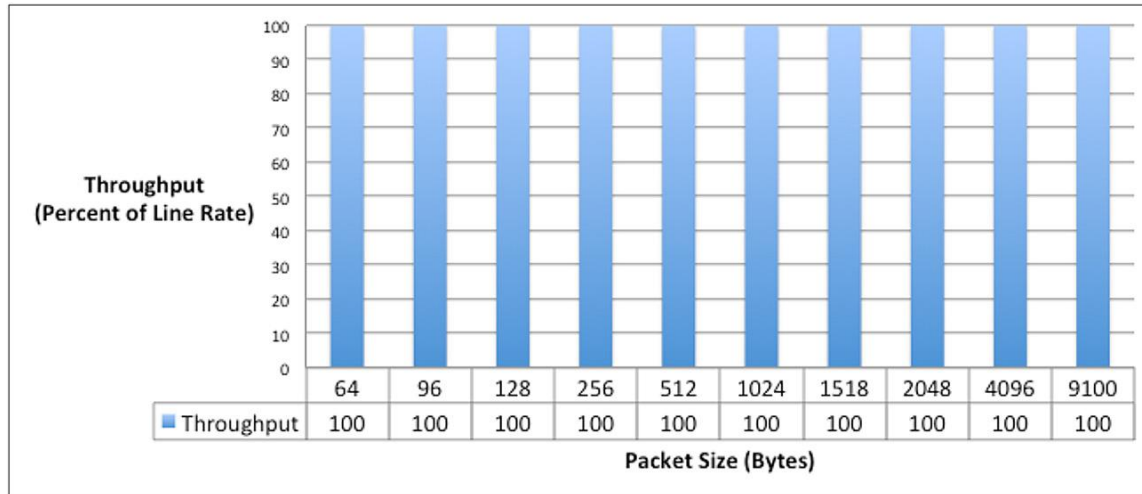


## Test Results

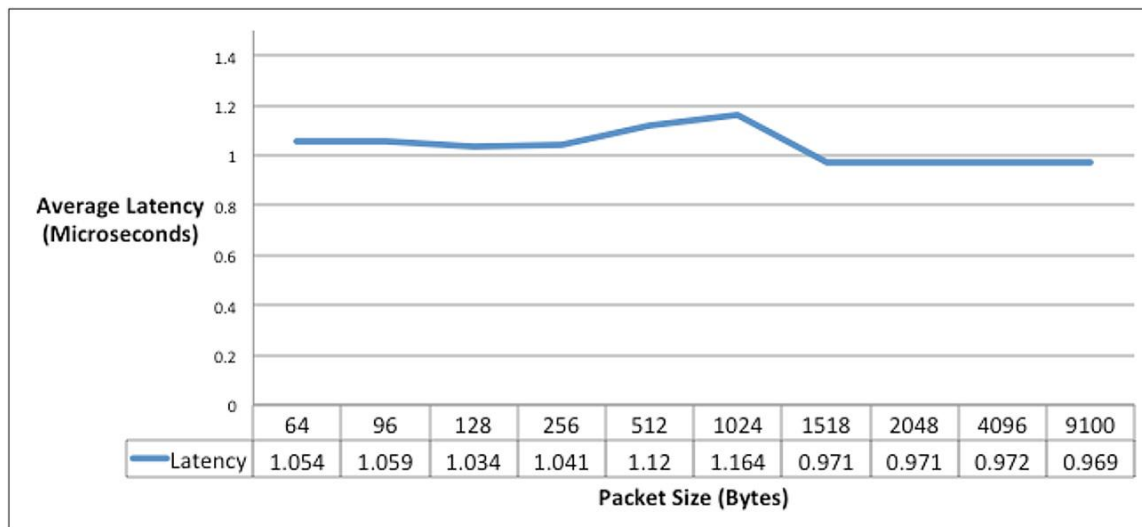
These results show that the Cisco Nexus 6004 can forward broadcast frames at 40 Gigabit Ethernet line rate for all packet sizes with an average latency of 1 microsecond.

Figures 19 and 20 show the test results.

**Figure 19:** RFC 2889 Broadcast Rate Measurement Throughput



**Figure 20:** RFC 2889 Broadcast Measurement Latency



## Conclusion

The Cisco Nexus 6004 Switch is a high-performance, high-density, low-latency 10 and 40 Gigabit Ethernet and FCoE switch that can forward packets at line rate on all ports with an average latency of 1 microsecond. It also has built-in line-rate high-scalability SPAN and multicast capabilities.

## For More Information

- Cisco Nexus 6000 Series Switches: <http://www.cisco.com/go/nexus6000>
- Public Cisco tests using Spirent equipment: <http://www.spirent.com/go/cisco>

© 2013 Cisco and/or its affiliates. All rights reserved. Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: [www.cisco.com/go/trademarks](http://www.cisco.com/go/trademarks). Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

© 2013 Spirent. All Rights Reserved.

All of the company names and/or brand names and/or product names referred to in this document, in particular, the name "Spirent" and its logo device, are either registered trademarks or trademarks of Spirent plc and its subsidiaries, pending registration in accordance with relevant national laws. All other registered trademarks or trademarks are the property of their respective owners.

The information contained in this document is subject to change without notice and does not represent a commitment on the part of Spirent. The information in this document is believed to be accurate and reliable; however, Spirent assumes no responsibility or liability for any errors or inaccuracies that may appear in the document.