



Cisco Nexus 3548 Switch Performance Validation

December 2012



Contents

Executive Summary	3
Test Bed	4
How Testing Was Performed	4
Test Results	6
RFC 2544: Unicast Latency, Throughput, and Jitter	6
RFC 2544: Unicast Latency, Throughput, and Jitter with ACL	11
RFC 2544: Unicast Latency, Throughput, and Jitter with NAT	13
RFC 3918: Multicast Latency, Throughput, and Jitter	15
Warp SPAN Latency	19
Conclusion	19
About Cisco	20
About Spirent	20



Executive Summary

Cisco, in partnership with Spirent, conducted performance benchmark tests for the Cisco Nexus[®] 3548 Switch in the Spirent Proof of Concept (SPOC) Lab. The testing focused on evaluating the switch performance with various traffic profiles and Cisco[®] NX-OS Software features up to and including the most stressful testing conditions.

The Cisco Nexus 3548 is an ultra-low-latency, high-density, high-performance switch. It provides 48 front panel ports supporting 100 Megabit Ethernet¹ and 1, 10, and 40 Gigabit Ethernet¹ in a one-rack-unit (1RU) form factor. Using cut-through technology, the Cisco Nexus 3548 can forward a packet through the switch in as little as 170 nanoseconds (ns).

To fully validate the Cisco Nexus 3548, rigorous testing needed to be performed. IETF RFC 2544² and 3918³ are widely accepted in the industry as standard benchmarking methodologies to evaluate switch performance by measuring throughput, packet forwarding rate, forwarding latency, and jitter at various conditions. Spirent TestCenter has built-in RFC 2544 and 3918 test suites, which were used during the tests.

In addition to standard RFC testing with a basic switch configuration, switch performance was tested when with access control lists (ACLs) and Network Address Translation (NAT) enabled on all interfaces, to evaluate whether those critical features have any impact on switch performance.

During the lab testing, the Cisco Nexus 3548 forwarded traffic in less than 250 ns with 100 percent throughput in a full mesh, full-load traffic profile with NAT and ACL enabled. Under most of the common network and traffic configurations, the Cisco Nexus 3548 forwards the packet in less than 200 ns.

Highlights of the Cisco Nexus 3548 performance test results include:

- 100 percent throughput without a single packet drop
- Switch latency of less than 200 ns under most common network and traffic configurations
- Identical latency characteristics for Layer 2 and 3 traffic
- Ultra low and consistent jitter under the most stressful test configuration (less than 8 ns)
- No performance impact when ACL and NAT are enabled on all interfaces

Spirent configured and validated Spirent TestCenter's 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 those validated configurations.

¹ 100M and 40G will be supported after FCS

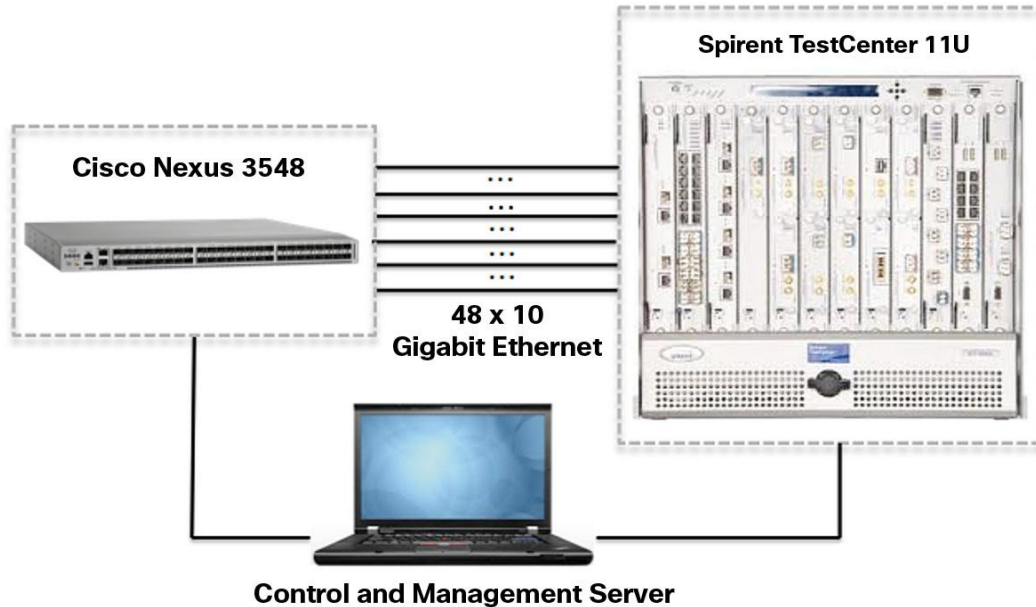
² IETF Network Working Group: Benchmarking Methodology for Network Interconnect Devices.

³ IETF Network Working Group: Methodology for IP Multicast Benchmarking.

Test Bed

Figure 1 provides an overview of the test bed.

Figure 1: Test Bed



How Testing Was Performed

The test topology and test bed includes a Cisco Nexus 3548 Switch, a Spirent TestCenter 11U with forty-eight 10 Gigabit Ethernet ports, and a control and management server on an out-of-band network. Table 1 shows the hardware and software matrix for the test configuration.

Table 1: Test Configuration Hardware and Software Matrix

Device	Software	Notes
One Cisco Nexus 3548 Switch with 48 x 10 Gigabit Ethernet interfaces	Cisco NX-OS 5.0(3)A1(1) (prerelease)	100 Megabit Ethernet and 40 Gigabit Ethernet interfaces are not supported at first customer shipment (FCS)
One Spirent TestCenter 11U chassis	Spirent Test Center Application 4.03 & 4.10	-
Six Spirent HyperMetrics CV test modules, each with 8 x 10 Gigabit Ethernet interfaces	Spirent Test Center Application 4.03 & 4.10	Total of 48 10 Gigabit Ethernet ports
48 x 1M Passive Twinax cables	-	-

The Cisco Nexus 3548 Switch's 48 ports are all connected to Spirent TestCenter HyperMetrics CV 10 Gigabit Ethernet test modules through a 1-meter 10GBASE-CU Enhanced Small Form-Factor Pluggable (SFP+) Twinax cable. The built-in RFC benchmark test suites from Spirent TestCenter Application Version 4.10 are used to validate the performance of the Cisco Nexus 3548.

First, the Cisco Nexus 3548 and Spirent TestCenter are configured to run RFC 2544 test suites to measure switch performance for unicast traffic, including throughput, latency, and jitter, for both Layer 2 and Layer 3 traffic.

Next, multicast performance is measured with RFC 3918 for both Layer 2 and Layer 3 traffic. Again, the throughput, latency, and jitter are all measured on each test.

Finally, the unicast performance test is repeated again after adding ACL and NAT configuration on all interfaces of the Cisco Nexus 3548 Switch. The results are compared with previous test results in which ACL and NAT were not enabled to see if the use of ACL and NAT degrades performance.

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

- Throughput is defined as the maximum forwarding rate the switch can achieve without losing a single packet. The throughput test measures the number of packets the switch can move at any given time interval, typically 1 second. For example, 100% line rate of 10GbE interface can move 14 millions packet per seconds.
- Switch forwarding latency is the length 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 like the Cisco Nexus 3548, this test measures using 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.
- Jitter is the statistical variance of the packet interarrival time. The lower the jitter, the better. The jitter test indicates how well the switch can deliver packets at a constant rate in very predicable intervals.


Packet-size, traffic-load, and traffic-pattern metrics are tested in all test cases to help ensure adequate test coverage, including during a 64-byte-packet full mesh, 100 percent line-rate performance test, which is the most stressful RFC test. During the full mesh test, all ports send traffic at the configured rate (up to 100 percent of line rate) to all other parts in a full mesh pattern.

- Packet size (bytes): 64, 128, 256, 512, 1024, 1518, 4096, and 9216 bytes.
- Traffic load (percent of interface line rate): 10, 50, 90, and 100 percent.
- Traffic pattern (48 ports): Port pair and full mesh for unicast, and 1-to-47 fan-out for multicast.

In the port pair configuration, all 48 ports are configured as 24 pairs; each pair consists of two ports to send and receive bidirectional traffic.

In the full mesh configuration, all 48 ports send and receive bidirectional traffic to all other ports. This is the most stressful performance test for evaluating switch performance. During the test, both the port pair and full mesh are tested (Figure 2).

Figure 2: Port Pair and Full Mesh Testing



The Cisco Nexus 3548 has a unique forwarding mode called warp mode. When warp mode is enabled in these tests, switch latency is reduced by up to 12 percent compared to normal mode. The most stressful tests were repeated (unicast Layer 2 and 3 full mesh at 100 percent of line rate, and multicast 1-to-47 fan-out at 100 percent of line rate) for both modes (normal mode and warp mode) to demonstrate the performance of warp mode forwarding. For detailed information about the technical differences between normal mode and warp mode, please go to <http://www.cisco.com/go/nexus3000>.

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

RFC 2544: Unicast Latency, Throughput, and Jitter

Description

RFC 2544 provides an industry-standard benchmark testing methodology to measure unicast packet latency, jitter, and throughput. Spirent TestCenter has a full-featured RFC 2544 test package to benchmark those measurements for both Layer 2 and Layer 3 traffic in any combination of traffic patterns, including port pair and full mesh configurations.

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 is measured for both port pair and full mesh configurations. Then the configuration is changed to Layer 3, and the same measurements (throughput, latency, and jitter) are taken. Spirent engineers confirmed that the Spirent TestCenter equipment was properly configured.

Test Results

The following test results were collected:

Warp mode:


- Layer 2 and Layer 3 unicast RFC 2544 throughput
- Layer 2 RFC 2544 port pair unicast latency with 10, 50, 90, and 100 percent traffic load
- Layer 2 RFC 2544 full mesh unicast latency with 10, 50, 90, and 100 percent traffic load
- Layer 3 RFC 2544 port pair unicast latency with 10, 50, 90, and 100 percent traffic load
- Layer 3 RFC 2544 full mesh unicast latency with 10, 50, 90, and 100 percent traffic load

Normal mode:

- Layer 2 and Layer 3 unicast RFC 2544 throughput
- Layer 2 RFC 2544 full mesh unicast latency with 100 percent traffic load
- Layer 3 RFC 2544 full mesh unicast latency with 100 percent traffic load

Figures 3 through 10 show the results.

Figure 3: Layer 2 and Layer 3 Unicast RFC 2544 Throughput in Warp Mode and Normal Mode



Note: Regardless the forwarding mode or traffic type, Cisco Nexus 3548 can forward the packet at 100 percent line rate for all sizes of packets.

Figure 4: Layer 2 Unicast RFC 2544 Port Pair Latency in Warp Mode




Figure 5: Layer 2 Unicast RFC 2544 Full Mesh Latency in Warp Mode




Figure 6: Layer 3 Unicast RFC 2544 Port Pair Latency in Warp Mode




Figure 7: Layer 3 Unicast RFC 2544 Full Mesh Latency in Warp Mode




Figure 8: Layer 3 Unicast RFC 2544 Port Pair Latency Comparison: Normal Mode Compared to Warp Mode



Figure 9: Layer 3 Unicast RFC 2544 Full Mesh Latency Comparison: Normal Mode Compared to Warp Mode





Figure 10: Layer 2 and Layer 3 Unicast RFC 2544 Full Mesh and Port Pair Jitter in Warp Mode



RFC 2544: Unicast Latency, Throughput, and Jitter with ACL

Description

IP ACL is an essential feature of any network device. It is important to measure the performance of the switch when hundreds of lines of ACL entries are applied to the system. The goal of this test is to determine whether there is any performance impact when the ACL feature is enabled. After enabling ACL, the RFC 2544 test package is used to perform the same measurement as in the previous RFC 2544 tests. The purpose of this test was not to demonstrate ACL scalability, only to evaluate performance of the feature. The Cisco Nexus 3548 supports up to 4,000 ACL entries in hardware.

Configuration

In the test, 201 lines of ACL entries are configured and applied to all 48 Layer 3 interfaces of the switch in the ingress direction. Then the Spirent TestCenter RFC 2544 test package is launched to measure the Layer 3 full mesh traffic pattern for all 48 ports. During the test, throughput, latency, and jitter are measured.

Test Results

Figures 11 through 13 show the test results.

Figure 11: Layer 2 and Layer 3 Unicast RFC 2544 Full Mesh Throughput with ACL in Normal Mode and Warp Mode






Figure 12: Layer 3 Unicast RFC 2544 Full Mesh Latency Comparison: With and Without ACL at 100 Percent Throughput in Warp Mode



Note: There is no latency impact at all when ACL is enabled on the system.

Figure 13: Layer 3 Unicast RFC 2544 Full Mesh Jitter Comparison: With and Without ACL in Warp Mode



RFC 2544: Unicast Latency, Throughput, and Jitter with NAT

Description

NAT operates on a network device connecting two networks together; one of these networks (designated as inside) is addressed with either private or obsolete addresses that need to be converted into legal addresses before packets are forwarded onto the other network (designated as outside). The translation operates in conjunction with routing, so that NAT can simply be enabled on a customer-side Internet access router when translation is desired.

The capability to perform NAT in hardware is a unique feature of the Cisco Nexus 3548 with Cisco NX-OS. The testing used an approach similar to that for ACL performance evaluation, measuring system performance when NAT is enabled. Testing then compared switch performance with and without NAT. The purpose of this test was not to demonstrate NAT scalability, only to evaluate performance of the feature. The Cisco Nexus 3548 supports up to 1,000 NAT translations in hardware.

Configuration

Different customers use NAT in different ways. From a switch performance point of view, the best way to measure whether NAT affects switch performance is to compare the performance results with and without NAT enabled in a very similar setup. Therefore, port pair RFC 2544 testing was performed with 24 inside interfaces and 24 outside interfaces. Among the 24 outside interfaces, static NAT was used for 12 of them, and static Port Address Translation (PAT) was used for the other 12 interfaces. Cisco NX-OS phase 1 supports only static NAT and PAT; dynamic NAT and PAT will be supported in a future release. After the RFC 2544 test with NAT is performed, the performance results, including throughput, latency, and jitter, are compared.

Because TCP header with special Spirent signature field was added for NAT testing to endure the received packet are indeed a legitimate packet from Spirent, the minimum packet size becomes 78 bytes during the NAT testing.

Test Results

Figures 14 through 16 show the results.

Figure 14: Layer 3 Unicast RFC 2544 Throughput with NAT in Warp Mode and Normal Mode




Figure 15: Layer 3 Unicast RFC 2544 Latency in Warp Mode With and Without NAT





Figure 16: Layer 3 Unicast RFC 2544 Latency in Normal Mode With and Without NAT



Notes: Figures 14 through 16 show that enabling the NAT function has no impact on switch performance.



RFC 3918: Multicast Latency, Throughput, and Jitter

Description

RFC 3918 provides an industry-standard benchmark testing methodology for measuring multicast latency, jitter, and throughput. 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. Similar to RFC 2544 Testing, switch throughput, forwarding latency and jitter are measured for each tests.

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 is tested against Layer 2 multicast traffic in a 1-to-47 fanout configuration in which receivers on all 47 interfaces all join the multicast groups. Then the configuration is changed to Layer 3, and the same measurements (throughput, latency, and jitter) are taken. Spirent engineers confirmed that the Spirent TestCenter equipment was properly configured.

Test Results

Warp mode:

- Layer 2 and Layer 3 multicast RFC 3918 throughput
- Layer 2 RFC 3918 1-to-47 latency with 10, 50, 90, and 100 percent traffic load
- Layer 3 RFC 3918 1-to-47 latency with 10, 50, 90, and 100 percent traffic load

Normal mode:

- Layer 2 and Layer 3 multicast RFC 3918 throughput
- Layer 2 RFC 3918 1-to-47 latency with 100 percent traffic load
- Layer 3 RFC 3918 1-to-47 latency with 100 percent traffic load

Figures 17 through 22 show the test results.

Figure 17: Layer 2 and Layer 3 Multicast RFC 3918 Throughput




Figure 18: Layer 2 Multicast RFC 3918 1-to-47 Latency in Warp Mode




Figure 19: Layer 3 Multicast RFC 3918 1-to-47 Latency in Warp Mode




Figure 20: Layer 2 Multicast RFC 3918 1-to-47 Latency Comparison: Normal Mode Compared to Warp Mode




Figure 21: Layer 3 Multicast RFC 3918 1-to-47 Latency Comparison: Normal Mode Compared to Warp Mode





Figure 22: Layer 2 and Layer 3 Multicast RFC 3918 1-to-47 Jitter





Warp SPAN Latency


Description

The Cisco Nexus 3548 Switch supports a mode of operation where traffic entering a single port is replicated out a set of egress ports with minimal processing.

Configuration

The test was configured using Spirent Test Center Application. Port E1/36 was connected to a traffic source sending multicast data, and Port E1/1 was connected to a multicast receiver. Ports E1/5 through E1/7 were configured for Warp SPAN, thus also receiving all traffic that was entering E1/36. The port-to-port latency on ports E1/5 through E1/7 was measured in Spirent TestCenter application v4.10. Since there is no standard RFC test package for this test, all test were done manually with different packet size.

Test Results



Conclusion

The Cisco Nexus 3548 Switch is a high-performance, ultra-low-latency Ethernet switch with NAT supported in hardware without any performance degradation. It can forward packets under the most stressful conditions in 243 ns in warp mode, and it can forward packets in less than 200 ns under most of the common network environment conditions at 90 percent of line rate. It also has extremely low jitter; even at 100 percent of line rate with a full mesh traffic configuration, jitter is only 7 ns and is very consistent across the packet size. In addition, there is no any performance degradation when ACL or NAT is enabled because all those critical features are all performed in the hardware switching path inside the switch.



About Cisco

Cisco is the leader in networking that transforms how people connect, communicate, and collaborate. Cisco has led in the innovation of IP-based networking technologies, including routing, switching, security, TelePresence, unified communications, video, and wireless. The company's responsible business practices help ensure accountability, business sustainability, and environmentally conscious operations and products.

For more information about Cisco and the products used, please visit:

- <http://www.cisco.com/go/algoboost>
- <http://www.cisco.com/go/nexus3548>
- <http://www.cisco.com/go/nexus>

About Spirent

Spirent Communications is a global leader in test and measurement, inspiring innovation in development labs, communication networks, and IT departments. At Spirent Communications, we work behind the scenes to help the world communicate and collaborate faster, better, and more often. The world's leading communications companies rely on Spirent to help design, develop, and deliver world-class networks, devices, and services. Spirent's lab test solutions are used to evaluate performance of the latest technologies.

As with any performance measurement, the latency test results can vary from different test iterations due to the margin of error introduced by testing equipment even under the same version of hardware and software. Spirent Test Center and CV-10G-S8 have 10ns latency resolution with +/- 20ns the margin of error.

For more information about Spirent and the products used, please visit:

- <http://www.spirent.com>
- <http://www.spirent.com/stc>

© 2012 Spirent Communications, Inc. 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 pending registration in accordance with relevant national laws. All rights reserved. Specifications subject to change without notice.

© 2012 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. 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)