Cisco Catalyst 3850

Stackable Aggregation Switch

Independent Performance Assessment

DR150225C

April 2015

Miercom
www.miercom.com

Contents

1 - Executive Summary

Miercom was engaged by Cisco Systems to independently assess the performance and key features of its latest switching system, Catalyst 3850, a stackable aggregation switch supporting several optional modules. Model WS-C3850-24S-E of the system were shipped to and tested at Miercom's main New Jersey lab in late 2014.

The testing focused on the ability of the switch to handle high data volumes with minimal loss and low latency. Key high-availability features were also exercised and assessed. The system was tested for throughput, latency, delay variance and stack convergence. All the performance testing was conducted in accordance with IETF standard RFCs 2544, 2889, and 3918.

In most tests data was delivered to and from the Catalyst via Gigabit/s fiber (SFP) connections. A full ring topology was used for convergence testing in the stacked configuration (see 'How We Tested' details).

Key Findings and Observations:

- Features designed into the Cisco Catalyst 3850 24S-E afford users high reliability and continued availability. Testing showed that in a stacked configuration, the Cisco StackWise-480 technology provides load balancing and resiliency to minimize the loss of a link or any stacked 3850 switch module.
- The Catalyst 3850 switch delivers full, bidirectional, line-rate throughput for all packet sizes, with zero packet loss, on every 1GE and 10GE port – in all scenarios tested, including port-pair, full-mesh, multicast, Layer 2 and Layer 3, IPv4 and IPv6 full-load traffic tests.
- Traffic traversing the Catalyst 3850 switch incurs consistent, low latency, even with maximum load on all ports.
- The switch applies very low jitter – latency variance – to all passing data, regardless of frame size, even with maximum load on all ports.

Based on the impressive results of our testing, we proudly award the *Miercom Performance Verified Certification* to the Cisco Catalyst 3850 24S-E, having turned in outstanding performance in Miercom's ongoing Gigabit-Switch Study.

Robert Smithers

CEO

Miercom

2 - Product Overview

The Cisco Catalyst 3850 Series is a new family of 1-RU (1.75-inch high) switch models. The models vary in port type (copper RJ-45, fiber SPF), number of ports (12,24 or 48 fixed ports), and support for link power (PoE, PoE+). All models include integral wireless controller support (wireless Access Points are separate modules), stackability support, (all switch modules need to be running the same operational software and features) built-in resiliency and stateful switchover (SSO) support, and a slot for an uplink option module. Uplink modules currently offered support: four 1GE ports, two 10GE/two 1GE ports, four 10GE ports and, included in the test configuration, a hybrid module with two 1GE and two 10GE ports.

Cisco Catalyst 3850 24S-E (front)

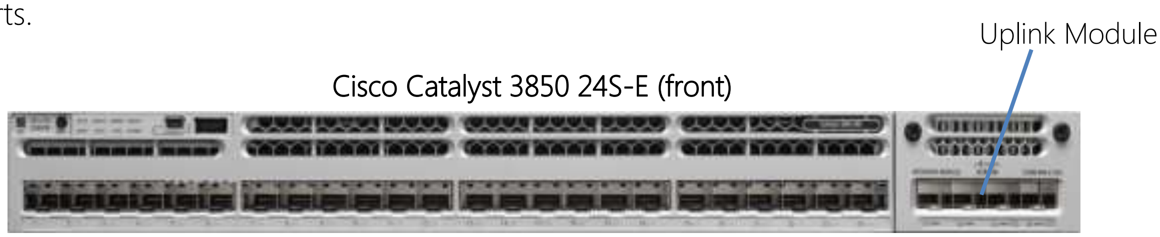

The integrated-wireless-controller support built into the Catalyst 3850 Series is impressive – handling up to 50 Access Points (APs) and 2,000 wireless clients per switch. A Cisco Catalyst 3850 Fiber switch can serve as a mobility controller for small to medium-sized deployments with Cisco 3850 Copper Switches configured as Mobility Agent. However, wireless support or functionality was not tested in this round.

The stackability of the Catalyst 3850 is another of its strengths, which was tested. Stacking the switches can deliver impressive load-sharing and resiliency via stateful switchover. The picture below shows the rear-side cabling for a four-switch stack, featuring Cisco's StackWise–480.

StackWise 480

These two features provide power-sharing and the ability of the stack to survive various failure scenarios.

Resiliency is further enhanced with multiple power supplies and fans. Dual redundant power supplies and dual redundant fans bolster survivability. In a stack, power can be shared and distributed among switch modules if one or more power supplies fail.

The Catalyst 3850 model tested was the stackable 24S-E, featuring 24 x 1GE fixed SFP (fiber) ports. In addition, our switch was equipped with an Uplink Option Module – in our case adding two more 1GE ports, for a total of 26, plus two 10GE SFP+ uplink ports.

In fully populated, all-port testing, Spirent STP 9000 test system collectively delivered 1GE of traffic to each of the Cisco Catalyst 3850 24S-E's 26 ports, as well as 10GE of traffic on each of the two 10GE uplink ports – a theoretical total maximum throughput of 46 Gbps. The Spirent system then check and confirm the number and state of packets forwarded, and the time delay (latency) added by the switch.

The results proved that all of the data delivered to the switch, even full load on all ports, was processed and forwarded by the Catalyst 3850 24S-E properly.

Since the test traffic was all IP packets, all of the regular IP traffic rules had to be obeyed – inter-frame gaps, packet-handling overhead and so on. This means that traffic applied never actually filled the full 1GE and 10GE clock rate of the optical link. This is true of any link carrying IP traffic for that matter. The entire spectrum of frame sizes was applied in testing the Catalyst 3850 24S-E, and in ***all*** cases the maximum load was accepted and forwarded – with no frame loss.

Cisco Catalyst 3850 switches are based on the Cisco Unified Access Data Plane (UADP) application-specific integrated circuit (ASIC). Various additional campus and branch LAN switching and converged mobility deployment options are offered. Fully stackable and with small footprint, the feature-rich and wired-wireless Cisco Catalyst 3850 aggregation switches reduce IT costs and complexity, deliver consistent wired and wireless user experience, enhanced security, and reduce TCO with lower OpEx and energy consumption.

3 – Test Bed

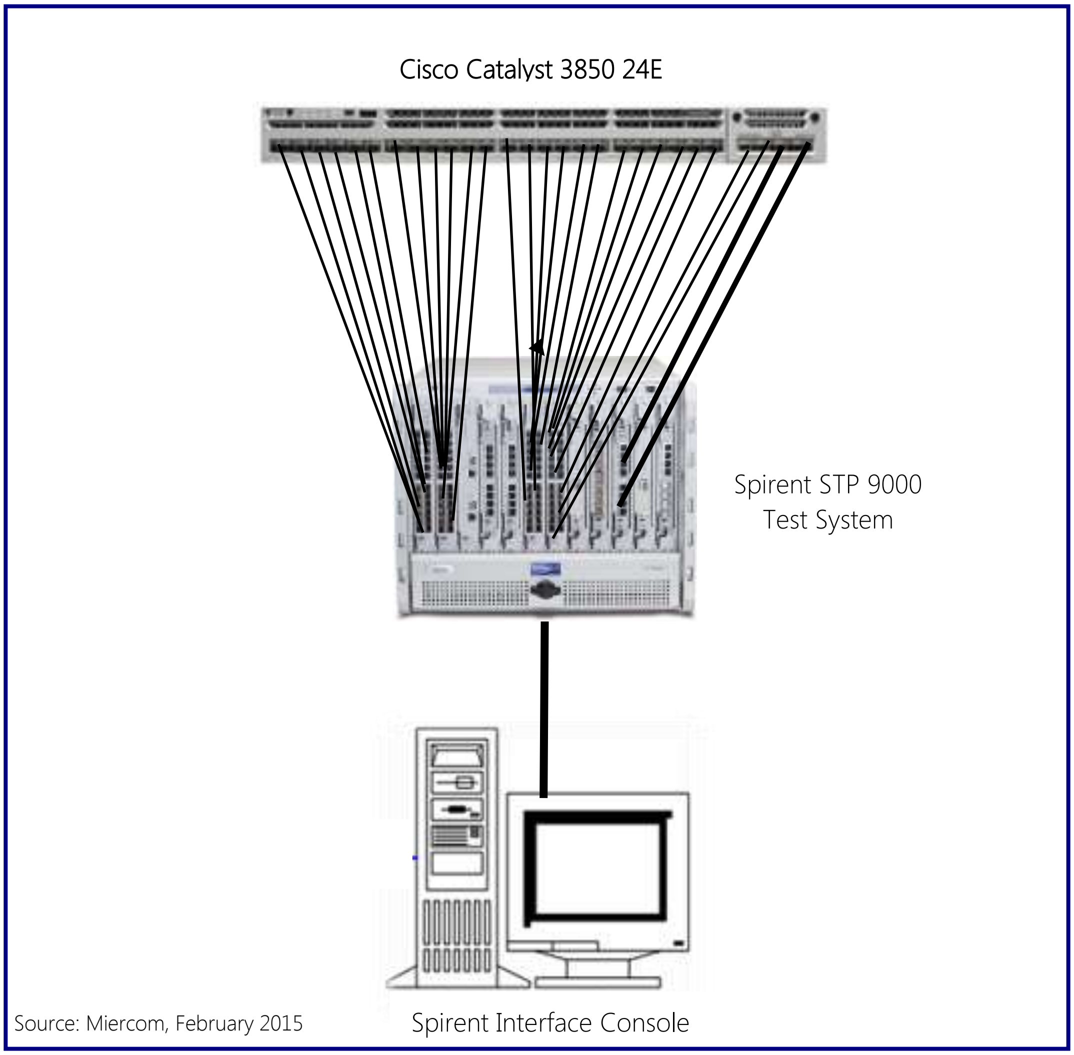

Spirent STP 9000 test ports were connected to every port on the Catalyst 3850 24S-E Switch. Three of the standards used in this testing, and which the Spirent test system incorporates, are detailed in publicly accessible RFCs. The ones most applicable to this testing include RFCs 2544 (port-pair), 2889 (full-mesh), and 3918 ("multicast) – for throughput and latency measurements of Layer 2 and Layer 3 (IPv4 and IPV6, tested separately) bidirectional traffic.

RFC 2544, describes how to conduct basic benchmark tests for throughput, Latency and frame loss measurement. Bidirectional Layer 2/3 (Frame/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, is a reference for conducting more stressful full mesh tests for throughput, latency and frame loss measurement. The Spirent test system provides fully meshed bidirectional traffic flow for these measurements, which fully stresses the switch fabric.

RFC 3918, addresses throughput, latency and packet loss measurement for IPV4/IPV6 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 frame/packet losses. The maximum throughput achievable through the switch was verified, as was proper multicast packet replication. If there is any loss, the test system incrementally ratchets down the delivered load until the maximum throughput without loss is determined. In the case of the Cisco Catalyst 3850, configured as it was for our testing, there was no loss in any of the scenarios tested.

How We Did It

The tests applied measured the latency and loss characteristics of high-volume, customizable traffic patterns traversing through a network platform. Several traffic profiles were simulated for this testing:

- One-to-many multicast
- Full Mesh unicast
- Port-pair unicast

All tests were performed on both Layer 2 (link-layer, switched frames) and Layer 3 (IPv4 and IPv6 packets) where possible. All of these environments are supported by the Cisco Catalyst 3850 Series and were tested.

Tests were separately conducted for 64-, 128-, 256-, 512-, 768-, 1024-, 1280-, 1518-, 2034-, 4068-, and 9216-byte frame sizes. The only exception is IPv6, which employs a 78-byte, minimum-size packet. The term packet and frame are used interchangeably in this report.

The setup used to conduct testing is shown in the previous test bed diagram. The DUT was:

The Cisco Catalyst 3850 24S-E switch running Cisco IOS Software CAT3K_CAA-UNIVERSALK9-M, Version 03.06.01.E RELEASE SOFTWARE (fc3). The switch supports 24 fixed 1-Gigabit Ethernet ports. It also had two more 1GE ports and two 10GE ports via an uplink expansion module.

Spirent: The Spirent TestCenter is an end-to-end testing solution that was used for this testing. Spirent has long been a pioneer in developing test software and hardware. Their systems are designed to fully and accurately assess device performance even at today's 100-Gigabit/s-per-second, core-switch port data rates.

The traffic generation used in this evaluation was based on Spirent standard testing software and processes, which are based on RFCs 2544, 2889, and 3918. The control software run was Spirent TestCenter.

The tests in this report are intended to be reproducible for customers who wish to recreate them with the appropriate test and measurement equipment. Contact reviews@miercom.com for additional details on the configurations applied to the device under test and test tools used in this evaluation. Miercom recommends customers conduct their own needs analysis study and test specifically for the expected network and traffic environment for product deployment before making a product selection.

4 – Port-Pair Throughput and Latency Performance RFC 2544

This throughput test determines the maximum rate at which the Cisco Catalyst 3850 24S-E switch receives and forwards traffic without frame loss. As a rule, frames are sent at a specified rate, which is then continually stepped up in subsequent iterations, using a binary search algorithm, until the maximum rate at which the switch forwards data without losing frames is determined.

Frames can be MAC only Layer 2, or Layer3 IPv4 or IPv6, with or without Extension Headers, or an IPv4/IPv6 mixture. IPv4 and IPv6 were used, in separate tests, and with different packet sizes, starting with minimum-size 64 byte packets (78 bytes for IPv6 packets), and increasing in increments up to a "jumbo" packet size of 9216 bytes.

Once the maximum traffic rate without loss is established for a particular frame size, latency through the switch is then calculated. This is done by subtracting the transmit time stamp from the receive time stamp. The minimum, maximum and average latencies for all packets sent and received are calculated and reported. Individual multi-port tests usually run for just a minute or two.

The Spirent test system's load generator was configured to forward traffic to and receive traffic from each port on the switch. This test is configured with a one-to-one, port-to-port traffic mapping – that is, all traffic arriving on one port are all delivered by the switch to the same outbound port, and vice versa. The following results show the maximum forwarding throughput the switch achieves without frame loss. When a switch accepts and successfully processes and forwards all traffic at the maximum theoretical rate of the port, the switch is said to perform at wire speed or full line rate for the particular packet size.

The Catalyst 3850 24S-E was configured for Layer 2 switching as well as Layer 3 forwarding (IP routing for IPV4 and IPV6). 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 28 of the Catalyst 3850 24S-E's fixed ports were connected to the Spirent load-generation system for these tests (26 ports at 1GE, and two ports at 10GE). The result: Traffic delivered to and received from each port on the Catalyst 3850 24S-E could be sent at wire-speed, at all packet sizes, without any data loss.

Configured in this manner, then – with 26 x 1GE ports and two 10GE ports – the Catalyst 3850 24S-E forwarded 68.45 Million packets per second (Mpps), for 64-byte packets, and delivered a cumulative throughput of 46 Gbps – wire-speed on all ports.

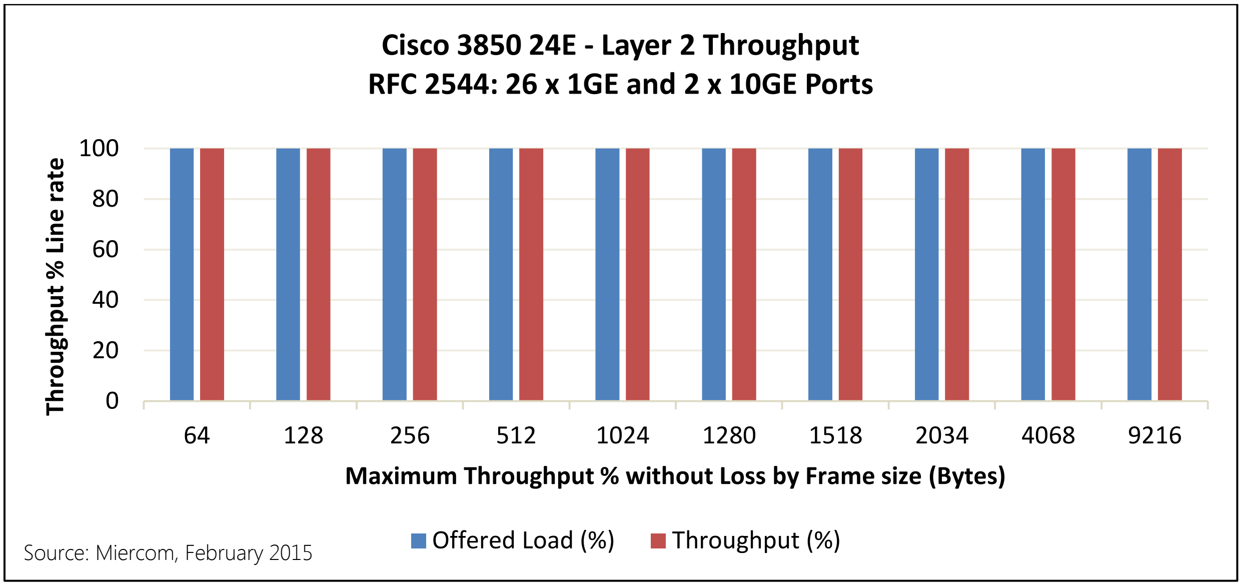

The test-load generator forwarded and received 1-Gbit/s traffic to and from each directly connected 1GE port plus 10-Gbit/s to each 10GE uplink port. The test was conducted for a range of frame sizes including 64, 128, 256, 512, 1024, 1280, 1518, 2034, 4068 and 9216 bytes. The Cisco Catalyst 3850 forwards Layer 2 frames at 100 percent line rate for all frame sizes, achieving a Maximum Forwarding Rate of 68.4 Mpps (Millions of packets per second), for 64-byte packets. The switch can forward frames at wire-speed without any loss.

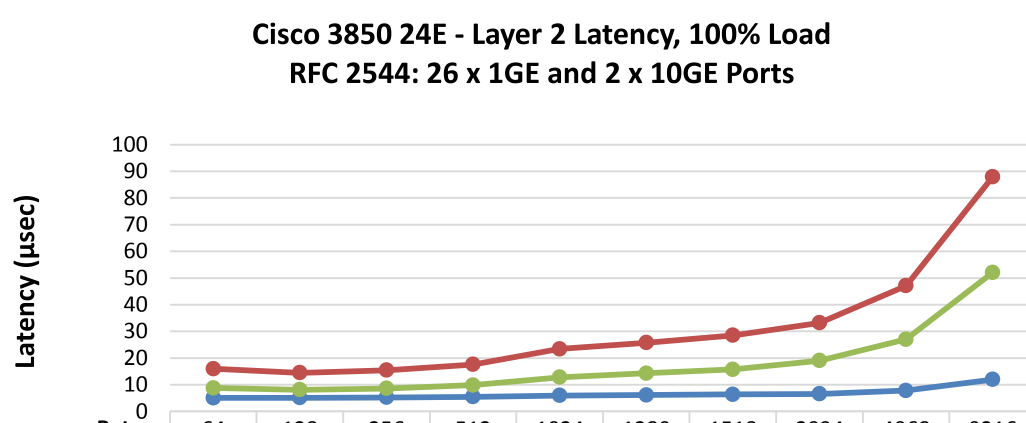

Bytes	64	128	256	512	1024	1280	1518	2034	4068	9216
Minimum Latency (µsec)	5.05	5.1	5.23	5.46	5.93	6.15	6.37	6.55	7.83	11.92
Maximum Latency (µsec)	15.95	14.46	15.45	17.61	23.43	25.74	28.5	33.15	47.11	87.91
Average Latency (µsec)	8.77	8.14	8.57	9.84	12.77	14.28	15.72	19.04	27.02	52.07

Minimum Latency (µsec) Maximum Latency (µsec) Average Latency (µsec)

Source: Miercom, February 2015

As shown in the previous chart, the minimum latency exhibited by the Cisco Catalyst 3850 is exemplary: ranging from just 5 to 6.37 microseconds (µsec) for packets from 64 to 1518 bytes, the normal range of packet size. The latency rose slightly to 11.92 µsec for 9216-byte jumbo packets. The ***average*** *latency ranged from about 9 to 52 µsec for the same packet-size range. The maximum latency by packet size rose to 16 to 88 µsec. For these latency measurements the switch was subjected to a 100-percent traffic load of Layer 2 unicast traffic on all 28 ports. Latency measurements were calculated in accordance with RFC 2544.*

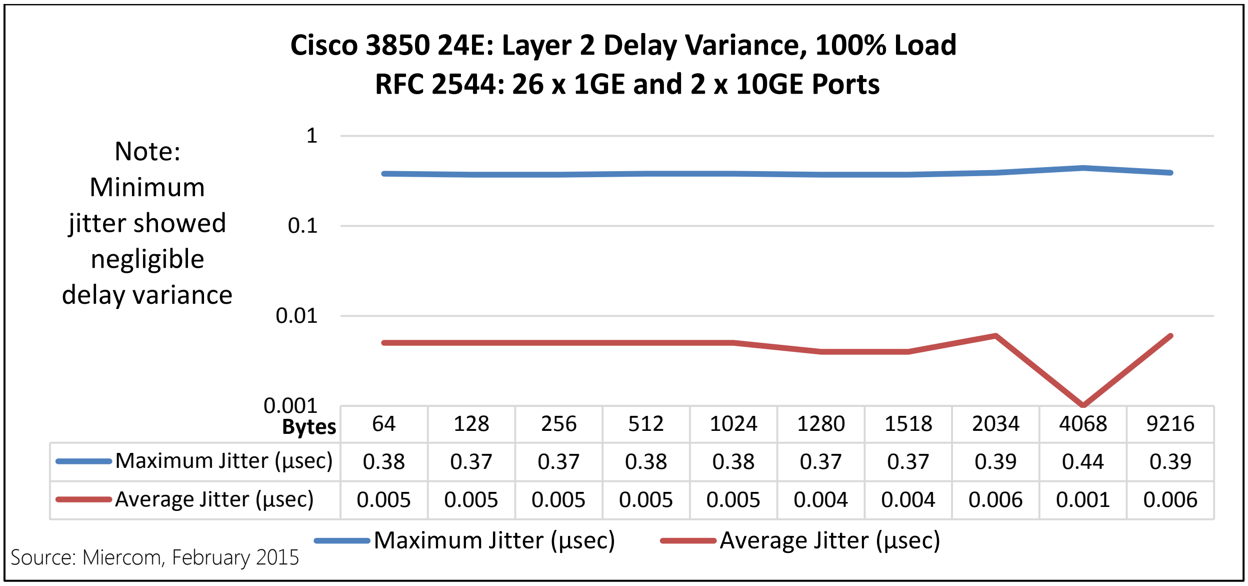

Bytes	64	128	256	512	1024	1280	1518	2034	4068	9216
Maximum Jitter (µsec)	0.38	0.37	0.37	0.38	0.38	0.37	0.37	0.39	0.44	0.39
Average Jitter (µsec)	0.005	0.005	0.005	0.005	0.005	0.004	0.004	0.006	0.001	0.006

The Cisco Catalyst 3850 24E exhibited very little variance in latency, also called jitter – less than one hundredth of one microsecond on average, for all packet sizes up to 9,216 bytes. The Cisco switch configured with 28 ports was subjected to a 100 percent traffic load of Layer 2 traffic. Measurements were in accordance with RFC 3393.

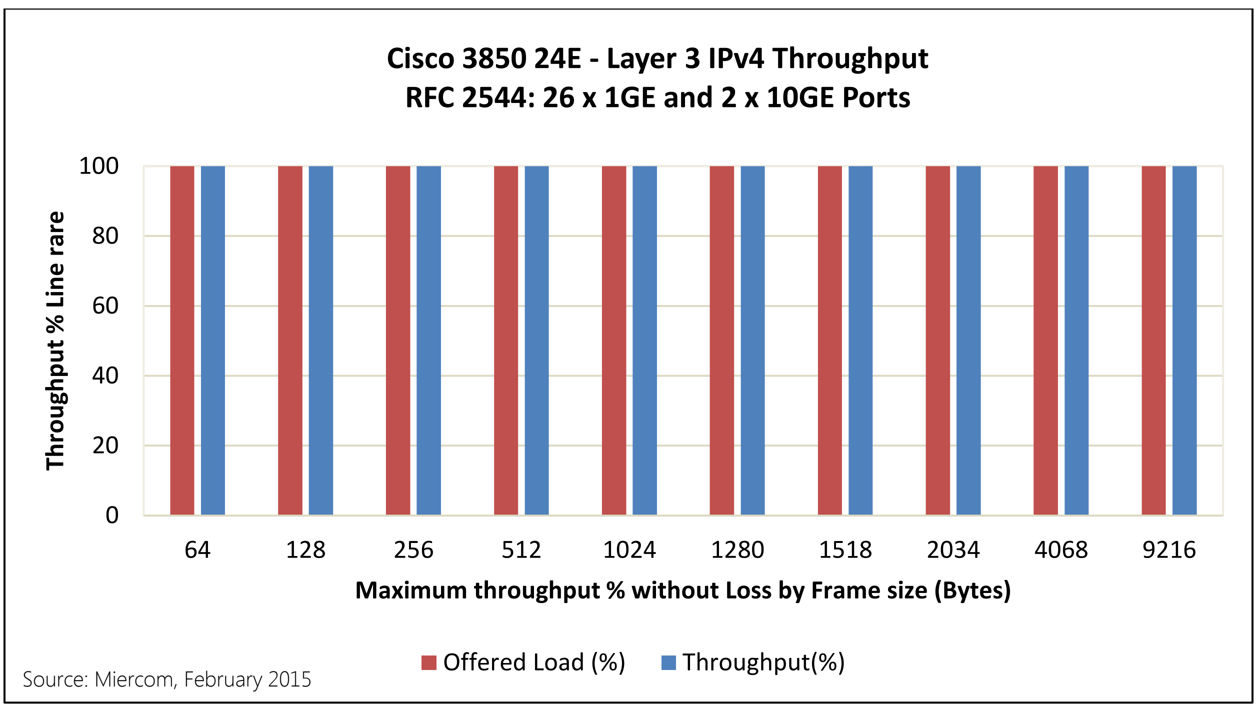

After Layer 2 traffic testing, throughput was measured for IPv4 packet traffic. As before, the test-load generator forwarded and received 1-Gbit/s traffic to and from each directly connected 1GE port plus 10-Gbit/s to each 10GE uplink port. The test was conducted for a range of frame sizes including 64, 128, 256, 512, 1024, 1280, 1518, 2034, 4068 and 9216 bytes. The Cisco Catalyst 3850 forwards IP packets at 100 percent line rate for all frame sizes, achieving a Maximum Forwarding Rate of 68.4 Mpps, for 64-byte packets. The switch can forward frames at wire-speed without any loss.

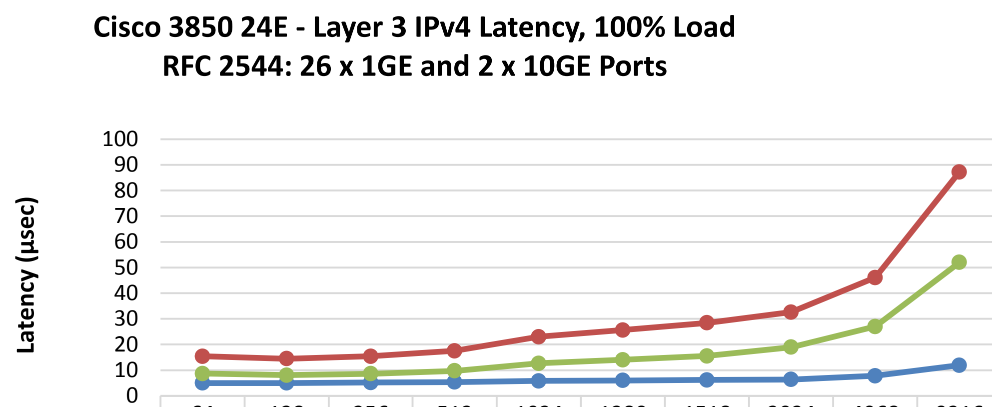

Bytes	64	128	256	512	1024	1280	1518	2034	4068	9216
Minimum Latency (µsec)	5	5.02	5.15	5.37	5.79	5.98	6.18	6.35	7.8	11.91
Maximum Latency (µsec)	15.43	14.49	15.38	17.5	22.99	25.62	28.4	32.6	46.06	87.24
Average Latency (µsec)	8.68	8.06	8.60	9.72	12.68	14.05	15.55	18.93	26.98	52.00

Minimum Latency (µsec) Maximum Latency (µsec) Average Latency (µsec)

Source: Miercom, February 2015

*Interestingly, the latency for maximum-rate Layer 3 packets through the switch is nearly identical to latency for Layer 2 frames. The minimum latency exhibited by the Cisco Catalyst 3850 ranges from just 5 to 6.18 (µsec) for packets from 64 to 1518 bytes, the normal range of packet size. The latency rose slightly to 11.91 µsec for 9216-byte jumbo packets. The **average** latency ranged from about 9 to 52 µsec for the same packet-size range. The maximum latency by packet size rose to 15 to 87 µsec. For these latency measurements the switch was subjected to a 100-percent traffic load of Layer 3 IPv4 unicast traffic on all 28 ports. Latency measurements were calculated in accordance with RFC 2544.*

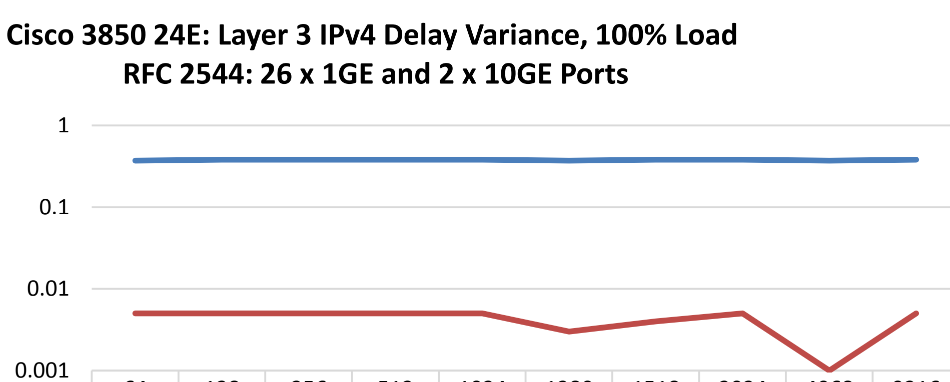

Bytes	64	128	256	512	1024	1280	1518	2034	4068	9216
Maximum Jitter (µsec)	0.37	0.38	0.38	0.38	0.38	0.37	0.38	0.38	0.37	0.38
Average Jitter (µsec)	0.005	0.005	0.005	0.005	0.005	0.003	0.004	0.005	0.001	0.005

Source: Miercom, February 2015

Maximum Jitter (µsec) Average Jitter (µsec)

As with Layer 2 traffic, the Cisco Catalyst 3850 24S-E exhibited very little variance in latency, or jitter – less than one hundredth of one microsecond on average, for all IPv4 packet sizes up to 9,216 bytes. The Cisco switch configured with 28 ports was subjected to a 100 percent traffic load of Layer 3 IPv4 traffic. Measurements were in accordance with RFC 2544.

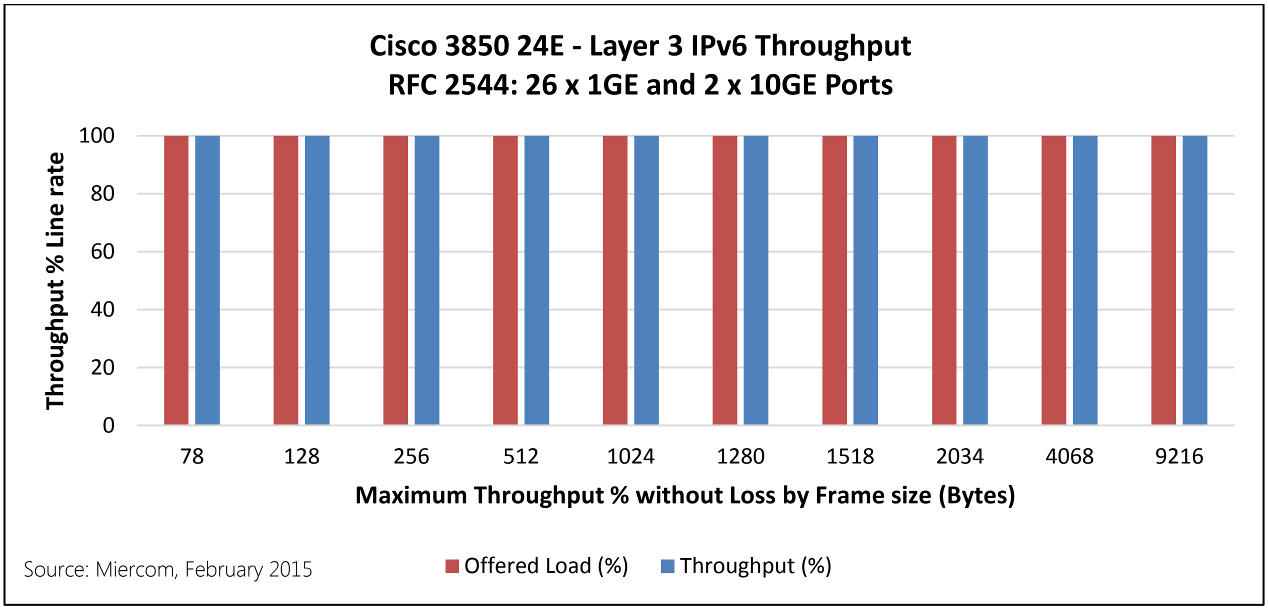

After IPv4 testing, throughput was measured for IPv6 packet traffic. As before, the test-load generator forwarded and received 1-Gbit/s traffic to and from each directly connected 1GE port plus 10-Gbit/s to each 10GE uplink port. The test was conducted for a range of frame sizes including 64, 128, 256, 512, 1024, 1280, 1518, 2034, 4068 and 9216 bytes. The Cisco Catalyst 3850 forwards IPv6 packets at 100 percent line rate for all frame sizes, achieving a Maximum Forwarding Rate of 58.67 Mpps for 78-byte packets. The switch can forward IPv6 packets at wire-speed without any loss.

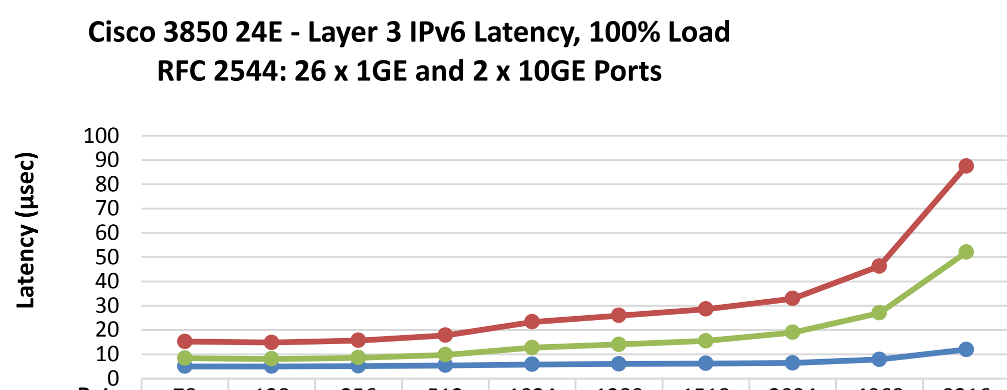

Bytes	78	128	256	512	1024	1280	1518	2034	4068	9216
Minimum Latency (µsec)	4.99	5.03	5.14	5.38	5.79	6	6.18	6.37	7.8	11.9
Maximum Latency (µsec)	15.21	14.8	15.66	17.78	23.31	25.94	28.59	32.94	46.29	87.46
Average Latency (µsec)	8.35	8.06	8.58	9.73	12.69	14.04	15.50	18.93	26.97	51.99

Minimum Latency (µsec) Maximum Latency (µsec) Average Latency (µsec)

Source: Miercom, February 2015

The latency for maximum-rate IPv6 packets through the switch is nearly identical to the latency for IPv4 and Layer 2 traffic. The minimum IPv6 latency exhibited by the Cisco Catalyst 3850 ranges from just 5 to 6.18 microseconds (µsec) for packets from 78 to 1518 bytes, the normal range of packet size. The latency rose slightly to 11.9 µsec for 9216-byte jumbo packets. The average latency ranged from about 8.3 to 52 µsec for the same packet-size range. The maximum latency by packet size rose to 15 to 87 µsec. For these latency measurements the switch was subjected to a 100-percent traffic load of Layer 3 IPv6 unicast traffic on all 28 ports. Latency measurements were calculated in accordance with RFC 2544.

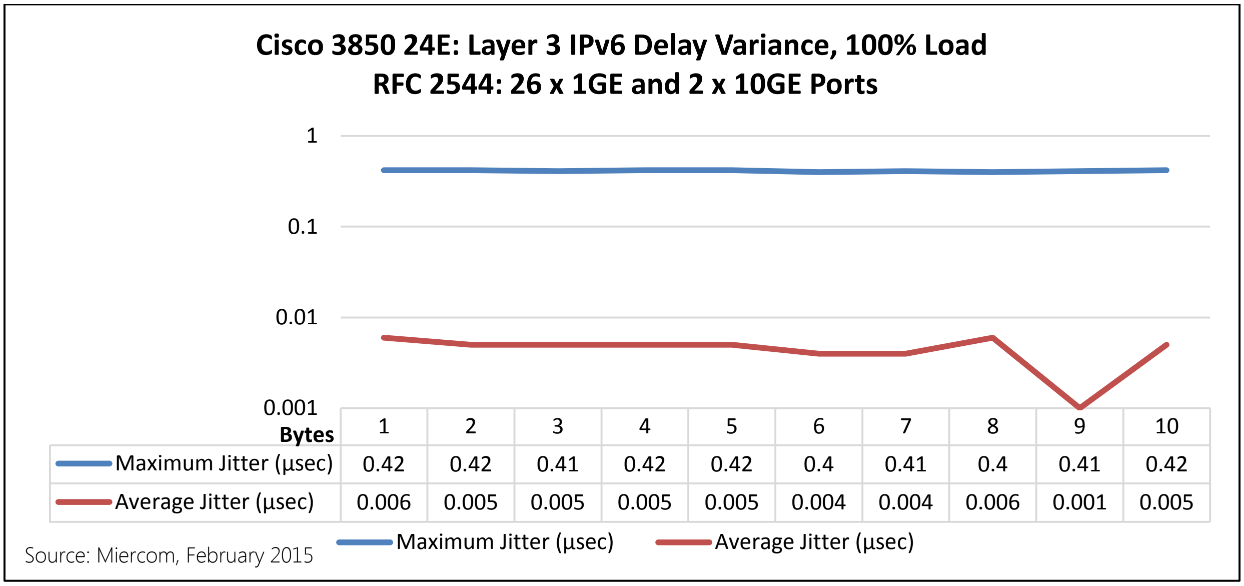

Bytes	1	2	3	4	5	6	7	8	9	10
Maximum Jitter (µsec)	0.42	0.42	0.41	0.42	0.42	0.4	0.41	0.4	0.41	0.42
Average Jitter (µsec)	0.006	0.005	0.005	0.005	0.005	0.004	0.004	0.006	0.001	0.005

As with Layer 2 and IPv4 traffic, the Cisco Catalyst 3850 24E exhibited very little variance in latency, or jitter – less than one hundredth of one microsecond on average, for all IPv6 packet sizes up to 9216 bytes. The Cisco switch configured with 28 ports was subjected to a 100 percent traffic load of Layer 3 IPv6 traffic. Measurements were in accordance with RFC 2544.

5 – Full-Mesh Throughput and Latency Performance RFC 2889

The meshed throughput performance test, as described in RFC 2889, determines the volume of traffic DUT can handle when forwarding packets to all other ports in round-robin "mesh-distribution" fashion. The results show the total number of frames transmitted to, and the total number of frames received on, all ports, for each frame size. The percentage loss of frames for each frame size is also determined and reported.

Mesh testing inherently stresses the switch by forcing traffic through all switch components and processes. The Cisco 3850 was configured for Layer 2 switching. A bidirectional traffic load is used for this test; each port is sending and receiving traffic simultaneously to and from all other ports of the same data rate.

In our testing, the Cisco 3850 was configured to route data received on each of the fixed 24 x 1GE ports to each of the other 23 ports. Therefore 24 ports were simultaneously switching maximum-rate, bidirectional frames to and from each of the other 23 ports in round-robin fashion.

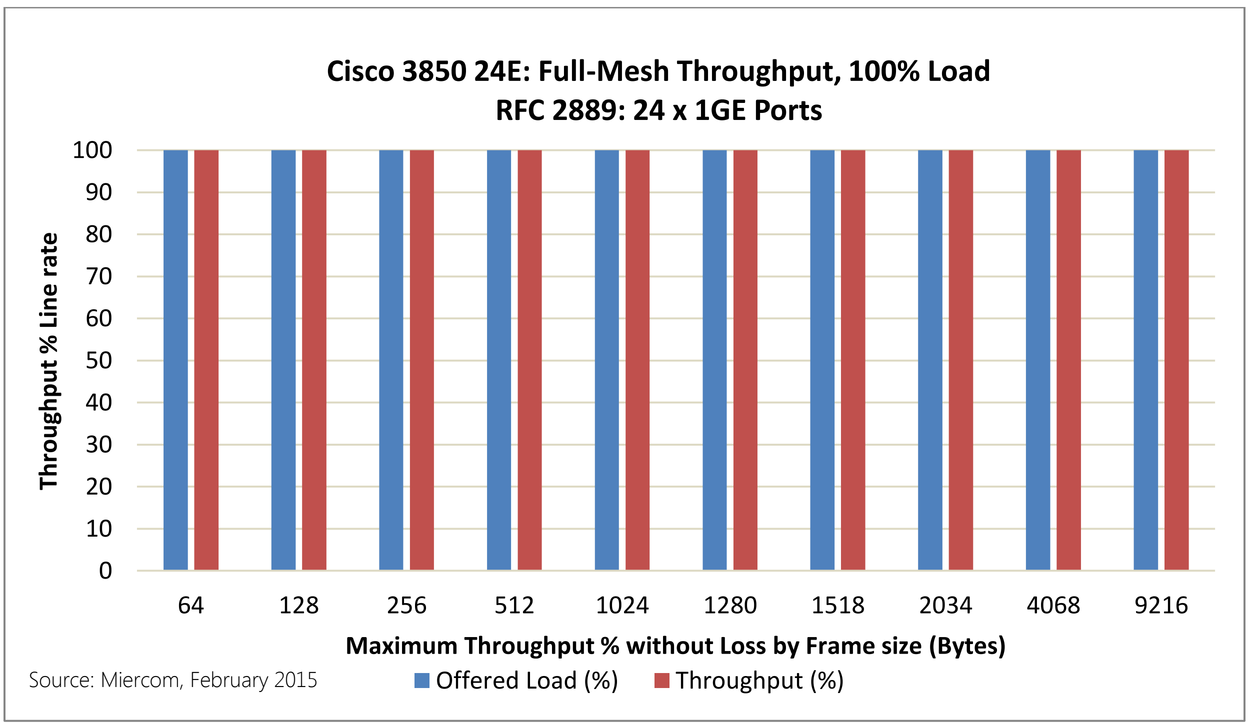

The Cisco Catalyst 3850 24S-E is able to forward full-line-rate traffic, bidirectionally, in full-mesh, round-robin fashion among all 24 fixed 1GE ports with no loss. Tests were conducted in accordance with RFC 2889.

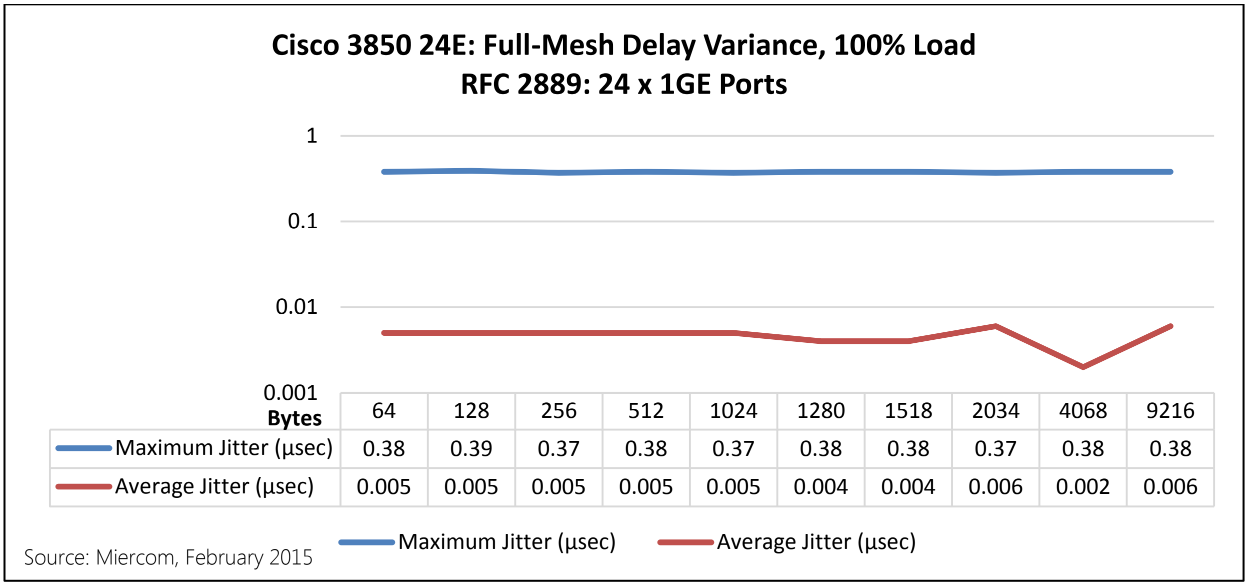

Bytes	64	128	256	512	1024	1280	1518	2034	4068	9216
Maximum Jitter (µsec)	0.38	0.39	0.37	0.38	0.37	0.38	0.38	0.37	0.38	0.38
Average Jitter (µsec)	0.005	0.005	0.005	0.005	0.005	0.004	0.004	0.006	0.002	0.006

In full-mesh testing the Cisco Catalyst 3850 24S-E exhibited very little variance in latency, or jitter – less than one hundredth of one microsecond on average – for all packet sizes up to 9216-byte. The Cisco switch configured with 24 x 1GE ports was subjected to a 100 percent traffic load. Layer 2 traffic was applied bidirectionally for each specified frame size.

6 - IP Multicast Throughput and Latency RFC 3918

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. It is also used in applications such as disseminating a real-time IP video feed to multiple recipients.

RFC 3918 describes the procedure for conducting a multicast forwarding and latency test. In multicasting, IP packets received on one port are replicated by the switch or router and forwarded to multiple output ports.

The Spirent test equipment has a standard program for conducting such multicast testing, measuring the forwarding throughput, and the minimum, maximum and average latency of multicast frames. The multicast frames are sent to clients on multiple subnets/ports, which are also configured via the test and measurement system.

The test results reveal how much processing overhead the DUT requires, on average, to replicate and forward multicast frames. The tester defines the multicast control protocol (our testing used IGMPv2; PIM-SM), 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 is then calculated from the results.

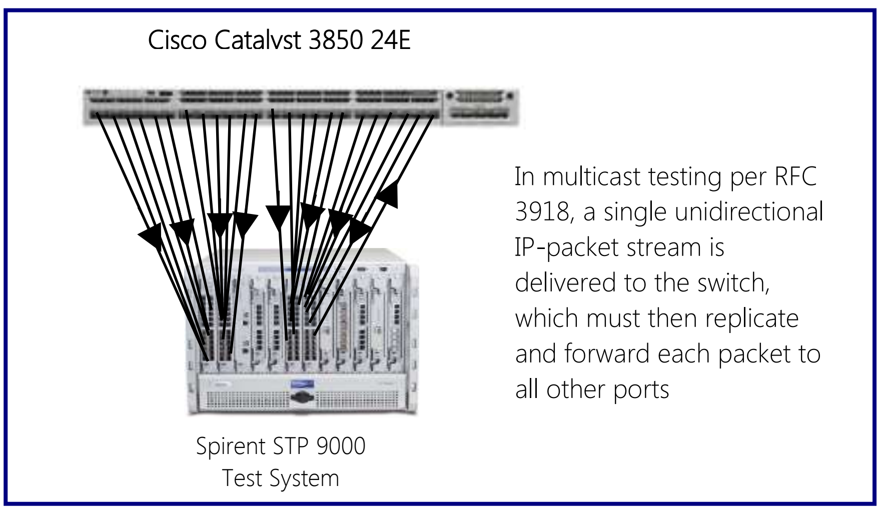

Unlike the other testing conducted up to this point, multicast traffic is one-way: Packets delivered to one port are replicated and distributed to the other 23 fixed 1GE ports. The total throughput, then, is the total replicated packets sent out by the 23 output ports. Latency is measured as before – by examining the sent timestamp that the test system places within test frames. The test system compares this timestamp, in frames received from the ports, with the current time and calculates the difference. The test system records the average, minimum and maximum latencies for each multicast group, in microseconds.

Multicast testing was conducted first using conventional IPv4 packets. Then the testing was repeated using IPv6 traffic. The results were the same: All packets in a line-rate IP stream were replicated and distributed to all other ports with no loss.

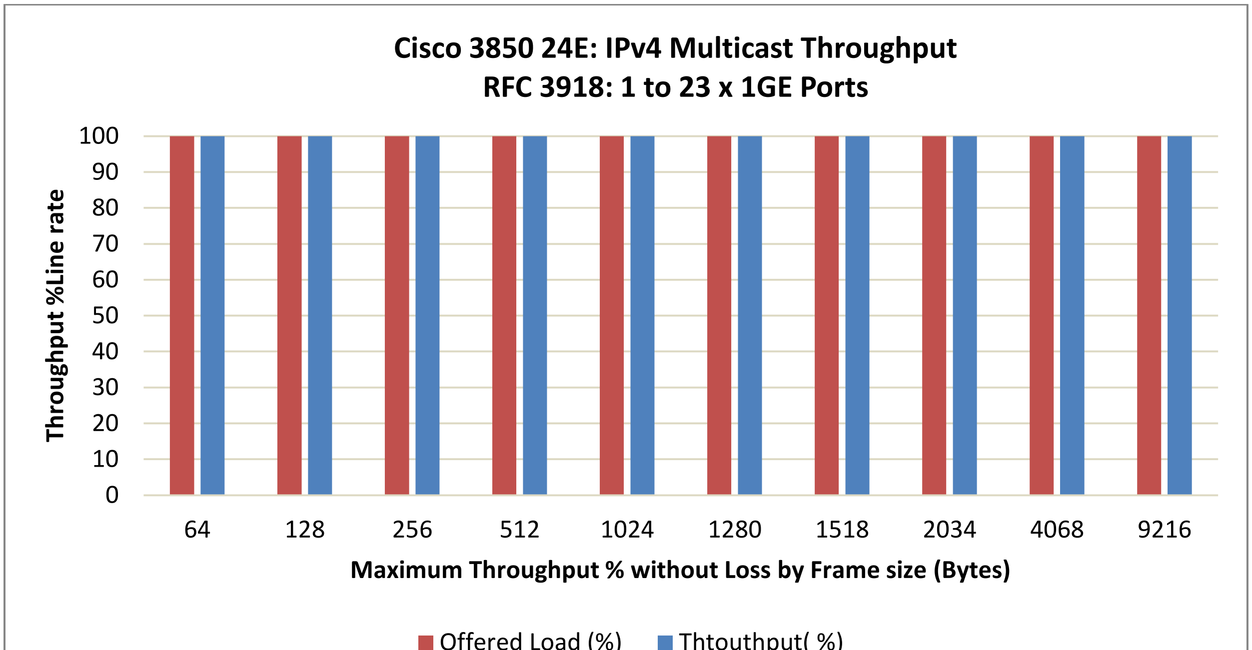

The Cisco Catalyst 3850 24S-E with 24 1GE SFP ports exhibits full-line-rate multicast handling – packet replication and forwarding of IPv4 traffic – from one input port to all 23 other 1GE ports. These tests were conducted in accordance with RFC 3918.

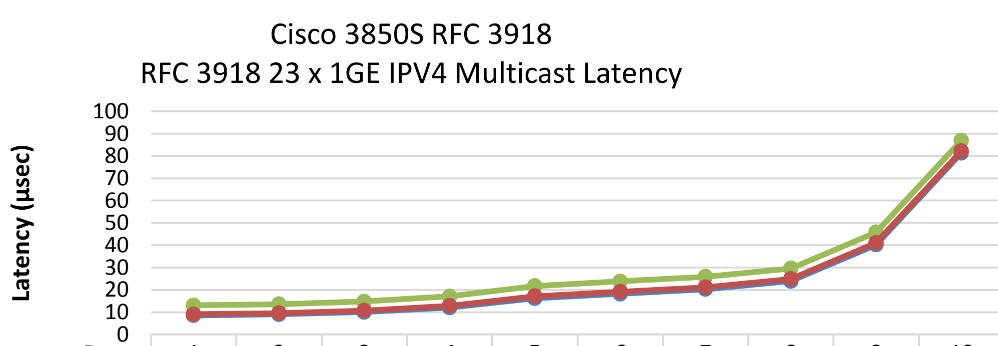

Bytes	1	2	3	4	5	6	7	8	9	10
Minimum Latency (µsec)	8.50	9.00	10.04	12.06	16.20	18.23	20.26	23.90	40.18	81.37
Maximum Latency (µsec)	13.05	13.57	14.78	17.09	21.70	23.76	25.85	29.60	45.77	86.95
Average Latency (µsec)	9.05	9.57	10.69	12.82	17.11	19.11	21.11	24.83	41.04	82.22

Minimum Latency (µsec) Maximum Latency (µsec) Average Latency (µsec)

Source: Miercom, February 2015

Even with a 100-percent input load, the Cisco Catalyst 3850 24S-E exhibits relatively low latency for IPv4 multicast traffic distribution across all of its 1GE ports. From the smallest packet size of 64 bytes to the largest 9216-byte jumbo frame, average latency ranges from 9.05 to 82.22 microseconds (µsec). These latency results are for IP multicast distribution of a line-rate packet stream to 23 x 1GE ports. These tests were conducted in accordance with RFC 3918.

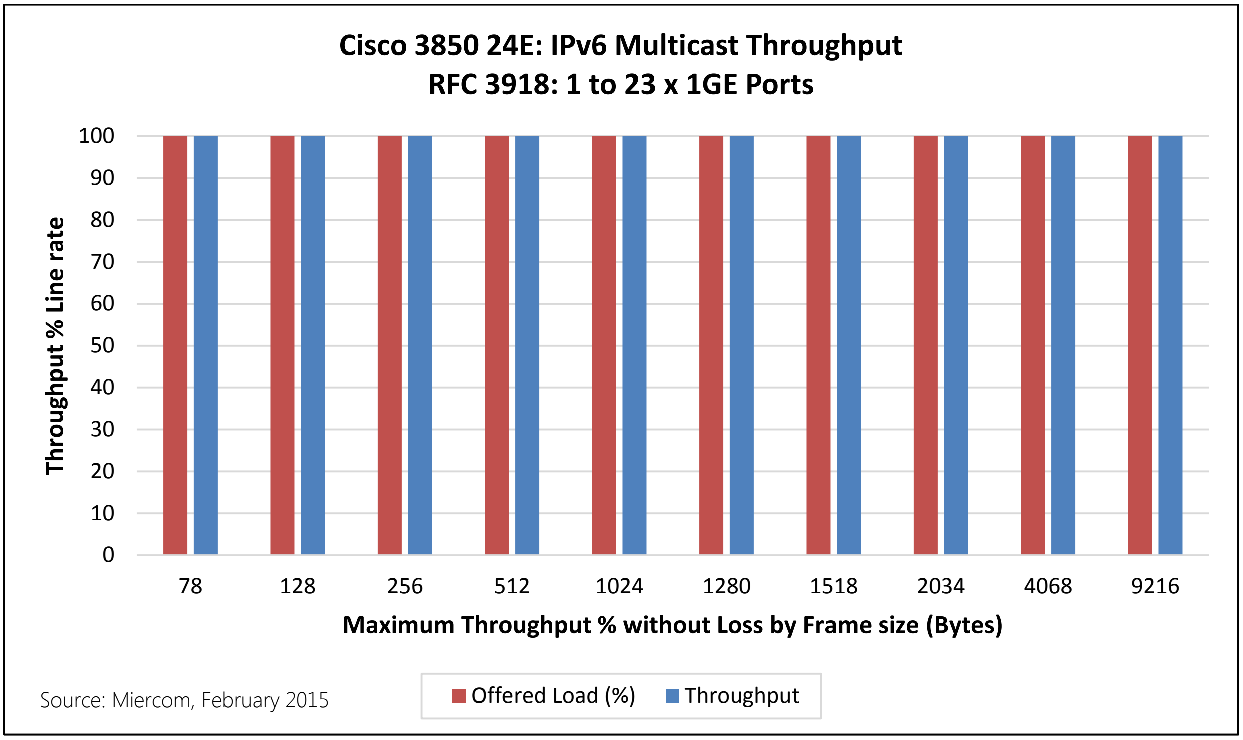

The Cisco Catalyst 3850 24S-E with 24 1GE SFP ports exhibits full-line-rate multicast handling – packet replication and forwarding of IPv6 traffic – from one input port to all 23 other 1GE ports. These tests were conducted in accordance with RFC 3918.

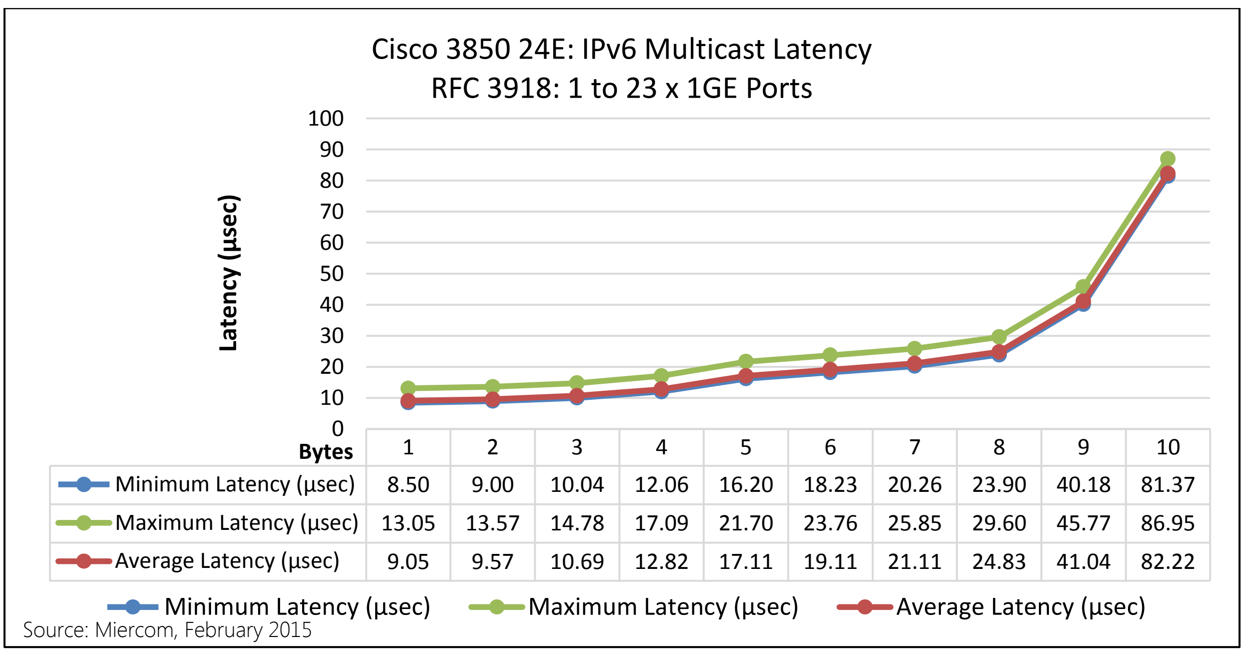

Bytes	1	2	3	4	5	6	7	8	9	10
Minimum Latency (µsec)	8.50	9.00	10.04	12.06	16.20	18.23	20.26	23.90	40.18	81.37
Maximum Latency (µsec)	13.05	13.57	14.78	17.09	21.70	23.76	25.85	29.60	45.77	86.95
Average Latency (µsec)	9.05	9.57	10.69	12.82	17.11	19.11	21.11	24.83	41.04	82.22

As with IPv4 multicast traffic handling, the Cisco Catalyst 3850 24S-E exhibits relatively low latency for IPv6 multicast traffic distribution across all of its 1GE ports. From the smallest packet size of 78 bytes to the largest 9216-byte jumbo frame, average latency ranges from 9.2 to 82.22 µsec. These latency results are for IP multicast distribution of a line-rate IPv6 packet stream to 23 x 1GE ports. These tests were conducted in accordance with RFC 3918.

7 - Stack Convergence Test

Stacking Ethernet switches provides the network administrator with three major operational benefits:

- Single point of management, since all switches in the stack are managed as one.
- Near instantaneous – on the order of just 37 microseconds – fail-over to alternate data path if the primary stack data path is interrupted.
- Built-in redundancy and high availability: The high-speed Stackwise 480 connections provide a redundant communication path for each stack member to every other member.
- Scalable to fit network needs: Adding another new switch to the stack is simple. As the need for additional access ports grows, adding a new switch to an existing stack is easier and faster than adding a new standalone switch to the network.

Stacking Diagram

Rear view

Front View

Cisco 3850 fiber-based switches can be stacked with Cisco 3850 copper-based switches, since the rear stacking connections are unrelated to the switches' actual port connectors. Cisco StackWise-480 technology delivers a stack bandwidth of 480 Gbps. StackWise-480 also features Cisco IOS-software-based Stateful Switch-Over (SSO) for providing resiliency within the stack if a switch fails or a link is lost. Additionally, the stack behaves as a single switching unit that is managed by an "active" switch elected by the member switches.

The stack-convergence test was performed to see how quickly the stack can recover and load-balance when one link is removed from the stack's ring-topology connectivity. To do this, four Cisco 3850 switches (fiber-based, copper-based, and PoE switches) were stacked in a typical full-ring topology. A 10-Gbps uplink port on the top switch was used as source port, and a 10-Gbps uplink port on the bottom switch was used as a receive port.

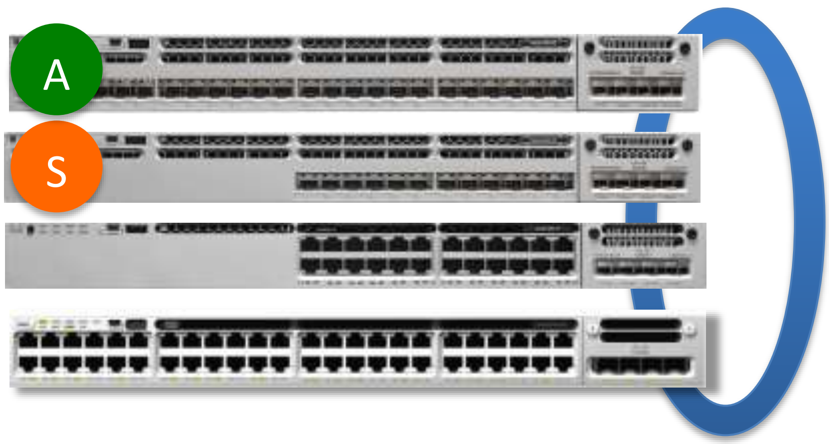

Traffic sent through the first switch is load-balanced across the stack ports. And there are redundant paths to allow traffic to circumvent an intermediate switch outage or link loss. With full-rate bidirectional 10-Gbps traffic running across the stack, from source port on the top switch to the receive port on the bottom, we disconnected the stack cable between Switch 1 and Switch 4 and measured the amount of frame loss that occurred before traffic was rerouted. A 9,216-byte jumbo packet size was used for the test traffic stream.

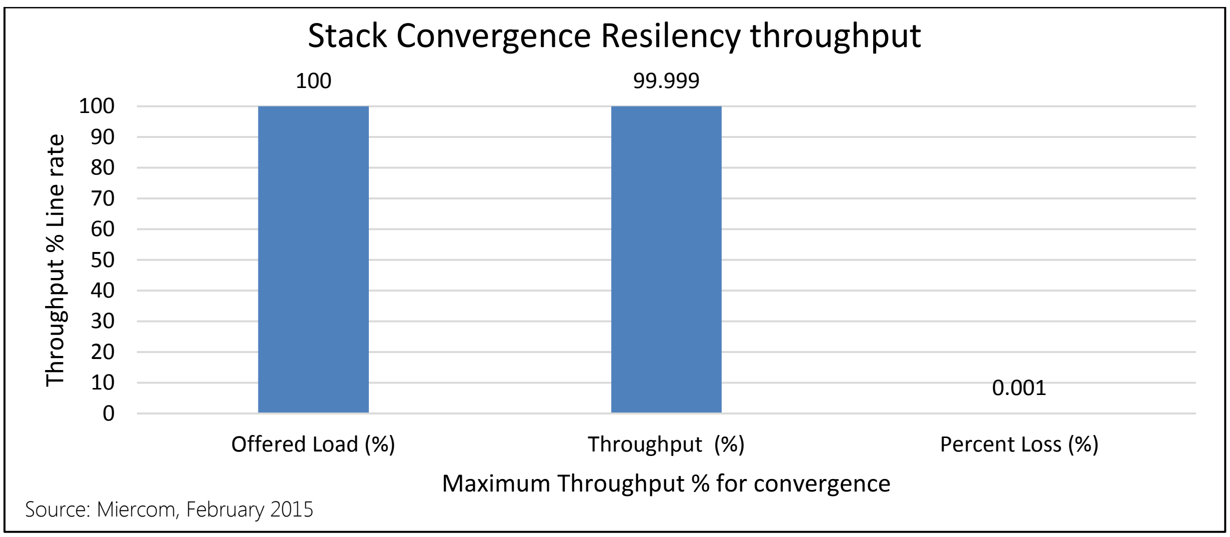

The Cisco Catalyst 3850 24S-E showed quick convergence – on the order of just 37 microseconds – when a stack link was intentionally disconnected. The test was conducted while a 10-Gbps bidirectional test traffic stream was being sent from the top stacked switch to a port on the bottom switch. The test was run for one minute, during which time the link was disconnected and traffic was re-routed. During the interruption just five of the jumbo packets were lost, out of 8.12 million packets sent. The loss represents 0.001 percent of the data sent during the one-minute test.

8 - About this Miercom Testing

Miercom has published hundreds of network-product-comparison analyses in leading trade periodicals and other publications. Miercom's reputation as the leading, independent product test center is undisputed.

Private test services available from Miercom 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.

9 - 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.