TEST REPORT

Tolly.

#221111

February 2021

Commissioned by

Dell EMC

Dell EMC Networking Data Center Switches

Dell EMC Z9332F-ON & Z9432F-ON 400GbE 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 400GbE 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 two of its data center switches: the Z9332F-ON (32x 400GbE + 2x 10GbE) and the Z9432F-ON (32x 400GbE + 2x 10GbE).

The Dell EMC switches delivered 100% line-rate RFC2544 Layer 2 and Layer 3 throughput with 32x 400GbE ports in full mesh topology at frame sizes of 512-bytes and greater. Tolly engineers also benchmarked Layer 2/Layer 3 latency, power consumption as well as MAC/FIB/ARP/ACL/IGMP/PIM table capacities.

THE BOTTOM LINE

The Dell EMC Networking data center switches demonstrated:

1. Z9332F-ON: 100% line-rate L2/L3 throughput at 512-byte and higher frame sizes with 32x 400GbE ports
2. Z9332F-ON: FIBv4 and IPv4 routing table capacity of 300,000 BGP routes
3. Z9432F-ON: 100% line-rate L2/L3 throughput at 512-byte and higher frame sizes with 32x 400GbE ports
4. Z9432F-ON: MAC table capacity of 155,650 MAC addresses; ARP table capacity of 160,000 ARP entries

Dell EMC 400GbE Data Center Switches

Z9332F-ON

Z9432F-ON

Source: Dell EMC

Figure 1

Test Results

Layer 2/3 RFC2544 Throughput

Tolly engineers benchmarked the throughput and latency using IETF RFC2544. The throughput tests used all 32x 400GbE ports on the switch simultaneously handling bi-directional traffic in a full-mesh topology. The maximum throughput with zero frame loss is reported for each test.

Z9332F-ON

As the results for both Layer 2 and Layer 3 tests were identical, they are represented with a single bar chart. At all frame sizes of 512-byte and above, including "jumbo" frames, the Dell EMC Z9332F-ON switch demonstrated line-rate 12.8Tbps throughput. With 64-byte frames, the switch delivered 5.3Tbps which is 41.41% of line-rate throughput. With 128-byte frames, the switch delivered 9.35Tbps which is 73.05% of line-rate throughput. With 256-byte frames, the switch delivered 11.68Tbps which is 91.21% of line-rate throughput. See Figure 2.

Dell EMC notes that the switch achieves 100% line-rate with a frame size 295 bytes. However, this was not tested due to time constraints.

Z9432F-ON

As the results for both Layer 2 and Layer 3 tests were identical, they are represented with a single bar chart. At all frame sizes of 512-byte and above, including "jumbo" frames, the Dell EMC Z9432F-ON switch demonstrated line-rate 12.8Tbps throughput. With 64-byte frames, the switch delivered 3.575Tbps which is 27.93% of line-rate throughput. With 128-byte frames, the switch delivered 6.35Tbps which is 49.61% of line-rate throughput. With 256-byte frames the switch delivered 11.82Tbps which is 92.38% of line-rate throughput. See Figure 3.

Dell EMC notes that the switch achieves 100% line-rate with a frame size 286 bytes. However, this was not tested due to time constraints.

Dell EMC

Data Center Switches Z9332F-ON and Z9432F-ON

400GbE Performance

Tested January 2021

Z9332-ON: Layer 2 & Layer 3 RFC 2544 Switch Throughput with 32x 400GbE Ports in Full-mesh Topology (as reported by Spirent TestCenter)

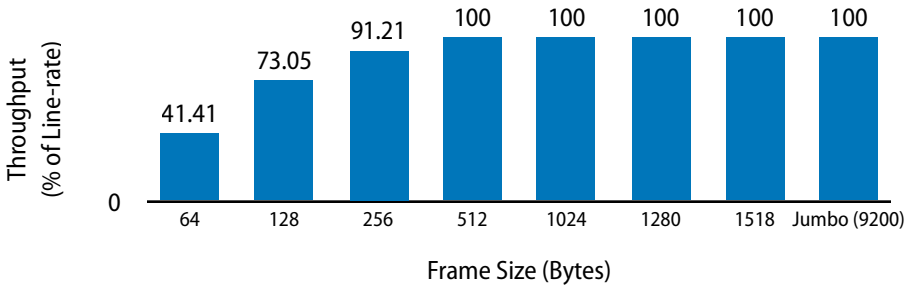

Note: Layer 2 and Layer 3 results were the same. There was no frame loss in the test with the reported throughput.

Source: Tolly, January 2021 Figure 2

Z9432-ON: Layer 2 & Layer 3 RFC 2544 Switch Throughput with 32x 400GbE Ports in Full-mesh Topology (as reported by Spirent TestCenter)

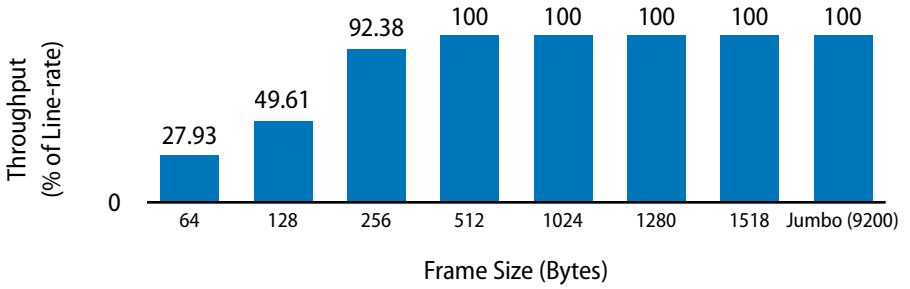

Note: Layer 2 and Layer 3 results were the same. There was no frame loss in the test with the reported throughput.

Source: Tolly, January 2021 Figure 3

Layer 2/3 RFC2544 Latency

Last-in, first out (LIFO) latency was benchmarked in a full-mesh topology with all 32x 400GbE ports on the switch. Z9332F-ON tests used the throughput results in Figure 2 as the traffic rate for each frame size. Z9432F-ON tests used the throughput results in Figure 3 as the traffic rate for each frame size.

Z9332F-ON

Tests measured latency on the 400GbE ports of the Z9332F-ON switch in store-and-forward mode. The results at Layer 2 and Layer 3 were almost identical. Latency averaging from 0.85 microseconds (µs) to 0.95µs was reported for different frame sizes. See the top chart in Figure 4.

Z9432F-ON

Tests measured latency on the 400GbE ports of the Z9432F-ON switch in store-and-forward mode. The results at Layer 2 and Layer 3 were identical. Latency averaging from 1.08 microseconds (µs) to 1.13µs was reported for different frame sizes. See the bottom chart in Figure 4.

Dell EMC Data Center Switches: Layer 2 & Layer 3 RFC 2544 Switch Latency
with 32x 400GbE Ports in Full-mesh Topology
(as reported by Spirent TestCenter)

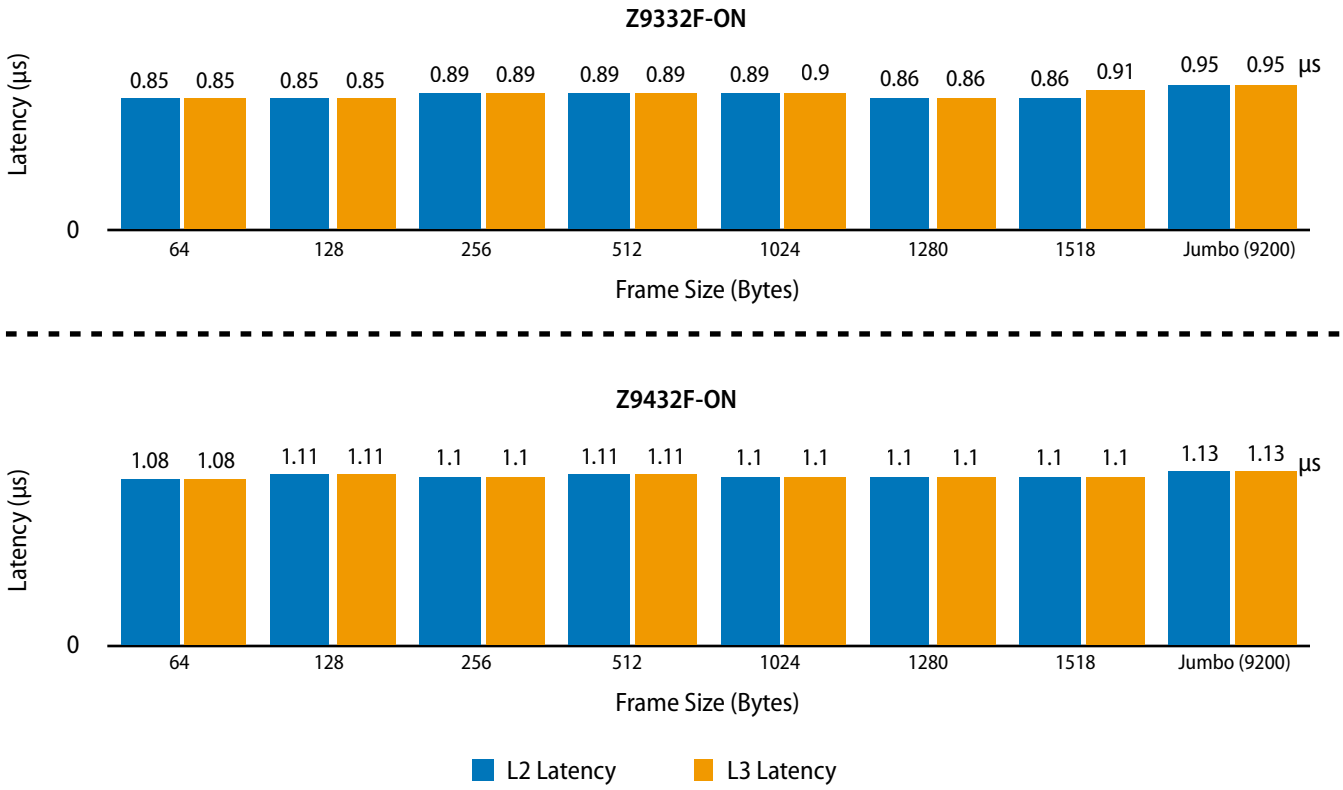

Note: Each switch has 32x 400GbE ports + 2x 10GbE ports. In each test, all 32x 400GbE ports on the switch were used in full-mesh topology. Traffic rate used the throughput results reported in Figure 2 and Figure 3. Last-In-First-Out (LIFO) latency type is reported.

Source: Tolly, January 2021

Figure 4

Power Consumption

Tolly engineers benchmarked the power consumption and calculated the ATIS weighted power of each switch according to the ATIS-0600015.03.2016 "Energy Efficiency for Telecommunication Equipment: Methodology for Measurement and Reporting for Router and Ethernet Switch Products" document. See Table 1.

The Z9332F-ON has ATIS weighted power of 329 Watts and ATIS TEER (Telecommunication Energy Efficiency Ratio) of 38.9Gbps/Watt.

The Z9432F-ON has ATIS weighted power of 332 Watts and ATIS TEER of 38.6Gbps/Watt.

Table Capacity

Tolly engineers also benchmarked the maximum capacity of several key tables on the Z9332F-ON and the Z9432F-ON switches. Please see Table 2 for verified capacity of each table.

Dell EMC Networking Data Center Switches Power Consumption

	ATIS Weighted Power (lower is better)	Idle Power (all ports connected and up)	30% Traffic	100% Traffic	ATIS TEER (Gbps/Watt)
Z9332F-ON	329W	299W	326W	382W	38.9
Z9432F-ON	332W	301W	329W	383W	38.6

Notes: 1. Z9332F-ON and Z9432F-ON each provides 32x 400GbE ports + 2x 10GbE ports. All 32x 400GbE ports were used in the test with 12,800Gbps maximum demonstrated throughput.

2. ATIS weighted power = 0.1 * (idle power) + 0.8 * (power with 30% traffic) + 0.1 * (power with 100% traffic). The default iMIX profile in Spirent TestCenter was used in the test with all 32 * 400GbE ports in full-mesh topology.

Source: Tolly, January 2021 Table 1

Dell EMC Networking Data Center Switches Table Capacity

	Z9332F-ON	Z9432F-ON
MAC Table	7,500	155,650
ARP Table	7,500	160,000
FIBv4 / IPv4 Routing Table - BGP	300,000	64,000
L2 MAC Ingress ACL Rules	2,302	4,094
L2 MAC Egress ACL Rules	510	2,046
IPv4 Ingress ACL Rules	2,300	4,092
IPv4 Egress ACL Rules	508	508
IPv6 Ingress ACL Rules	764	2,044
IPv6 Egress ACL Rules	252	508
IGMP Snooping Groups	8,000	8,000
IGMP Groups	8,000	8,000
PIM	8,000 (S,G), 8,000 (*,G)	8,000 (S,G), 8,000 (*,G)

Source: Tolly, January 2021 Table 2

Test Methodology

Test Bed

The test bed includes one Dell EMC PowerSwitch Z9332F-ON data center switch, one Dell EMC PowerSwitch Z9432-ON data center switch, and four Spirent pX3 400G appliances. The four Spirent appliances provide 32x 400GbE ports in total which were all directly connected to the Dell EMC switch under test.

RFC2544 Throughput and Latency

All L2/L3 performance benchmarks were completed using the RFC2544 wizard in Spirent TestCenter for throughput and latency. A "jumbo" frame size of 9200-byte was tested that is beyond the RFC specification but commonly tested in any case. Each frame size was tested for three iterations with one minute per iteration. In all test runs, all three iterations had the same result.

Full-mesh traffic topology was used in the test with traffic from each port to all other ports evenly using unicast traffic. L2/L3 learning was performed before each test run. Maximum throughput with zero frame loss was reported for each frame size.

In L2 tests, the switch ports worked in L2 (switch port) mode. In L3 tests, the switch ports worked in L3 (no switch port) mode.

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 performance tests with iMIX packets. When all 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 weighted power (Watts) = 0.1 * (idle power) + 0.8 * (power with 30% traffic) + 0.1 * (power with 100% traffic) as a data center switch.

Dell EMC 400GbE Data Center Switches Test Bed

1x Dell EMC PowerSwitch Z9332F-ON data center switch,
1x Dell EMC PowerSwitch Z9432F-ON data center switch,
4x Spirent pX3 400G appliances (qx3-qsfp-dd-8)

For each switch, 32x 400GbE ports are used in full-mesh traffic topology

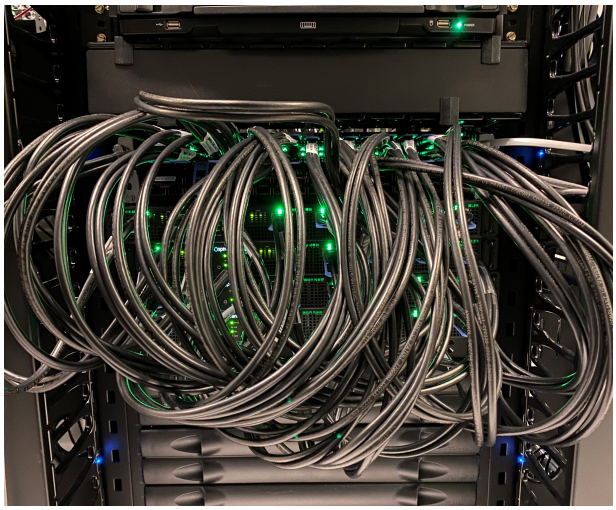

Source: Tolly, January 2021

Figure 5

Table Capacity

Capacity of each table was evaluated individually in the default switch mode.

In the MAC table test, engineers created devices on one Spirent port with different MAC addresses and sent traffic to the Dell EMC switch to let it learn the MAC addresses. Then engineers sent traffic from the other Spirent port to the devices through the Dell EMC switch. All traffic was forwarded by the MAC table without broadcasting. A third Spirent port was connected to the Dell EMC switch to verify that there was no broadcast test frames.

In the ARP test, engineers created devices on one Spirent port and let the devices resolve gateway's ARP. The switch which was the gateway for the Spirent emulated devices then learned ARP records of all emulated devices. Test streams were then sent in from another Spirent port to verify all ARP records on the switch worked properly.

In the FIBv4 test, engineers sent in BGP routes from one Spirent port to the Dell EMC switch and verified that all routes worked properly on the switch to forward traffic with traffic from the other Spirent port.

In the ACL tests, ACL rules were enabled on the ingress or egress direction of a port. Half of the rules were "deny" rules and the other half were "allow" rules. Then test traffic matching all the rules were sent in. Engineers verified that traffic matching the "deny" rules were all dropped by the switch while traffic matching the "allow" rules were all forwarded by the switch without loss.

In the multicast tests, one Spirent port sent in IGMPv2 packets to join the multicast groups. Another Spirent port sent in L2 or L3 multicast traffic matching all multicast groups. Tolly engineers verified that the switch only forwarded the multicast traffic to the Spirent port sending in IGMP packets to subscribe.

About Tolly

The Tolly Group companies have been delivering world-class ICT services for more than 30 years. Tolly is a leading global provider of third-party validation services for vendors of ICT products, components and services.

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

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

Test Equipment Summary

Vendor	Product	Web
Spirent	TestCenter	https://www.spirent.com

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.

ef-u-yx-wt-2021-02-15-VerD2