

Dell EMC Networking S4100-ON Series Switches

Dell EMC S4148T-ON and S4148F-ON 100GbE Multirate Performance Evaluation

Executive Summary

Data center architects require high-capacity, high-performance switches at the core of the network to provide the best application performance for users. Dell EMC offers a complete set of data center purpose built Layer 2 and Layer 3 switches that not only provide multi-rate 10/25/40/50/100GbE interfaces but are also part of the open networking innovation supporting multiple Network Operating Systems (NOS) options.

Dell EMC commissioned Tolly to benchmark the overall performance, latency and power consumption of its S4100 Switch Product family, specifically the S4148T-ON (48x10GBASE-T) and the S4148F-ON (48xSFP+) switches. The Dell EMC Open Networking (ON) switches provide 880Gbps of backplane bandwidth and were tested with a full complement of 48 10GbE ports plus four 100GbE ports with over 800MPPS of packet throughput. The systems also support 25/40/50GbE ports.

The Dell EMC switches delivered 100% line-rate RFC2544 and 2889 Layer 2 and Layer 3 throughput running various port-pairing configurations of 10GbE and 100GbE at frame sizes of 128-bytes and greater. Tolly engineers also benchmarked 10GbE and 100GbE Layer 2/Layer 3 latency, power consumption, port buffering as well as route, ARP and MAC table capacities.

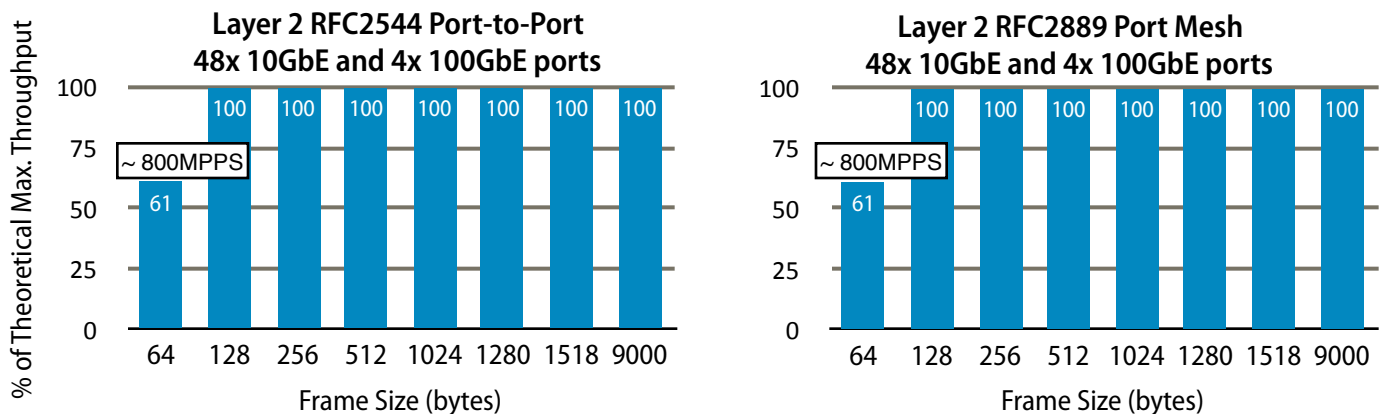
The Bottom Line

The Dell EMC Networking S4148T-ON and S4148F-ON Data Center Switches demonstrated:

- 1 100% line-rate Layer 2 and Layer 3 RFC2544 (port-to-port) throughput with 48x10GbE and 4x100GbE ports at 128-bytes and greater. Approx. 800MPPS processing rate
- 2 100% line-rate RFC2889 (full-mesh) Layer 2 and Layer 3 throughput with 48x 10GbE and 4x100GbE ports at 128-bytes and greater. Approx. 800MPPS processing rate
- 3 ARP table capacity - 135,000+; MAC table capacity - 256,000; Route Table IPv4 capacity - 16,000+ (to 128,000)

Dell EMC S4100 Series Layer 2 Switch Throughput: S4148T-ON & S4148F-ON

(as reported by Ixia IxNetwork v8.30)



Note: Dell EMC S4148T and S4148F reported identical throughput in test conditions and are presented as a single series. 10GbE and 100GbE ports tested concurrently. Layer 2 RFC2544 test run in "port-to-port" configuration. Layer 2 RFC2889 test run in "Dual Full-Mesh" configuration. 64-bytes used to determine max PPS rate.

Source: Tolly, August 2017

Figure 1



Test Results

Layer 2 Throughput: Port-to-Port & Port Mesh

For completeness, Tolly engineers benchmarked the throughput and latency using two similar and related Internet Engineering Task Force (IETF) RFC benchmarks: RFC2544 and RFC2889¹. The throughput tests used 48x10GbE and 4x100GbE simultaneously handling bi-directional traffic. All communication was between ports of the same topology: 10GbE to 10GbE and 100GbE to 100GbE. For RFC2544 tests, port-to-port pairs were used. For RFC2889 tests, traffic was transmitted in a mesh configuration where each port transmitted to and received from every other port of the same topology.

As the results for both switches were identical, they are represented with a single bar chart. At all frame sizes of 128-bytes and above, including "jumbo" 9000-byte frames, both switches demonstrated line-rate throughput. With 64-byte frames, both switches delivered 61% of line-rate

throughput or approximately 800MPPS. See Figure 1.

Layer 3 Throughput: Port-to-Port & Port Mesh

As the results for both switches were identical, they are represented with a single bar chart. At all frame sizes of 128-bytes and above, including "jumbo" 9000-byte frames, both switches demonstrated line-rate throughput. With 64-byte frames, both switches delivered 77% of line-rate throughput or approximately 800MPPS. (Note: Throughput is somewhat higher at L3 as the actual packet size is 84-bytes with IP headers.) See Figure 2.

Layer 2/3 Port-to-Port Latency: 10GbE & 100GbE

Last-in, first out (LIFO) latency was benchmarked in port-to-port configuration (RFC2544) and port-mesh configuration (RFC2889) separately at 10GbE and 100GbE. Results are discussed in separate sections.

For each of the switches, the Layer 2 and Layer 3 latency results were very similar. The

Dell EMC

Data Center Switches S4148T-ON and S4148F-ON



100GbE Multirate Performance

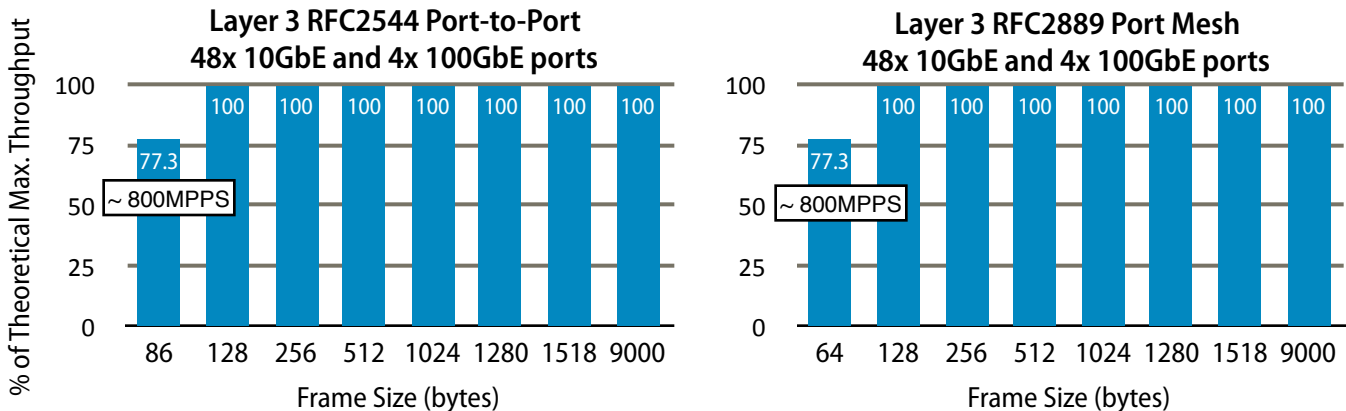
Tested August 2017

copper-based S4148T-ON latency ranged from 2,284 to 2,644ns across the range of frames tested at Layer 2 and from 2,319 to 2,649ns at Layer 3.

The fiber-based S4148-ON latency ranged from 779 to 1,141ns at Layer 2 and from 810 to 1,142ns at Layer 3. See Figure 3 for all port-to-port 10GbE results.

Across the 100GbE fiber ports, the results were, again, very similar at Layer 2 and Layer 3 and also from switch to switch.

Dell EMC S4100 Series Layer 3 Switch Throughput: S4148T-ON & S4148F-ON (as reported by Ixia IxNetwork v8.30)



Note: Dell EMC S4148T and S4148F reported identical throughput in test conditions and are presented as a single series. 10GbE and 100GbE ports tested concurrently. Layer 2 RFC2544 test run in "port-to-port" configuration. Layer 2 RFC2889 test run in "Dual Full-Mesh" configuration. 64-bytes used to determine max PPS rate.

Source: Tolly, August 2017

Figure 2

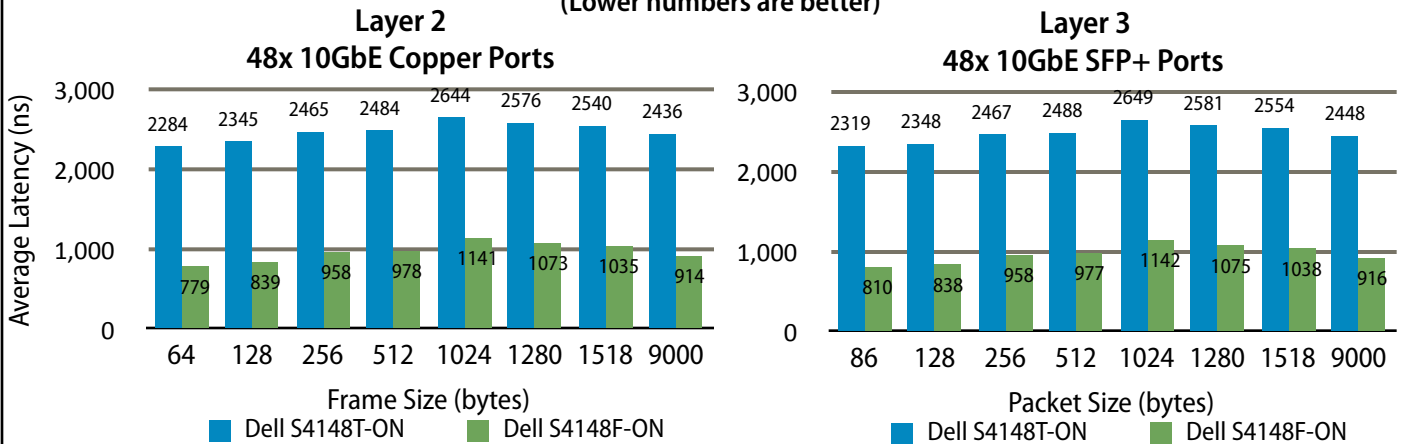
¹ IETF RFC2544 is the earlier benchmark. It was focused on routers and port-to-port traffic flows. RFC2889 was focused on LAN switches and mesh traffic flows. Those RFCs specify sizes to 1518. 9000 bytes is customarily tested as well.



Dell EMC S4100 Series Layer 2/3 Port-to-Port (RFC2544) 10GbE Latency

(as reported by Ixia IxNetwork v8.30)

(Lower numbers are better)



Note: Latency measured between all ports running bidirectional traffic at 100% line rate (port-to-port), including 4x 100GbE ports. LIFO latency measured, time to store data not included in measurement. Dell EMC S4148T was connected to Ixia via 2 meter Cat 5e cables, S4148F was outfitted with SFP+ optics and connected via 2 meter multi-mode fiber.
Source: Tolly, August 2017

Figure 3

The S4148T-ON latency ranged from 822 to 978ns across the range of frames tested at Layer 2 and from 822 to 975ns at Layer 3.

The S4148F-ON latency ranged from 812 to 968ns at Layer 2 and from 815 to 971ns at Layer 3.

Layer 3. See Figure 4 for all port-to-port 100GbE results.

Layer 2/3 Port-Mesh Latency: 10GbE & 100GbE

For each of the switches, the Layer 2 and Layer 3 latency results were very similar. The

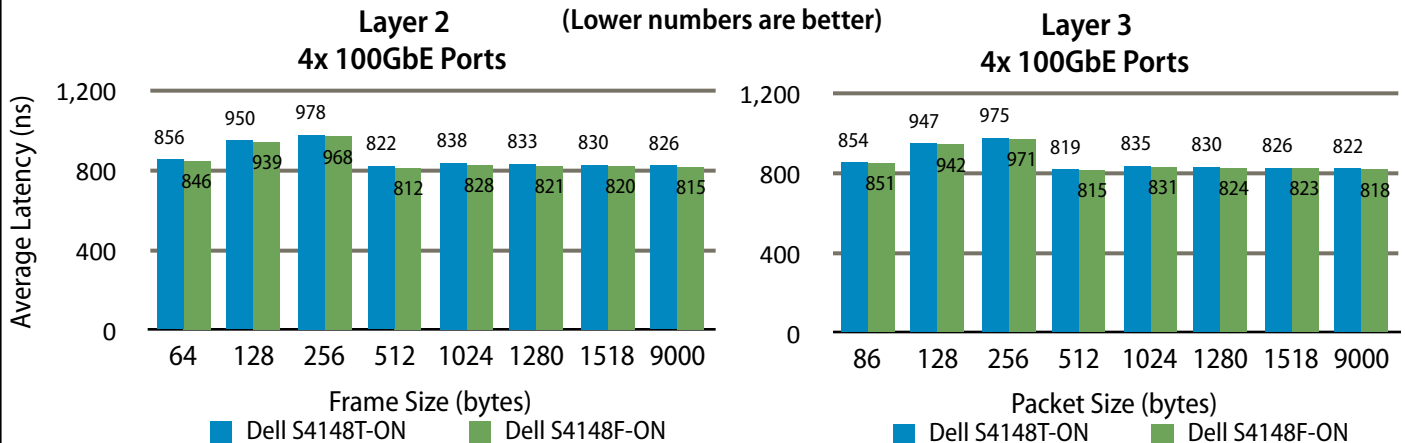
copper-based S4148T-ON latency ranged from 2,289 to 2,658ns across the range of frames tested at Layer 2 and from 2,318 to 2,674ns at Layer 3.

The fiber-based S4148F-ON latency ranged from 784 to 1,153ns at Layer 2 and from 813 to 971ns at Layer 3.

Dell EMC S4100 Series Layer 2/3 Port-to-Port (RFC2544) 100GbE Latency

(as reported by Ixia IxNetwork v8.30)

(Lower numbers are better)



Note: Latency measured between all ports running bidirectional traffic at 100% line rate (port-to-port), including 48x 10GbE ports in test. LIFO latency measured, time to store data not included in measurement.

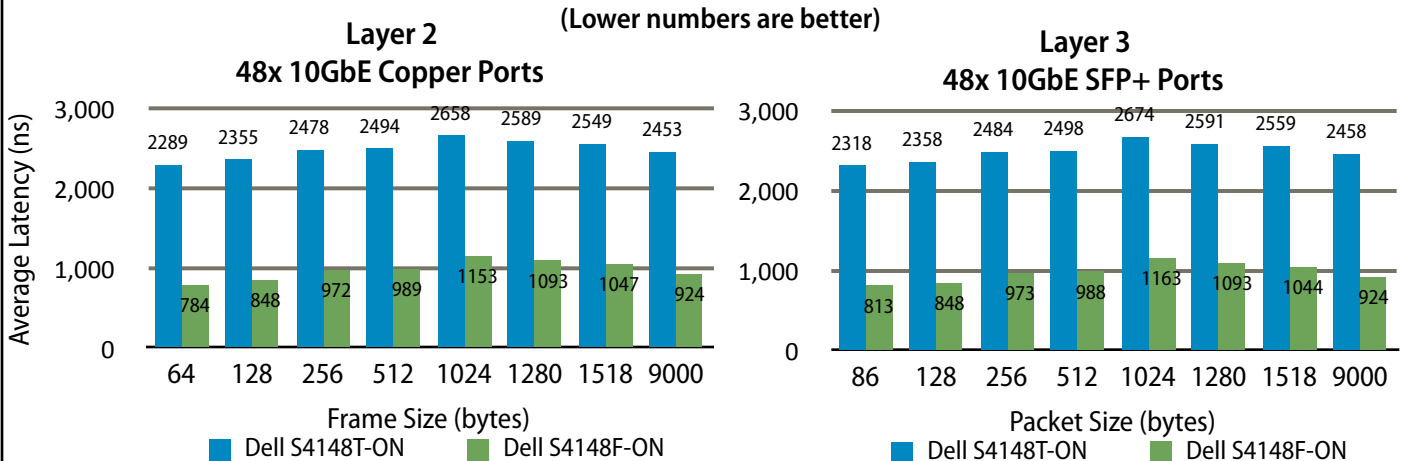
Source: Tolly, August 2017

Figure 4



Dell EMC S4100 Series Layer 2/3 Port Mesh (RFC2889) 10GbE Latency

(as reported by Ixia IxNetwork v8.30)



Note: Latency measured between all ports running bidirectional traffic at 100% line rate (Dual Full-mesh), including 4x 100GbE ports. LIFO latency measured, time to store data not included in measurement. Dell EMC S4148T was connected to Ixia via 2 meter Cat-5e cables, S4148F was outfitted with SFP+ optics and connected via 2 meter multi-mode fiber.

Source: Tolly, August 2017

Figure 5

to 970ns at Layer 3. See Figure 5 for all port-to-port 10GbE results.

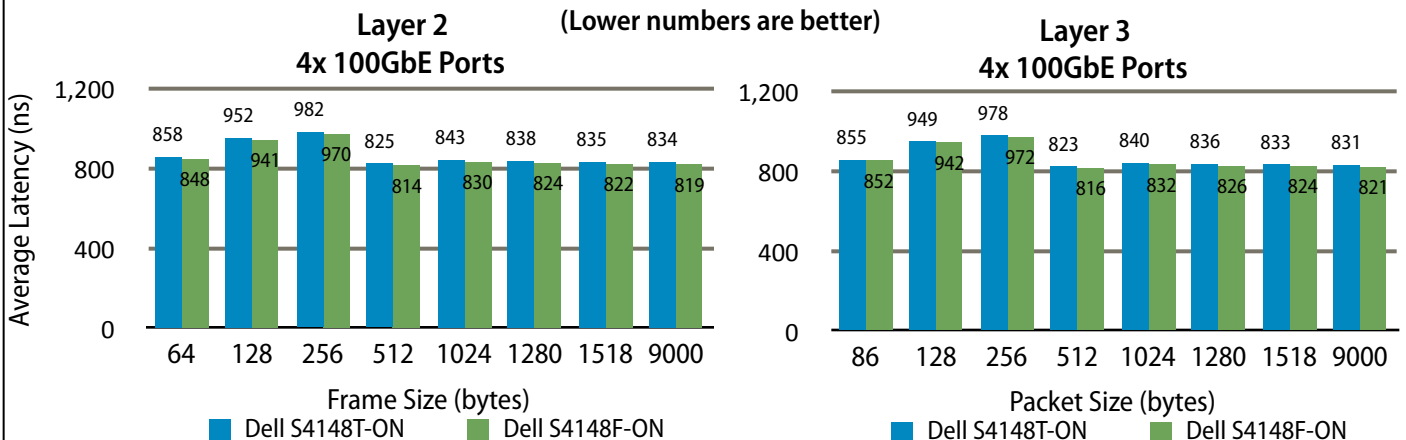
The S4148T-ON latency ranged from 834 to 982ns across the range of frames tested at Layer 2 and from 922 to 975ns at Layer 3.

Across the 100GbE fiber ports, the results were, again, very similar at Layer 2 and Layer 3 and also from switch to switch.

The S4148F-ON latency ranged from 834 to 970ns at Layer 2 and from 816 to 972ns at Layer 3. See Figure 6 for all port-to-port 100GbE results.

Dell EMC S4100 Series Layer 2/3 Port Mesh (RFC2889) 100GbE Latency

(as reported by Ixia IxNetwork v8.30)



Note: Latency measured between all ports running bidirectional traffic at 100% line rate (Dual Full-Mesh), including 48x 10GbE ports in test. LIFO latency measured, time to store data not included in measurement.

Source: Tolly, August 2017

Figure 6



Port Buffering

The Dell EMC switches use a common packet buffer that is shared across all ports of a given switch on a dynamic basis. These buffers can be used to prevent packet loss by buffering packets where an output port is momentarily oversubscribed. Tolly engineers determined that the port buffer capacity of each of the switches was eight MB.

Power Consumption

Tolly engineers benchmarked the power consumption and calculated the ATIS weighted power of each switch.

The S4148T has an ATIS weighted value of 217 and four Gbps/Watt. The S4148F-ON has an ATIS weighted value of 95 (lower numbers are better) and nine Gbps/Watt. See Table 1.

Table Capacity

Tolly engineers also benchmarked the maximum capacity of several key tables. The results were the same for both switches. The ARP table could handle over 135,000 in

hardware and, additionally, was verified to handle over 220,000 addresses when the table was expanded into software.

The MAC table was verified to hold over 275,000 addresses. See Table 1 for all table capacity results.

The IPv4 route table was verified to be able to handle 16,000 entries in its default mode. Dell EMC notes that the route table capacity is 128,000 in extended mode. (Not tested because of time constraints.)

Many-to-One (10x10GbE to 100GbE)

Finally, a basic test was run to verify that ten 10GbE links could have their traffic all directed to a single 100GbE output link. This test was successful.

Test Environment

Devices Under Test

See Table 2 for details about the models and version levels tested. Tests were run using ~2.5 meter lengths of multi-mode mode fiber cables and SFP+ SR between the test tool and the SUT. For the Dell EMC S4148T-

ON, 3 meter lengths of Category 5e cable were used to connect the DUT to the test tool.

For all 100GbE tests, 4x QSFP28 DAC cables were used.

Traffic Generation & Measurement

All test traffic was generated using an Ixia XGS12 Chassis. The chassis was equipped with one 8-Port Novus QSFP28 High-density 100/25GE Load Module, and 3x 16-port Novus 10GbE Load Modules. All benchmarks were run using Ixia IxNetwork V8.30.1077.21.

Test Methodology

Layer 2/Layer 3 Performance Test

All L2/L3 performance benchmarks were completed using RFC2544 and/or 2889 (as appropriate) as implemented by Ixia in its IxNetwork benchmarking suites for throughput and latency. A single "jumbo" frame size of 9000-bytes was tested that is beyond the RFC specification but commonly tested in any case.

Dell EMC Networking S4100 Series Power Consumption & Table Capacity

	Dell EMC S4148T-ON & S4148F-ON Power Consumption with 48x 10GbE and 4x 100GbE Ports				
	ATIS Weighted	No Traffic (all ports connected and up)	30% Traffic	100%Traffic	ATIS Gbps/Watt
Dell EMC S4148T-ON	217W	207W	216W	244W	4
Dell EMC S4148F-ON	95W	88W	95W	104W	9

	Dell EMC Networking S4140 Series Table Capacity		
	IPv4 Route Table (in Default Switch Mode)	ARP (in Default Switch Mode)	MAC (in Scaled L2 Switch Mode)
Maximum Table Entries	16,000*	135,168+	278,528

Note: Power measurements average traffic profile across all frame sizes as reported by BayTech MRP-21 Modular Rack Power. Hardware ARP table capacity of 135k, ARP table can be expanded into software. Tolly Engineers verified 220k+ addresses successfully. * Dell EMC notes the IPv4 routes table capacity is 128K in scaled/extended mode. Not tested due to time constraints.

Source: Tolly, August 2017

Table 1



Devices Under Test Details

Vendor	Product	Software Version	Port Configuration
Dell EMC	Dell S4148T-ON	Dell EMC Networking OS10 Enterprise Version: 10.3.1E(149)	48x 10GBASE-T Ports 2x QSFP+ Ports 4x QSFP28 Ports
Dell EMC	Dell S4148F-ON	Dell EMC Networking OS10 Enterprise Version: 10.3.1E(149)	48x 10GbE SFP+ Ports 2x QSFP+ Ports 4x QSFP28 Ports

Source: Tolly, August 2017

Table 2

Many-to-One

As part of the test cases, engineers benchmarked the switch capability to forward multiple 10GbE links to a single 100GbE link. Traffic profiles were configured to send traffic between 10x10GbE ports, and one 100GbE port.

Port Configuration

Both the Dell EMC Networking S4148T-ON and S4148F-ON are equipped with 54 physical ports, and has a combined 960Gbps of port-connectivity (See Table 2). However, the capacity of the switch fabric is 880Gbps. As a result, the switch is equipped with additional ports, and six port operating profiles to fit a variety of use cases. For the Tolly tests, a total of 52 ports were used.

This backplane is utilized fully by the configuration involving 48x10GbE ports and 4x100GbE (QSFP28) ports.

Layer 2

RFC2544 Port-to-Port Test

For the RFC2544 test, ports were configured as 'pairs' with like PHYs. Two groups, consisting of 48x10GbE ports, and 4x100GbE ports, were independently placed in pairs.

Layer 2, Bidirectional traffic was configured to pass on each of the designated port groups. The binary search feature of IxNetwork was utilized to determine the zero-loss throughput for each of the frame sizes: 64, 128, 256, 512, 1024, 1280, 1518, and 9000 bytes.

Latency and throughput were recorded during the 1 minute test iterations, were run three times at each frame size to ensure accuracy.

RFC2889 Port-Mesh Test

Maintaining the same physical connections as the previous tests, engineers provisioned the switch into two different full-mesh

groups, consisting of the 48x10GbE ports, and the 4x100GbE ports.


Engineers configured traffic from each port to every other port in each group. The RFC2889 tests were run for the same duration and packet sizes as the RFC2544 tests.

Latency Measurements

The store-and-forward option was used in each RFC test creation to measure the latency as last in first out (LIFO) latency. Results as reported by Ixia.

Test Equipment Summary

The Tolly Group gratefully acknowledges the providers of test equipment/software used in this project.

Vendor	Product	Web
Ixia	Ixia XGS12 Chassis, 1x NOVUS™ QSFP28 High-Density 100/25GE Load Module, 3x Novus 3-Speed 10G/1G/100M IxNetwork V8.00.1027.17EA	 www.ixiacom.com

Layer 3

RFC2544 and RFC2889

The same physical port connections were used for the Layer 3 tests. An IP address was assigned for each Switch port from a different /24 subnet. An interface was created on each Ixia port, and a valid IP was assigned for each port.

Port mappings for the RFC2544 and RFC2889 tests were identical to the Layer 2 tests. Engineers configured Ixia to generate valid IP traffic between port pairs, and results were recorded in a similar fashion. Due to the IP header overhead, the minimum packet size for Layer 3 was increased to 86 bytes, all other packet sizes remained consistent with Layer 2 tests.

Port Buffering

To saturate the switch port buffers, engineers configured Ixia to provide one line-rate traffic stream destined for a single port. Engineers created a separate traffic burst destined to this fully-loaded port, and verified that after traffic was stopped, that all packets were received.

Power Consumption

Power consumption data was recorded by a BayTech MRP-21 Modular Rack Power Supply. The circuit for each switch was independent from any other operating equipment. The same test setup was used as for the RFC2544 and RFC2889 performance tests. When all 48x10GbE 4x100GbE ports were connected and up but without traffic, the power consumption was recorded.

Then Tolly measured the power consumption with 30% line-rate traffic and 100% line-rate traffic. Data was recorded after 1 minute of steady state operation for each the 'loaded' test cases. The ATIS weighted power consumption was calculated as $ATIS(w) = (Idle) * 0.1 + (30\%) *$

Dell EMC Networking S4148T-ON



Source: Dell EMC

$0.8 + (100\%) * 0.1$ (Watts) as a data center switch.

Table Capacity

FIBv4, ARP table, and MAC table capacity were evaluated individually in different switch modes.

In the FIBv4 test, engineers sent in 16,000 OSPF routes from the Ixia test tool to the Dell EMC switch and verified that all routes worked properly to forward traffic. According to the specifications, the number of routes is 128,000. Due to time constraints this test was not performed".

In the ARP table test, engineers sent in 220,000 ARP replies from Ixia to the Dell EMC switch and verified that all ARP replies were learnt by the DUT.

In the MAC table test, engineers created 132,000 hosts on one Ixia port with different MAC addresses and sent traffic to the Dell

Then engineers sent traffic from the other Ixia port to the 132,000 hosts through the Dell EMC switch. All traffic was forwarded by the MAC table without broadcasting.

The S4100 Series Switch has an option for a 'Scaled Layer 2 Mode', which increases the MAC table to 278K addresses. Engineers applied this settings, and re-ran the test with an increased number of addresses. The switch was again able to learn all MAC addresses.

Dell EMC Networking S4148F-ON



Source: Dell EMC

EMC switch to let it learn the MAC addresses.



About Tolly...


The Tolly Group companies have been delivering world-class IT services for over 25 years. Tolly is a leading global provider of third-party validation services for vendors of IT products, components and services.

You can reach the company via E-mail at sales@tolly.com, or via telephone at +1 561.391.5610.

Visit Tolly on the Internet at: <http://www.tolly.com>

About Dell EMC Networking

Visit www.dell.com/networking for more information on for more information on the datacenter switches and additional future-ready networking solutions from Dell EMC Networking.



Terms of Usage

This document is provided, free-of-charge, to help you understand whether a given product, technology or service merits additional investigation for your particular needs. Any decision to purchase a product must be based on your own assessment of suitability based on your needs. The document should never be used as a substitute for advice from a qualified IT or business professional. This evaluation was focused on illustrating specific features and/or performance of the product(s) and was conducted under controlled, laboratory conditions. Certain tests may have been tailored to reflect performance under ideal conditions; performance may vary under real-world conditions. Users should run tests based on their own real-world scenarios to validate performance for their own networks.

Reasonable efforts were made to ensure the accuracy of the data contained herein but errors and/or oversights can occur. The test/audit documented herein may also rely on various test tools the accuracy of which is beyond our control. Furthermore, the document relies on certain representations by the sponsor that are beyond our control to verify. Among these is that the software/hardware tested is production or production track and is, or will be, available in equivalent or better form to commercial customers. Accordingly, this document is provided "as is", and Tolly Enterprises, LLC (Tolly) 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 herein. By reviewing this document, you agree that your use of any information contained herein is at your own risk, and you accept all risks and responsibility for losses, damages, costs and other consequences resulting directly or indirectly from any information or material available on it. Tolly is not responsible for, and you agree to hold Tolly and its related affiliates harmless from any loss, harm, injury or damage resulting from or arising out of your use of or reliance on any of the information provided herein.

Tolly makes no claim as to whether any product or company described herein is suitable for investment. You should obtain your own independent professional advice, whether legal, accounting or otherwise, before proceeding with any investment or project related to any information, products or companies described herein. When foreign translations exist, the English document is considered authoritative. To assure accuracy, only use documents downloaded directly from Tolly.com.

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