

Dell Networking Z9500 Fabric Switch DR140402G



Contents

i. Executive Summary	3
ii. About the Dell Networking Z9500 10/40GE Fabric Switch	5
iii. Test Bed Setup	6
iv. How We Did It	8
1 Throughput and Latency Performance Tests (RFC 2544 and RFC 3393)	9
1.1 RFC 2544 Throughput Test System and DUT Configuration	9
1.2 RFC 2544 Cut-Through 40GE Layer 2 Throughput	10
1.3 RFC 2544 Store-and-Forward 40GE Layer 2 Throughput.....	11
1.4 RFC 2544 Store and Forward 2 x 40GE X-Fabric Layer 3 Latency - 100 percent load	12
1.5 RFC 2544 Store and Forward 2 x 40GE Local Switching Layer 2 Latency - 100 percent load	13
1.6 RFC 2544 Store and Forward 2 x 40GE Local Switching Layer 3 Latency - 100 percent load	14
1.7 RFC 2544 Cut-Through 40GE Layer 2 Latency – 100 percent load	15
1.8 RFC 2544 Store-and-Forward 40GE Layer 2 Latency – 50 percent load	16
1.9 RFC 2544 Store-and-Forward 40GE Layer 2 Latency – 100 percent load	17
1.10 RFC 3393 40GE Layer 2 Delay Variance - 100 percent load	18
2 Fully Meshed Throughput and Latency Test (RFC 2889)	19
2.1 RFC 2889-based Fully Meshed Test System and DUT Configuration.....	19
2.2 RFC 2889 40GE Layer 3 Throughput	20
2.3 RFC 2889 Cut-Through 40GE Layer 2 Latency – 100 percent load	21
2.4 RFC 2889 Cut-Through 40GE Layer 3 Latency with TCP – 100 percent load.....	22
2.5 RFC 2889 Cut-Through 40GE Layer 3 with TCP – 100 percent load	23
2.6 RFC 2889 Store and Forward 40GE Layer 3 Latency with TCP – 100 percent load ..	24
3 RFC 3918 Layer 3 Multicast Throughput	25
3.1 RFC 3918 Multicast Throughput Configuration	25
3.2 RFC 3918 Layer 3 Multicast Throughput	26
3.3 RFC 3918 Multicast Layer 3 Latency	27
3.4 RFC 3918 Multicast Layer 3 Latency – Packets Traversing Two Line Cards and Three ASICs between Ports	28
3.5 RFC 3918 Layer 3 Group Join Delay and Group Leave Delay	29
3.6 RFC 3918 Multicast Group Capacity	30
4 Overnight Attrition – Burn In – Packet Loss.....	31
5 Dell Z9500 Power Consumption	32
6 Competitive Power Consumption Analysis.....	34
7 About the Miercom 40GE Switch Industry Study.....	35
8 About Miercom.....	35
9 Use of This Report.....	35

i. Executive Summary

Miercom was engaged by Dell Networking to independently test its latest data-center switch – the Z9500 10/40 Gigabit Ethernet Fabric Switch, the high-end model in the vendor's Z Series and the highest-density, highest-throughput switch produced by Dell Networking to date.

Introduced in March 2014, the Z9500 supports up to 132 x 40GE ports – four times the capacity of the Z9000, the other member of the Z Series. The Z9500 delivers 10.4 Tbps (Terabits/s) full duplex of switching throughput and a non-blocking switching fabric more than four times more than the predecessor Z9000.

Previous testing of a Dell Networking switch occurred in 2013. This was the S6000 10/40 Gigabit Ethernet Top-of-Row/End-of Rack switch, the top model in the vendor's S Series and the first supporting high density 40GE ports.

The tests of both the Z9500 and S6000 are part of our ongoing evaluation of switches that support 40GE – the Miercom 40GE Switch Industry Assessment.

We tested the Z9500 in both store-and-forward and cut-through operating modes. The earlier S6000 was successfully tested in store-and-forward mode only; cut through mode was not available at the time of testing. Dell states that cut-through mode has since been released.

Our testing of the Z9500 primarily addressed performance and energy efficiency.

For performance testing, traffic was directed across the switching fabric. The results of the following performance tests are included here:

- RFC 2544 Layer 2 throughput and latency operating in store-and-forward mode
- RFC 2889 Layer 2 and 3 throughput and latency operating in both store-and forward mode and cut-through mode
- RFC 3393 Layer 2 jitter (latency variance) operating in store-and-forward mode

Among the key test results:

- In store-and-forward mode, the Z9500 transmits Layer 2 packets of all sizes (64 to 9,216 bytes) at full line-rate on all 132 x 40GE ports, with zero loss and with low latency.
- Operating in both store-and-forward and cut-through modes, the Z9500 readily transmits Layer 3 packets of all sizes (70 to 9,216 bytes) at full line-rate with zero loss and with low latency.
- The Z9500 exhibits impressively low Layer 2 jitter (latency variance) while operating in store-and-forward mode.
- The switch consumes a comparatively respectful 15.65 watts per 40GE port, with all 132 x 40GE ports handling a random mix of Layer 2 packet sizes, 64 to 9,216 bytes, at 100 percent full load.
- Experienced zero loss of Layer 2 minimum-size (64-byte) packets during a 15-hour test that applied traffic across the switching fabric on all 132 x 40GE ports.

The review also verified these characteristics of the Z9500:

- The 3RU (three rack units, about 5 ¼ inches) chassis houses a switching fabric and 11 IO modules. Each IO module supports 12 x 40GE ports.
- The switching fabric includes six Broadcom Trident II chipsets while each IO module has one – a system total of 17 Trident II chipsets.
- Load balancing is achieved by dynamically directing traffic across the switching fabric.
- The switch functions perfectly with only two of its four load-balancing, hot-swappable power supplies in operation.
- Internal sensors monitor the temperature to ensure the switch stays within operating range, and dynamically vary the speed of the cooling fans to save energy.

Miercom independently substantiates the impressive performance and energy efficiency of the Dell Networking Z9500 10/40 Gigabit Ethernet Fabric Switch exhibited in testing. The Z9500 is awarded the **Miercom Performance Verified Certification** for operation in both store-and-forward and cut-through modes.

Robert Smithers
CEO
Miercom



ii. About the Dell Networking Z9500 10/40GE Fabric Switch

The Dell Networking Z9500 is a compact Layer 2 and Layer 3 Ethernet switch/router designed for 10/40GE aggregation in all types of data centers – such as traditional, cloud, hyperscale and provider-hosted.

Real-world deployment scenarios of the Z9500 include:

- For active-fabric 10/40GE switching in enterprise, high-performance computing (HPC) and cloud computing data centers, where 1/10GE servers require maximum bandwidth for maximum performance.
- As a switching device for non-blocking Clos (multistage) architectures in hyperscale data centers.
- As a high-density, end-of-row switch for blade server aggregation.
- As a small-scale, active-fabric spine switch, along with Dell Networking 1/10GE S Series top-of-rack switches for cost-effective aggregation of 10/40GE uplinks.

Introduced in March 2014, the Z9500 is the high-end model in the Dell Networking Z Series, joining the Z9000. The 3RU Z9500 is the highest-density, highest-throughput switch offered to date by Dell Networking, supporting up to 132 x 40GE (QSFP) ports or up to 528 x 1/10GE (SFP+) ports with breakout cables.

The “pay as you go” approach to scaling up 40GE ports is a key selling feature. A Z9500 chassis can be purchased with a license for 36 or 84 x 40GE ports. Additional licenses can be purchased to increase the 40GE port count to the maximum of 132 x 40GE ports.

This approach gives network administrators the flexibility to scale costs and switching infrastructure capacity as future requirements demand. An initial purchase can handle the current traffic level of the enterprise network and, as network load grows, as it inevitably does, it is quick and easy to upgrade the 40GE port license.

The Z9500 delivers high performance – 10.4 Tbps (Terabits/s) of full-duplex switching I/O bandwidth with a non-blocking switching fabric – in a compact 3U form factor that conserves rack space. Full Layer 2 and Layer 3 switching functionality include Quality of Service (QoS) traffic prioritization and a complement of IPv4 and IPv6 features, including support for OSPF (Open Shortest Path First) and BGP (Border Gateway Protocol) routing control protocols.

Support for Layer 2 Multipath is provided via Virtual Link Trunking (VLT), a Dell Networking proprietary Layer 2 link aggregation protocol. VLT offers servers connected to different access switches a redundant, load-balancing connection to the network core in a loop-free environment with benefits surpassing the aging Spanning Tree Protocol.

The Z9500 also supports Multi-domain Virtual Link Trunking (mVLT), a proprietary Dell Networking technology for multi-chassis link aggregation that allows multiple VLT domains to be linked with a VLT link aggregation group (LAG).

iii. Test Bed Setup

Testing of the Dell Z9500 employed the Ixia XG12, the highest port density test system available for Ethernet. The XG12 can be used for Layer 1-to-7 testing and features a 12-slot modular chassis. A chassis can support up to 384 x 10 Gigabit Ethernet (GE) ports, 192 x 1GE ports, 20 x 100GE Higher Speed Ethernet (HSE) CFP MSA ports, 48 x 40GE HSE QSFP+ ports, 24 Packet over SONET (POS) ports or 24 Asynchronous Transfer Mode (ATM) ports.

The centrally controlled XG12 test system is shown below. In testing of the Z9500 the XG12 was configured with 40GE ports.

To generate load on the full 132 x 40GE Z9500 ports, the Dell switch was configured to route traffic through all the ports chained together (traffic sent out on egress ports was delivered back into the switch on ingress ports). For cross-fabric tests, VLANs were configured in the Z9500 to route traffic across the fabric between ports.

Test System:

Ixia XG12 multi-slot chassis with line cards for 24 ports



24 x 40GE
QSFP+
ports

Device under Test (DUT):

Dell Z9500 10/40GE Fabric Switch

3RU chassis with fabric and 11 fixed Network Processor Units, each of which houses 12 40GE ports



132 x 40GE
QSFP+ ports



Once connectivity between the Ixia XG12 and the Dell Z9500 was confirmed, we did not have to re-cable for the duration of the testing.

Each of the tests conducted is discussed individually in this report. Collectively they examined the performance and energy efficiency of the Dell Z9500, referred to throughout this report as the Device Under Test, or DUT.

The procedures for testing Layer 2 and Layer 3 switches and routers have been standardized in recent years, and this testing employed a number of those procedures. Naturally the high number and speed of switch ports and the incredible volumes of test traffic needed to fully load them are orders of magnitude more today for data center switches than just a few years ago. Even so, the same procedures still apply.

Three of the standards we applied in this testing, which the Ixia IxNetwork QuickTests incorporate, are publicly available as Internet Requests for Comments, or RFCs. The ones applicable in this testing include RFC 2544 and 2889 for throughput and latency measurements of Layer 2 and Layer 3 traffic.

RFC 2544, issued in 1999, delineates 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 Ixia solution provides a fully meshed bidirectional traffic flow for these measurements, which fully stresses the switch fabric.

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

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.

The Ixia hardware is controlled by the IxNetwork software application. IxNetwork is used for network topology testing and traffic analysis. Additionally, it contains many built in QuickTests, which can be used to execute the RFC tests mentioned above to test the performance of the Dell Z9500.

iv. How We Did It

The tests were conducted in stages, starting with two Ixia ports connected to two Dell switch ports. The remaining 130 ports were connected in a line with fiber optic cables and via a single VLAN, to yield a cut-through configuration for our initial testing. This did not go through the switch fabric's IO modules.

In the next stage we tested again with two Ixia ports, but reconfigured the Dell switch to have many VLANs that criss-crossed the line cards, thereby routing the forwarded traffic through the six fabric IO modules.

For the RFC 2889 full-mesh testing, we first connected 12 Ixia XG12 ports to 12 Z9500 ports, all on one IO module, in order to validate connectivity. Frames were forwarded through the switch in the full-mesh traffic pattern. Each port sent and received traffic from each of the other 11 ports to which it was connected.

After this configuration was successfully validated, we connected 24 Ixia ports to 24 Z9500 ports – 12 each on two IO modules – and ran full-mesh traffic. Each port sent traffic to, and received it from, the other 23 ports to which it was connected.

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 assistance.

1 Throughput and Latency Performance Tests (RFC 2544 and RFC 3393)


This test determines the maximum rate at which the Dell Z9500 switch receives and forwards traffic without frame loss. The rate at which MAC or IPv4 frames are sent is progressively increased to determine the maximum rate at which the switch does not lose frames. A spectrum of frame sizes was used for testing the Dell Z9500.

Once the maximum 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, IxNetwork 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.

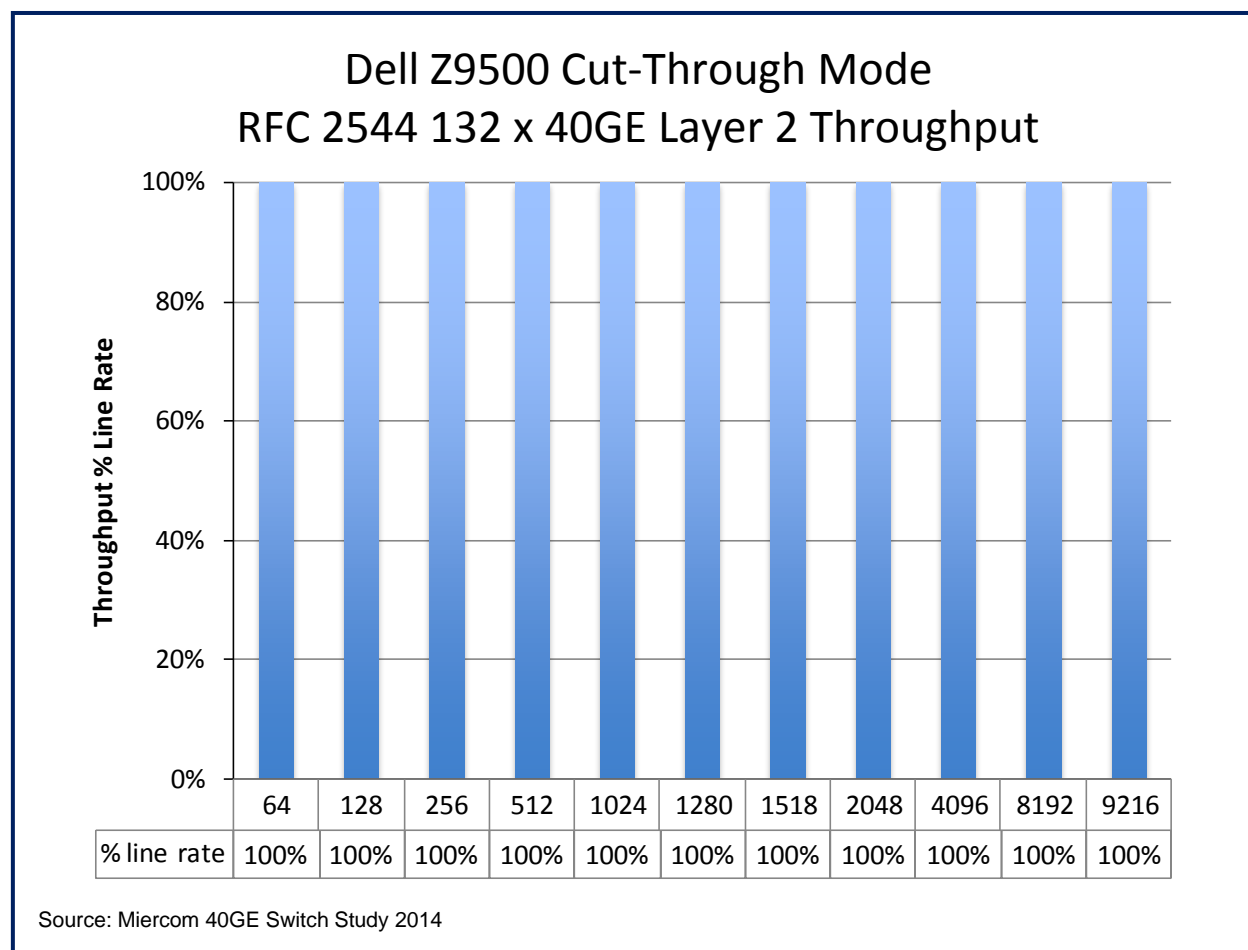
The Dell Z9500 was configured for Layer 2 switching in store-and-forward mode. Port-pair combinations were assigned within the test system so that bidirectional traffic was transmitted between line-card ports across the switching fabric modules, in accordance with RFC 2544. All 24 of the Dell Z9500's 40GE ports were connected to the Ixia load-generation system for this test. The test system established that traffic delivered to and received from each port on the Dell Z9500 could be sent at "wire speed," at all tested packet sizes, with zero data loss. The packet sizes ranged from 64 to 9,216 bytes.

1.1 RFC 2544 Throughput Test System and DUT Configuration



1.2 RFC 2544 Cut-Through 40GE Layer 2 Throughput

“Operating in cut-through mode with all 132 40GE ports handling a 100 percent Layer 2 traffic load, the Dell Z9500 delivered line-rate performance for all frame sizes with zero loss.”



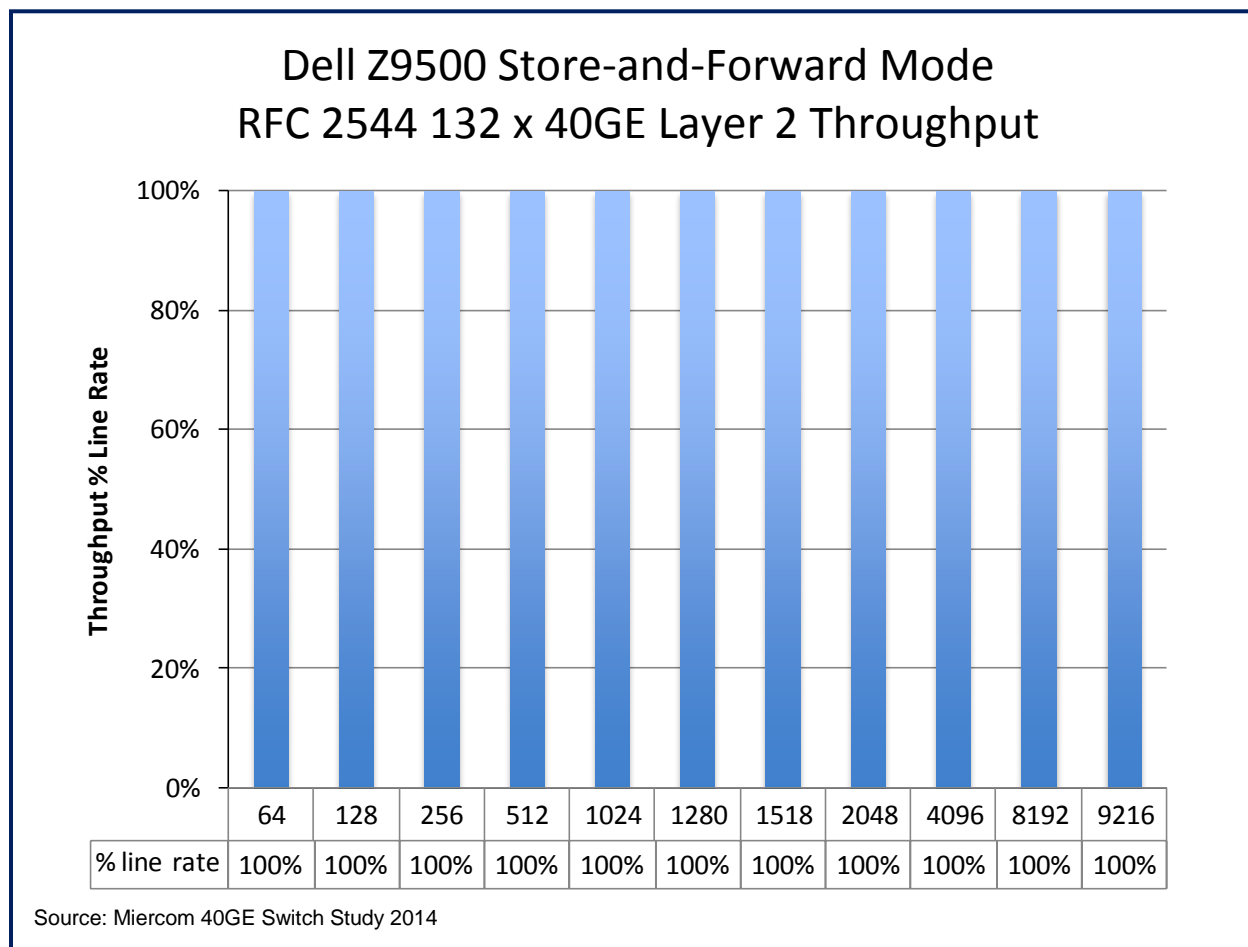
The Dell Z9500 exhibited full line-rate forwarding performance for all packet sizes in this throughput test, which was conducted in accordance with RFC 2544. No frame loss occurred as the Z9500 operated in cut-through mode, with all 132 40GE ports fully loaded with Layer 2 traffic.

Observations

Maximum throughput was verified for the full range packet sizes, ranging from 64 to 9,216 bytes. The Dell Z9500 could handle full line-rate traffic on all ports for all frame sizes without incurring any loss.

1.3 RFC 2544 Store-and-Forward 40GE Layer 2 Throughput


“Operating in store-and-forward mode, the Dell Z9500 forwarded frames of all sizes at full line-rate with zero loss while handling a 100-percent Layer 2 traffic load on all 132 x 40GE ports.”



The Dell Z9500 exhibited full line-rate forwarding performance for all packet sizes in this throughput test, conducted in accordance with RFC 2544. No frame loss occurred as the Z9500 operated in store-and-forward mode, with all 132 x 40GE ports fully loaded with Layer 2 traffic.

1.4 RFC 2544 Store and Forward 2 x 40GE X-Fabric Layer 3 Latency - 100 percent load


"The Dell Z9500 with 2 ports at 100 percent load with traffic between 2 IO modules (cross fabric) exhibits the following latency for all packet sizes tested during RFC 2544-based latency tests."



Average latency for the Dell Z9500 x-fabric store and forward mode ranged from 1.65 to 3.50 μ sec.

1.5 RFC 2544 Store and Forward 2 x 40GE Local Switching Layer 2 Latency - 100 percent load

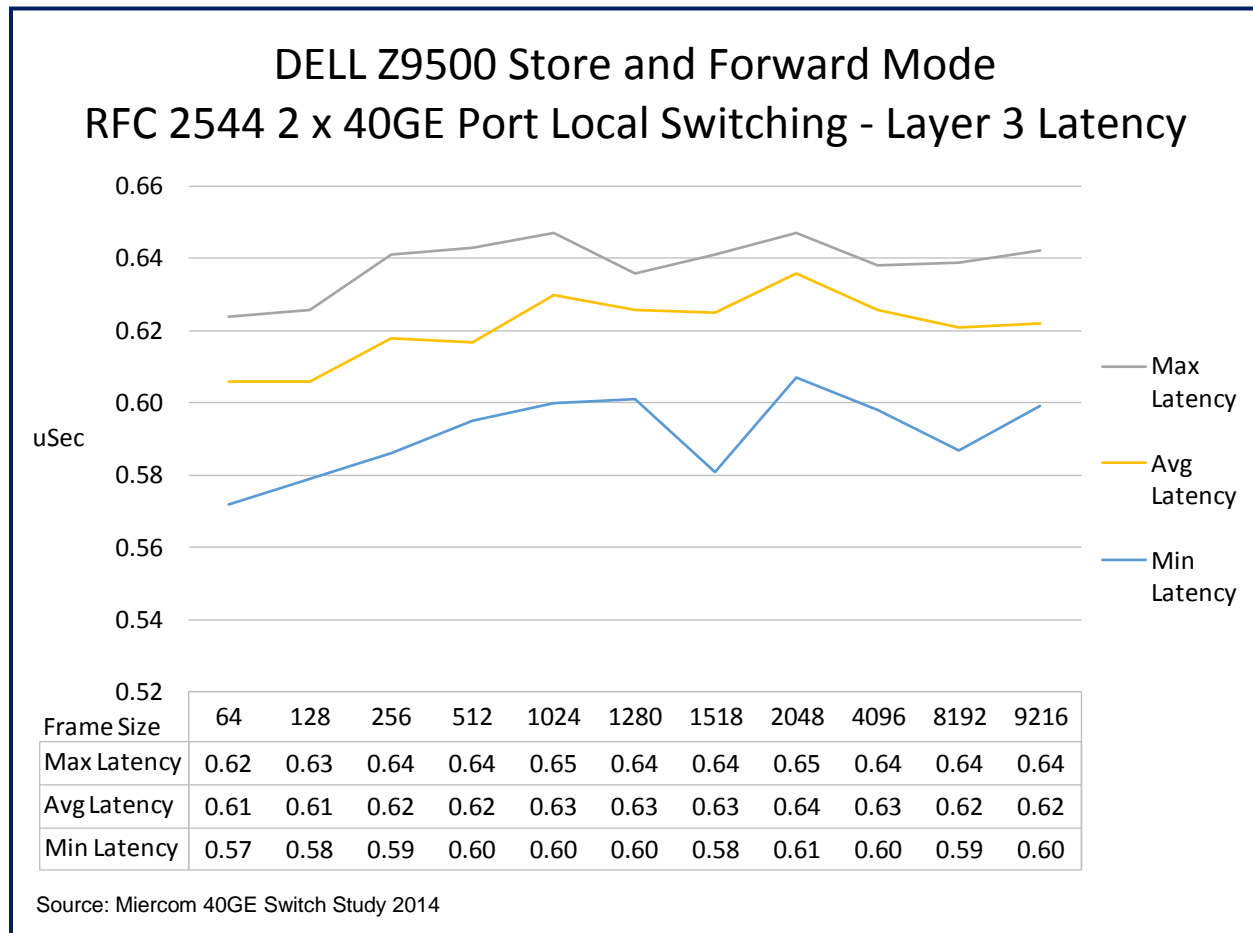
"The Dell Z9500 with 2 ports at 100 percent load with traffic locally switched through 1 IO module exhibits the following latency for all packet sizes tested during RFC 2544-based latency tests."



Average latency for the Dell Z9500 in locally switched Layer 2 store and forward mode ranged from only 0.60 to 0.64 μ sec.

1.6 RFC 2544 Store and Forward 2 x 40GE Local Switching Layer 3 Latency - 100 percent load


"The Dell Z9500 with 2 ports at 100 percent load with traffic locally switched through 1 IO module exhibits the following latency for all packet sizes tested during RFC 2544-based latency tests."



Average latency for the Dell Z9500 in locally switched Layer 3 store and forward mode ranged from only 0.61 to 0.64 μ sec.

1.7 RFC 2544 Cut-Through 40GE Layer 2 Latency – 100 percent load

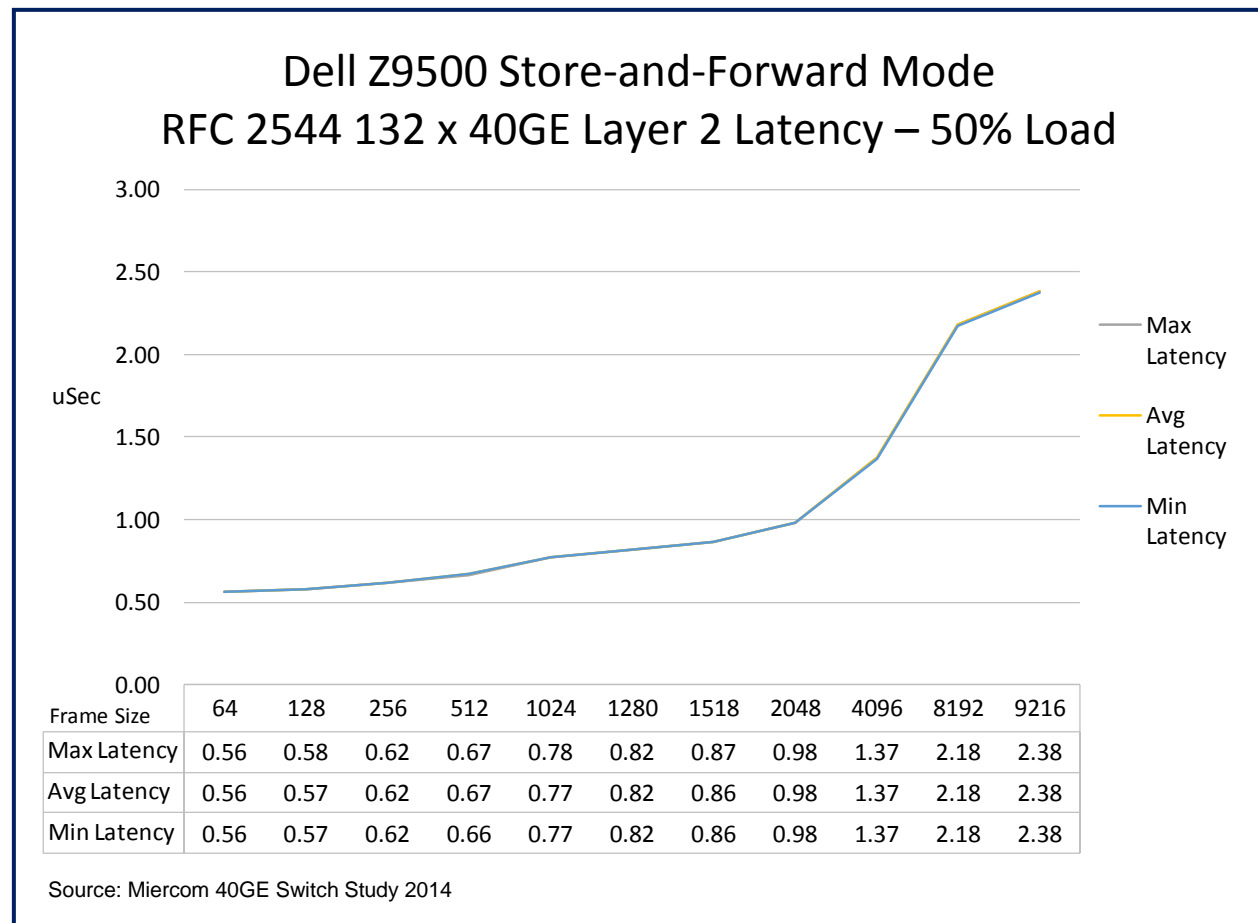
"The Dell Z9500 with 132 ports at 100 percent load exhibits the following latency for all packet sizes tested during RFC 2544-based latency tests."



Average latency for the Dell Z9500 in cut-through mode ranged from only 0.51 to 0.60 μ sec.

1.8 RFC 2544 Store-and-Forward 40GE Layer 2 Latency – 50 percent load


"Operating in store-and-forward mode with all 132 x 40GE ports handling a 50 percent traffic load, the Dell Z9500 exhibited consistently low average latency for all frame sizes."



The chart shows minimum, average and maximum latency for each frame size validated in RFC 2544 testing of the Dell Z9500, operating in store-and-forward mode and handling a 50 percent load of Layer 2 traffic on all 132 ports. For all frame sizes, average latency was consistently low, ranging from a minimum of 0.56 μ sec for the smallest frame size, 64 bytes, to 2.38 μ sec for the largest 'jumbo' frame size, 9,216 bytes. Also, for all frame sizes, 0.01 μ sec was the largest difference between minimum and maximum latency.

1.9 RFC 2544 Store-and-Forward 40GE Layer 2 Latency – 100 percent load


"Operating in store-and-forward mode with all 132 40GE ports handling a 100 percent Layer 2 traffic load, the Dell Z9500 exhibited consistently low average latency for all frame sizes."



The chart show minimum, average and maximum latency for each frame size validated in RFC 2544 testing of the Dell Z9500 operating in store-and-forward mode and handing a 100 percent load of Layer 2 traffic on all 132 ports. For all frame sizes, average latency was consistently low, ranging from a minimum of 0.58 μ sec for the smallest frame size, 64 bytes, to 2.39 μ sec for the largest, 9,216 bytes. Also for all frame sizes, 0.02 μ sec was the largest difference between minimum and maximum latency. This is why the min, average and max latencies all lie on the same line. Latency variance is virtually nil.

1.10 RFC 3393 40GE Layer 2 Delay Variance - 100 percent load

"The Dell Z9500 operating in store-and-forward mode with all 132 40GE ports handling a 100 percent Layer 2 traffic load exhibits little delay variance (jitter) for all frame sizes."



The Dell Z9500 with eleven 12 Port 40G line cards 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 bytes. The Dell switch configured with 132 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.

Observations

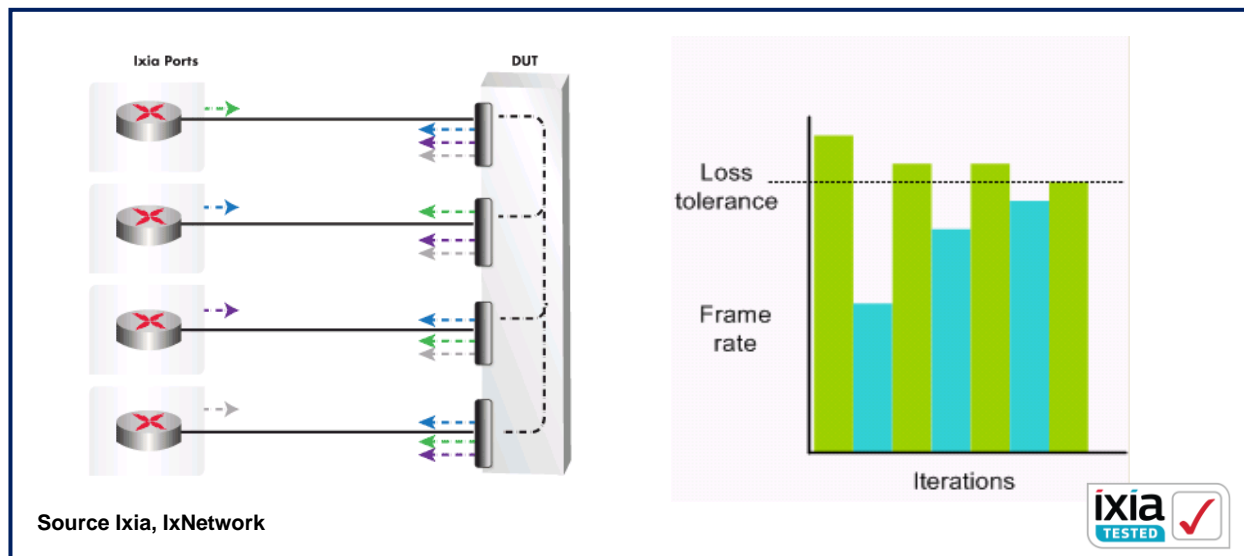
Delay variance and latency were tested for different packet sizes in the range of 64 bytes to 9,216 bytes. The Dell Z9500 could handle full line-rate traffic on all ports, at all frame sizes, with very little latency.

2 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, 24 ports of the Ixia test system are connected, and traffic flows of fixed packet sizes are sent in a mesh-distribution fashion between each active port and the other 23 active ports. 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).

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



2.2 RFC 2889 40GE Layer 3 Throughput


“The Dell Z9500 with 24 of 132 ports running Full Mesh at 100 percent line-rate throughput had zero packet drops.”



The Dell Z9500 had 12 connected ports each on two separate line cards running full mesh traffic over 24 ports from the Ixia test system. The switch fabric exhibits full line-rate forwarding performance with ports fully loaded without loss for Layer 3 (IP) traffic for all frame sizes. Testing was conducted in accordance with RFC 2889.

2.3 RFC 2889 Cut-Through 40GE Layer 2 Latency – 100 percent load

"The Dell Z9500 with 24 ports across 2 IO modules exhibits consistent low latency for all packet sizes tested during RFC 2889-based latency tests, with the switch fully loaded on all active ports."

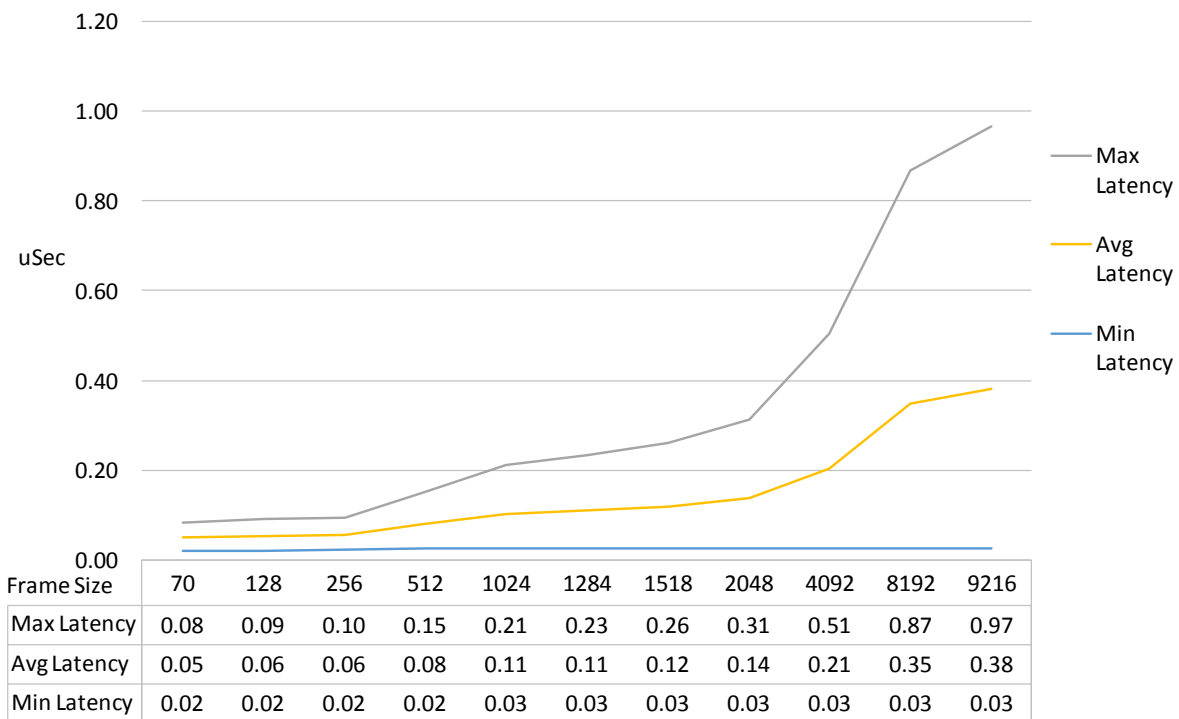


The Dell Z9500 with 24 active ports with fully meshed traffic, exhibits consistent low latency up through 2048-byte packets on all connected ports. From the smallest packet size of 64-bytes to the largest of 9,216-bytes, the average latency ranges from 0.05 to 0.42 microseconds. Tests were conducted in accordance with RFC 2889.

2.4 RFC 2889 Cut-Through 40GE Layer 3 Latency with TCP – 100 percent load

"The Dell Z9500 with 24 ports across 2 IO modules 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."

Dell Z9500 Cut-Through Mode
RFC 2889 24 x 40GE Layer 3 Latency with TCP – 100% Load




Source: Miercom 40GE Switch Study 2014

The Dell Z9500 with 24 active ports with fully meshed traffic in cut-through, exhibits consistent low latency up through 2048-byte packets on all connected ports. From the smallest packet size of 70-bytes to the largest 9,216-bytes, the average latency ranges from 0.05 to 0.38 μ secs. Tests were conducted in accordance with RFC 2889.

2.5 RFC 2889 Cut-Through 40GE Layer 3 with TCP – 100 percent load


"The Dell Z9500 with 24 Ports across 2 IO modules 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 24 ports."



The Dell Z9500 with 24 active ports with fully meshed traffic in cut-through, exhibits consistent low latency up through 2048-byte packets on all connected ports. From the smallest packet size of 64-bytes to the largest 9,216-bytes, the average latency ranges from 0.05 to 0.35 μ secs. Tests were conducted in accordance with RFC 2889.

2.6 RFC 2889 Store and Forward 40GE Layer 3 Latency with TCP – 100 percent load

"The Dell Z9500 with 24 Ports across 2 IO modules 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 24 ports."



Dell Z9500 with 24 active ports with fully meshed traffic in store and forward mode, exhibits consistent low latency up through 2048-byte packets on all connected ports. From the smallest packet size of 70-bytes to the largest 9,216-bytes, the average latency ranges from 0.05 to 0.31 μ secs. Tests were conducted in accordance with RFC 2889.

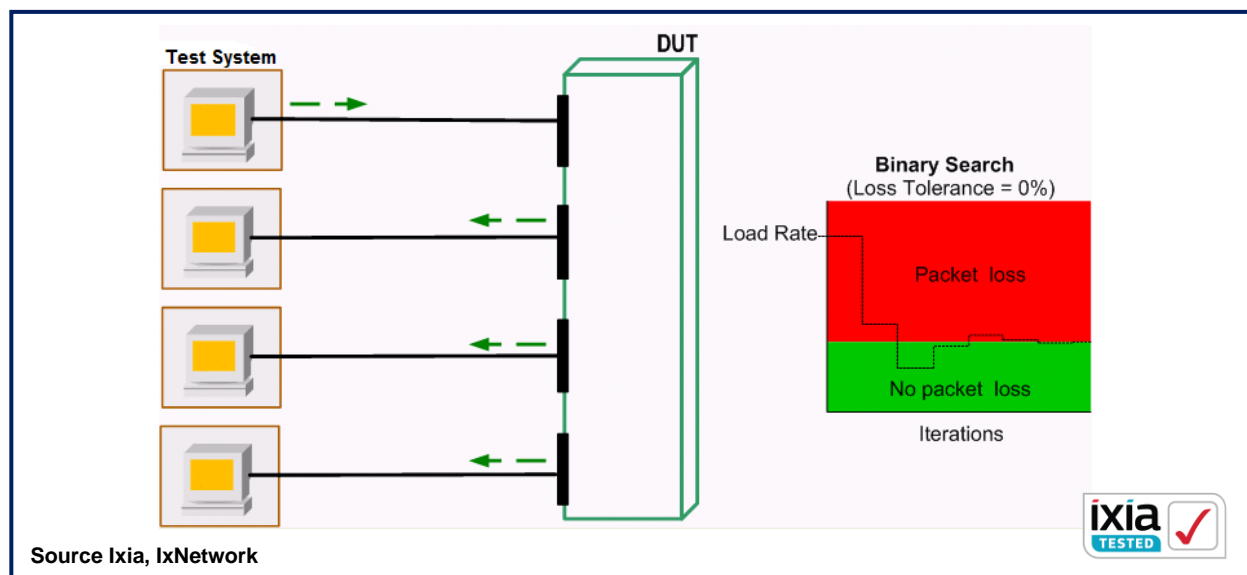
3 RFC 3918 Layer 3 Multicast Throughput

The objective of the RFC 3918 Layer 3 Multicast Throughput Test is to validate the maximum rate of Layer 3 IPv4 multicast traffic that can be handled by the Z9500.

A binary or a linear search to detect maximum loss free loads and group capacity and one-to-many traffic mapping, on a minimum of three ports is required by the RFC. This test was run in accordance with the RFC as shown in the figure below. IGMP snooping, based on the IGMPv2 protocol, was enabled on the Z9500 switch in order to learn multicast groups and their members.

Multicast Traffic was tested using four ports. 1,024 Multicast groups were used and all packets were transmitted from one port on the Z9500 and received on three other ports. The IxAutomate application running on the Ixia XM12 injected IGMPv2 multicast traffic.

3.1 RFC 3918 Multicast Throughput Configuration




The Z9500 successfully transmitted traffic to all multicast member ports at 100% line rate for frame sizes ranging from 64 to 9252 bytes.

Testing verified that the switch was capable of snooping the multicast groups and then properly transmitting multicast traffic at 100% line rate with zero loss to each multicast group member.

3.2 RFC 3918 Layer 3 Multicast Throughput


"The Dell Z9500 achieved 100% line-rate throughput for all frame sizes."



The Dell Z9500 switch exhibited line rate throughput for Layer 3 IPv4 traffic using RFC 3918 standard tests. The Z9500 achieved 100% line rate throughput for Layer 3 IPv4 multicast traffic for all frame sizes from 64 bytes to 9252 bytes. An Ixia XM12 conducted the test on the Z9500.

3.3 RFC 3918 Multicast Layer 3 Latency

"Consistent and low latency was observed with 2048-byte frames."




The Dell Z9500 Switch exhibited consistently low latency values in the RFC 3918 Multicast Latency Test. Average latency ranged from a low of 0.65 μ s for 64- and 128-byte frames to a high of 0.77 μ s for 2048-byte frames. The Z9500 was configured in store-and-forward mode. An Ixia XM12 using RFC standard benchmark suites conducted the test of the switch.

A latency value was captured for each frame size utilized in RFC 3918 Layer 3 throughput testing. The Z9500 switch exhibited low latency for all frame sizes tested.

Latency was determined in two ways: 1. Using Tx and Rx ports on one line card (as above), and 2. on different line cards as shown on the next page. The difference is that on one line card, one ASIC is used to route packets from any port to another. When a port-to-port route traverses two line cards, the packet is processed by three ASICs, one for each line card and one to route between the line cards. Therefore the latency going through two line cards would be expected to be approximately three times that of packets traversing one line card.

3.4 RFC 3918 Multicast Layer 3 Latency – Packets Traversing Two Line Cards and Three ASICs between Ports

"Dell Z9500 had an average latency that ranged from 2.24 to 4.11 microseconds."



As measured, the latency difference varied from about three times for smaller frames to six times for the larger packets.

3.5 RFC 3918 Layer 3 Group Join Delay and Group Leave Delay

The Group Join Delay Test determines how long it takes a switch to register multicast clients to a new or existing group in its forwarding table. The duration between the time a switch receives a group of IGMP/MLD Join requests and the time the multicast clients begin receiving traffic for the groups they joined is measured. The impact of different frame sizes on the duration is recorded.

The Group Leave Delay Test determines how long it takes a switch to remove a client from its multicast table. The duration between the time a switch receives a group of IGMP/MLD Leave requests and the time the multicast clients stop receiving traffic for the groups they left is measured. The impact of different frame sizes on the duration is recorded.

RFC 3918 Multicast Group Join Delay Test

Frame Size (Bytes)	Group Join Delay (ns)
64	12
128	18
256	32
512	58
1024	110
1280	136
1518	160
2048	214
4096	423
8192	841
9252	949

"In the Group Join Delay Test, the Z9500 exhibited a gradual increase in latency that corresponded with the increase in frame size."

RFC 3918 Multicast Group Leave Delay Test


Frame Size (Bytes)	Avg Leave Delay (sec)
64	4.99
128	4.47
256	6.52
512	3.91
1024	4.24
1280	4.41
1518	6.22
2048	4.44
4096	4.50
8192	3.63
9252	4.57

"In the Group Leave Delay Test, the Z9500 exhibited a fractional decrease in latency as the frame size increased."

With all three receivers subscribed to 1,024 multicast groups, the Group Join Delay latency of the Z9500 ranged from 12 to 949 nanoseconds compared to an average Group Leave Delay latency of 4.72 seconds.

3.6 RFC 3918 Multicast Group Capacity

"The maximum multicast group capacity of the Z9500 proved to be the vendor-stated figure of 8,000."



Group capacity was measured by the XM12 by changing the number of multicast groups at a given frame size at line rate multicast traffic to find the maximum number that can sustain 0% packet loss.

4 Overnight Attrition – Burn In – Packet Loss

"A 15-hour cross-fabric burn-in test was conducted. Two 40GE ports on the Ixia test system were configured to route traffic through all 132 40GE ports of the Dell Z9500. The test was completed successfully with zero frames lost."

Test Results

15Hour-64 Byte Across Fabric Snake Test.rtf		
Name	10.0.0.1:01.01	10.0.0.1:01.02
Link State	Link Up	Link Up
Line Speed	40GE	40GE
Frames Sent	3,256,785,971,792	3,256,785,971,792
Frames Sent Rate	0	0
Valid Frames Received	3,256,785,971,792	3,256,785,971,792
Valid Fraames Received Rate	0	0
Bytes Sent	208,434,302,194,688	208,434,302,194,688
Bytes Sent Rate	0	0
Bytes Received	208,434,302,194,688	208,434,302,194,688
Bytes Received Rate	0	0
Fragments	0	0
Undersize	0	0
Oversize	0	0
CRC Errors	0	0
Vlan Tagged Frames	0	0
DROPPED	0	0
Flow Control Frames Received	0	0
Oversize and CRC Errors	0	0
User Defined Stat 1	0	0
User Defined Stat 2	0	0
Capture Trigger(UDS 3)	3,256,785,971,792	3,256,785,971,792
Capture Filter(UDS 4)	3,256,785,971,792	3,256,785,971,792
User Defined Stat 5	0	0
User Defined Stat 6	0	0
Protocol Server Transmit	0	0
Protocol Server Receive	0	0
Transmit Arp Reply	0	0
Transmit Arp Request	0	0
Transmit Ping Reply	0	0
Transmit Ping Request	0	0


Observations

During the test, the Dell Z9500 registered temperature rises, which triggered increased fan speed in order to keep the temperature below 90 degrees Fahrenheit and within the operating range.

5 Dell Z9500 Power Consumption

We measured power consumption externally to the Z9500 at the PDU into which the Dell Z9500 was connected. Power consumption was measured at each stage after letting it settle for 10 minutes in each. The baseline was measured with the Ixia test box and the Dell Z9500 up and running, but with no traffic. We then pulled the plug on two of the four Z9500 power supplies and took another baseline. Then with four power supplies plugged in, we created traffic across 132 ports at increasing frame size, and let each settle for 10 minutes and recorded these measurements.

Dell Z9500 Power Consumption



Two power supplies are more than sufficient to operate the Z9500 switch, with no degradation. Four power supplies are supported and we found that the Dell switch consumed the same wattage at the baseline whether two or four power supplies were being used. Four power supplies add redundancy and automatically load balance power requirements. For a randomly selected packet size, the switch drew 2,066 Watts.

Dell Z9500 Power Consumption per port

For a random frame size traffic mix ranging from 64 to 9,216 bytes forwarding across all 132 ports, the power draw per 40GE port was 15.65 watts and per 10GE port was 3.91 watts.

Frame Size	Total Watts	Watts per 40GE Port	Watts per 10GE Port
iMIX (64-9216)	2,066	15.65	3.91


Observations

The power in the lab at the time was measured at 202.2 Volts, somewhat lower than the 230 Volts that is typical for such environments. The reduced voltage will tend to cause the device's power supplies to operate at less than peak efficiency.

6 Competitive Power Consumption Analysis

Miercom confirmed in testing that the Dell Z9500 consumes 15.6 watts of power per 40GE port, while handling a full load of IMIX traffic (IP traffic of varied frame sizes).

As shown in the chart, the Z9500's power consumption per 40GE port positions it between two other switches previously tested in the Miercom 40GE Switch Industry Assessment.



Power consumption by the Dell switch per 40GE port was measured to be 15.6 watts, placing it on the heels of the Cisco Nexus 9508 switch for lowest power consumption. The Arista 7508E consumes 23.1 percent more power per 40GE port; the Cisco Nexus 9508 consumes 3.8 percent less. The Cisco Nexus results are based on data sheet specifications for the N9K-X9636PQ aggregation line card, and not the Application Centric Infrastructure (ACI-ready) leaf line card. All vendor line cards in this testing used Broadcom Trident II chipsets.

7 About the Miercom 40GE Switch Industry Study

This report was sponsored by Dell 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 completed to comment on the test 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.)

8 About Miercom

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.

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.

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 Dell 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.