TEST REPORT

Tolly.

#220114

April 2020

Commissioned by Broadcom, Inc.

Emulex® LPe35002 Host Bus Adapter

32G Fibre Channel HBA Performance vs. Marvell QLogic QLE2772

EXECUTIVE SUMMARY

Innovations in server and storage technologies such as all-flash arrays, NVMe, and NVMe over Fabrics (NVMe-oF) are delivering record speeds and lower latency to improve application performance. As a result, enterprise datacenters require high performance Gen 7 32G Fibre Channel (32GFC) throughput for server connectivity that also provides extreme reliability and scalability under load. Additionally Gen 7 Fibre Channel maintains backward compatibility with existing Fibre Channel infrastructure to ensure streamlined performance upgrades.

Broadcom commissioned Tolly to benchmark the performance of its Emulex Gen 7 LPe35002 32GFC host bus adapter (HBA) against the performance of Marvell's QLogic QLE2772 32GFC HBA (branded as "enhanced 32GFC"), running database applications and synthetic benchmarks. Both products are second generation, dual-port, 32GFC HBAs. The series of enterprise datacenter-oriented tests performed included synthetic IO, latency, bandwidth, and database OLTP tests.

The Emulex LPe35002 HBA demonstrated superior performance across all test scenarios including up to 5x higher IOPS, as much as 1/2 the latency, and approximately 20% more transactions per minute (TPM) for Microsoft SQL Server and Oracle Database OLTP workloads.

THE BOTTOM LINE

The Emulex LPe 35002 FC HBA delivered:

1. IOPS performance 4.9x that of QLogic for single-port and 3.2x QLogic for dual-port
2. Latency is half that of QLogic
3. SQL Server database TPM throughput 19% higher than QLogic
4. Oracle 19c database TPM throughput 22% higher than QLogic

32G Fibre Channel HBA Maximum IOPS
(as reported by Viavi Medusa)

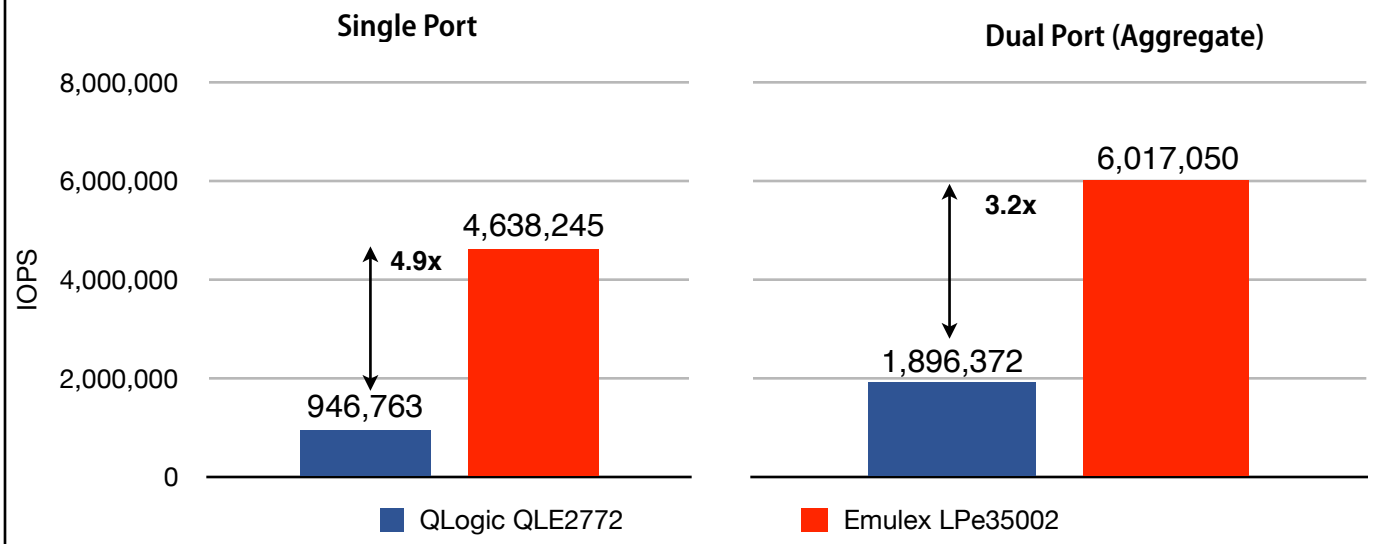

Note: Connected via Brocade switch. Payload size of 512-bytes. Publicly available software/firmware. Default settings. Average of three 30-second runs.

Source: Tolly, April 2020

Figure 1

Overview

The test scenarios are designed, in part, to prove the benefits of two aspects of the Emulex architecture: 1) Dynamic multi-core, and 2) Fast path.

Dynamic multi-core enables the Emulex LPe35002 to apply ASIC resources to ports as needed. Since many dual-port HBAs are used in active-standby configuration, this allows all ASIC resources to be allocated to the single, active port while, at the same time, being able to reallocate those resources to the second port should it become the active port.

Fast Path processes IO requests in hardware in most instances with firmware only handling the exceptions. The result is significantly reduced latency.

Test Results

Three sets of test scenarios were run: IOPS, latency and database transactions. Details are found in the Test Setup & Methodology section of this report.

Maximum IOPS

This test benchmarked the maximum IOPS performance of the HBA using a single port and then using both ports simultaneously.

In the single port test, Emulex demonstrated the performance benefits of dynamic ASIC use. The Emulex HBA delivered 4.9x the throughput of QLogic with an average of 4,638,245 IOPS for Emulex compared with 946,763 IOPS for QLogic.

In the dual port test, the Emulex HBA delivered 3.2x the throughput of QLogic with an average across both ports of 6,017,050 IOPS for Emulex compared with 1,896,372 IOPS for QLogic. See Figure 1.

Latency

This test benchmarked the round-trip latency of the HBA using a single port between two systems that were directly connected.

In this test, Emulex demonstrated the benefits of its hardware acceleration. The Emulex latency was 50% that of QLogic. The QLogic latency was 21.6 µs compared with only 10.6 µs for Emulex. (With latency, lower is better.) See Figure 2.

Database Transactions

While low-level component benchmarks are instructive, ultimately system architects are rightly most interested in how low-level improvements can translate into application performance.

To provide an application context, relational database transaction performance tests were run using both Microsoft SQL Server 2019 and Oracle Database 19c database environments.

Emulex

LPe35002 HBA

32G Fibre Channel Host Bus Adapter Performance

Tolly. Certified

Tested April 2020

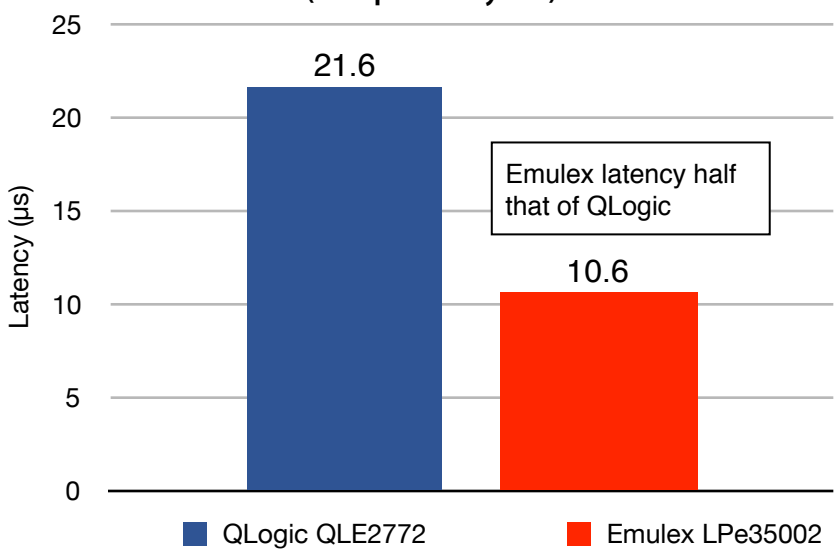

Note: Direct connection. Average of 3 runs. Emulex used as target for both tests. QLogic results calculated by factoring out Emulex latency.

Source: Tolly, April 2020

Figure 2

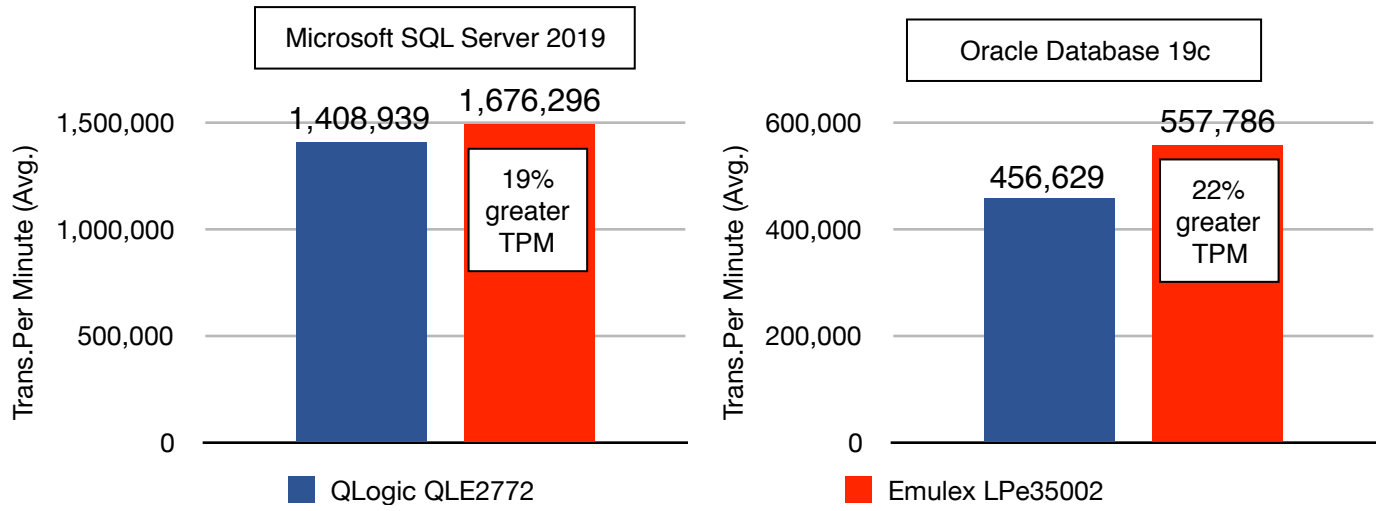

Note: HammerDB TPC-C load. Better of two runs reported. All flash NVMe array.
Source: Tolly, April 2020

Figure 3

This test benchmarked the maximum transactions per minute (TPM) throughput of each HBA using a single port and tested separately with the Microsoft and Oracle database solutions.

In the test of Microsoft SQL Server 2019, the Emulex HBA system delivered 19% greater average transaction throughput than QLogic with an average of 1,676,296 TPM for Emulex compared with an average of 1,408,939 TPM for QLogic.

Similarly, in the test of Oracle Database 19c, the Emulex HBA system delivered 22% greater average transaction throughput than QLogic with an average of 557,786 TPM for Emulex compared with an average of 456,629 TPM for QLogic. See Figure 3.

Test Setup & Methodology

The HBAs under test used current production drivers that were publicly available. Dell-branded and HPE-branded versions of the Marvell QLogic HBA were used along with drivers provided by each of the OEM vendors respectively. Default settings were used.

Maximum IOPS Test

The goal of this test was benchmark the maximum input/output performance of each HBA. An initiator server was configured with the HBA under test. This HBA drove 33 target HBAs in five servers. While each server had four, dual-port LPe35002 HBAs not all ports were required to achieve maximum IOPS.

As each HBA under test was a dual-port adapter, the test was run in both a single-port and a dual-port configuration. The VIAVI Solutions Medusa 7.2 test tool was used for benchmarking with a 512-byte block size. Test runs were 30-seconds each. Tests were run three times and the results were averaged. Additional details of the test environment are found in Tables 1-4.

Latency Test

The goal of this test was benchmark the latency performance of each HBA. With latency tests, lower latency is better. An initiator server was configured with the HBA under test. This server was directly connected (i.e. no FC switch) to a target server. Because no appropriate target server drivers were available for the QLogic HBA, the Emulex HBA was used as the target for both tests. The test was run with a single port.

The QLogic results were calculated based on the measured latency of the Emulex HBA. The following formula was used to calculate the QLogic round trip (RT) latency: ((QLE2772 Raw RT latency) – (LPe35002 / 2)) * 2 = QLE2772 calculated RT latency.

The open source FIO 3.7 test tool was used for benchmarking with a 512-byte block size. Test runs were 30-seconds each. Tests were run three times. The average latency was recorded and these results were averaged.across the test runs. Additional

Test Bed Topology: IOPS Performance

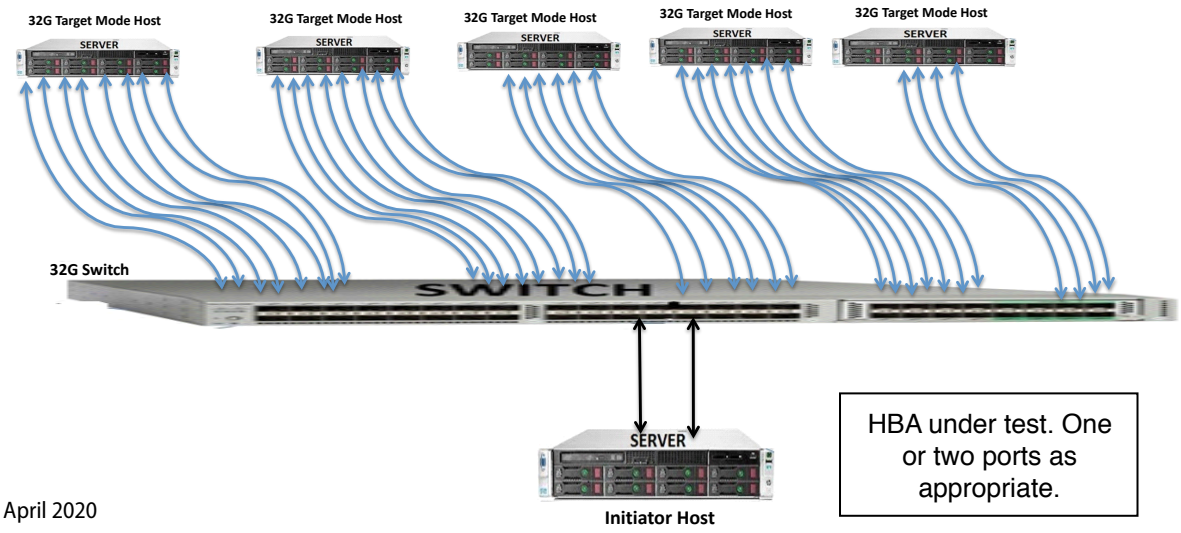

Source: Tolly, April 2020

Figure 4

Maximum IOPS Test Configuration Summary

32G HBAs Under Test

Vendor	Product Name	Driver Version	Firmware Version	Notes
Broadcom	Emulex LPe35002 (32G)	12.6.182.4	12.6.182.8	
Marvell	QLogic QLE2772 (32G)	10.01.00.54.08.0-k	9.02.05 (d0d5)	Dell branded

Table 1

Server Configuration

Number of Nodes	1 Initiator, 5 Target
CPU	2 socket - Intel(R) Xeon(R) CPU E5-2699 v3 @ 2.30GHz (72 CPUs)
Hyperthreading	Enabled
Memory (RAM)	128 GB
BIOS Power Mode	Performance
OS	Red Hat Ent. Linux 8 (RHEL8)

Table 2

Test Tool

Vendor	VIAVI Solutions
System	Medusa Labs Test Tools Suite v.7.2
Configuration	Threads: 1, Queue Depth: 30, Total Namespaces: 33
Test Parameters	maim 32768b -b512B -t1 -Q30 -n -u -r -o -%f100 -d30 -Y1 -M0 --perf-mode -f/root/targets.dat

Table 3

Target Configuration - NVMe

RAMD Driver	LPe35002 FW: 12.6.182.8 SPDK_fc_nvme_pkg_12.6.240.6 OCS_SDK_PKG_12.6.240
Target Ports (32G)	33
Namespace	1 namespace per target port
Network Switch	Brocade G620 Gen 6 32GFC switch

Source: Tolly, April 2020

Table 4

Latency Test Configuration Summary

32G HBAs Under Test

Vendor	Product Name	Driver Version	Firmware Version	Notes
Broadcom	Emulex LPe35002 (32G)	12.6.182.8	12.6.182.1 (OOB)	Both HBAs Installed in initiator server
Marvell	QLogic QLE2772 (32G)	10.01.00.55.08.0-k	9.02.05 (d0d5)	

Table 5

Server Common Configuration

Number of Nodes	1 Initiator, 1 Target
CPU	2 socket - Intel(R) Xeon(R) Platinum 8176 CPU @ 2.10GHz (56 CPUs)
Hyperthreading	Disabled
Memory (RAM)	64 GB
BIOS Power Mode	Performance

Table 6

Initiator Software

OS	RHEL 8.0
Kernel	4.18.0-80.el8_0.x86_64
Test Tool	Fio 3.7
Test Parameters	fio --name=test --ioengine=libaio --iodepth=1 --numjobs=1 --bs=512b --rw=read --size=64M --direct=1 --runtime=30 --group_reporting --time_based --filename=/dev/sdd --cpus_allowed=6

Table 7

Target Software & Hardware

OS	CentOS 7
Kernel	3.10.0-1062.el7.x86_64
Data Plane Dev. Kit (DPDK)	16.11.3
HBA	LPe35002 FW: 12.6.182.8 SPDK_fc_nvme_pkg_12.6.240.6 OCS_SDK_PKG_12.6.240

Note: Emulex HBA used for Target for tests of both vendors. Initiator and target were direct connected. No FC switch was used.

Source: Tolly, April 2020

Table 8

details of the test environment are found in Tables 5-8.

OLTP Database Test

The goal of this test was benchmark the database transaction performance of each HBA. An initiator server was configured with the HBA under test. This HBA connected to a NetApp storage target system via a Brocade 32GB FC switch. The test utilized a single port.

This test was run using two different database solutions: 1) Microsoft SQL Server 2019, and 2) Oracle Database 19c.

The open source HammerDB 3.2 test tool was used to populate the database schema and run the workload using its TPC-C load test option. Test runs were 10 minutes each, including a 5 minute ramp up time. Tests were run two times and the better result was used.

Additional details of the test environment are found in Tables 9-14. There were slight differences in the database and test configuration between the Microsoft and Oracle database tests and those details are noted in the aforementioned tables.

OLTP Database Test Configuration Summary

32G HBAs Under Test

Vendor	Product Name	Driver Version	Firmware Version	Notes
Broadcom	Emulex LPe35002 (32G)	12.4.243.20	12.4.282.0	
Marvell	QLogic QLE2772 (32G)	10.01.00.52.08.0-k	9.02.02 (d0d5)	HPE branded SN 1610Q

Table 9

Server Configuration

Number of Nodes	1 Initiator
CPU	2 socket - Intel(R) Xeon(R) Gold 6148 CPU @ 2.40GHz (80 CPUs)
Hyperthreading	Enabled
Memory (RAM)	128 GB
Power Mode	Performance
OS	Red Hat Ent. Linux 8 (RHEL8)
Kernel	4.18.0-80.11.2.el8_0.x86_64

Table 10

Microsoft Database Configuration

Database	SQL Server 2019
Storage	XFS file system
Dataset Size	100 GB
DB Memory Allocation	1:10 relationship is maintained between dataset and the amount of memory allocated for SQL DB

Table 11

Test Tool

Vendor	Open Source
System	HammerDB 3.2
TPC-C Load Settings	Warehouses: 1000 (Microsoft), 5000 (Oracle) Virtual Users: 80 Ramp up time: 5 min. Run time: 5 min.

Table 12

Storage Configuration

Vendor/Device	NetApp AFF A800 2 Nodes/Clusters
Release	9.7
Namespace/Lun (Microsoft SQL)	500 GB mapped through one 32G target port
Namespace/Lun (Oracle 19c)	1.5 TB mapped through one 32G target port
Network Switch	Brocade G620 Gen 6 32GFC switch

Table 13

Oracle Database Configuration

Database	Oracle Database 19c
Storage	XFS file system
Dataset Size	500 GB
DB Memory Allocation	1:10 relationship is maintained between dataset and the amount of memory allocated for Oracle DB

Table 14

Source: Tolly, April 2020

About Tolly

The Tolly Group companies have been delivering world-class IT services for 30 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 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

Emulex LPe35000

The Emulex® LPe35000 Gen 7 FC HBAs by Broadcom are designed for demanding mission-critical workloads and emerging applications. The family of adapters features Silicon Root of Trust security, designed to thwart firmware attacks aimed at enterprises and governments.

Emulex LPe35000-series HBAs are available with single, dual or quad 32GFC optics to tackle the toughest workloads and NVMe deployments. Gen 7 32GFC provides seamless backward compatibility to 16GFC and 8GFC networks.

Source: Broadcom, April 2020

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.

220114 czse2 -wt 2020-05-01-VerH