TEST REPORT

Tolly.

#220117

April 2020

Commissioned by Broadcom, Inc.

Emulex® Gen 7 LPe35002 Host Bus Adapter

NVMe over Fibre Channel vs. SCSI Performance in VMware ESXi 7 Environments

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.

Fibre Channel leads the industry as the first transport with complete NVMe over Fabrics ecosystem support. Every major server vendor is shipping NVMe standard and offers Emulex NVMe over Fibre Channel (NVMe/FC) capable Host Bus Adapters (HBAs).

Broadcom commissioned Tolly to benchmark the performance of its Emulex Gen 7 LPe35002 32GFC HBAs running NVMe/FC against the performance of the same adapters running in SCSI mode. The series of tests included database TPC-C-like tests using Microsoft SQL Server 2017 and Oracle Database 19c running on Red Hat Enterprise Linux 8 (RHEL 8). The Emulex LPe35002 HBA running in NVMe/FC mode demonstrated superior performance in both test scenarios.

THE BOTTOM LINE

The Emulex LPe35002 HBA using NVMe/FC running on VMware ESXi 7.0 delivered:

1. 2.4x higher transactions compared to SCSI for Microsoft SQL Server 2017
2. 2.1x higher transactions compared to SCSI for Oracle 19c

32G Fibre Channel HBA OLTP Database Workloads
NVMe/FC vs. SCSI: VMware ESXi 7.0 Server
(as reported by HammerDB)

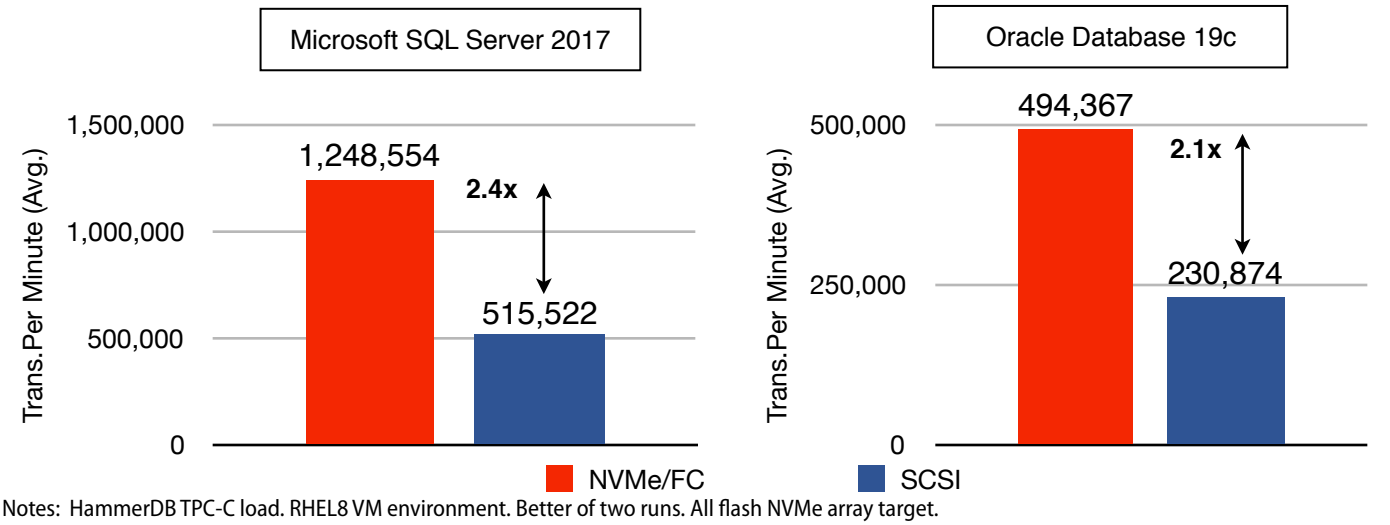

Notes: HammerDB TPC-C load. RHEL8 VM environment. Better of two runs. All flash NVMe array target.

Source: Tolly, April 2020

Figure 1

Overview

Storage vendors are currently shipping NVMe/FC capable arrays and newer Fibre Channel switches fully support NVMe/FC. NVMe/FC leverages the parallelism and performance benefits of NVMe with the enterprise-class storage area network (SAN) technology of Fibre Channel.

A key feature of Emulex Gen 7 HBAs is the ability to support concurrent NVMe/FC and SCSI operation, meaning users can run these technologies simultaneously or they may choose to run either independently. This means that existing Fibre Channel networks can be easily transitioned to NVMe/FC without requiring the purchase of new hardware, unlike other protocols that require investment in new equipment including adapters.

Test Results

Tests were run using SCSI and NVMe for both Microsoft SQL Server 2017 and Oracle 19c database environments. Details are found in the Test Setup & Methodology section of this report.

VMware ESXi 7.0 Server

Benchmarking Microsoft SQL Server 2017 over NVMe/FC resulted in 2.4x higher transaction rate of SCSI. With NVMe/FC, 1,248,554 TPMs achieved versus 515,522 for SCSI.

Similarly, benchmarking Oracle 19c over NVMe/FC resulted in 2.1x higher transaction rate of SCSI. With NVMe/FC, 494,367 TPMs were achieved versus 230,874 for SCSI. See Figure 1.

Emulex

LPe35002 HBA

32G Fibre Channel Host Bus Adapter

VMware ESXi 7.0 NVMe/FC Performance

Tested April 2020

Composite Test Bed Topology

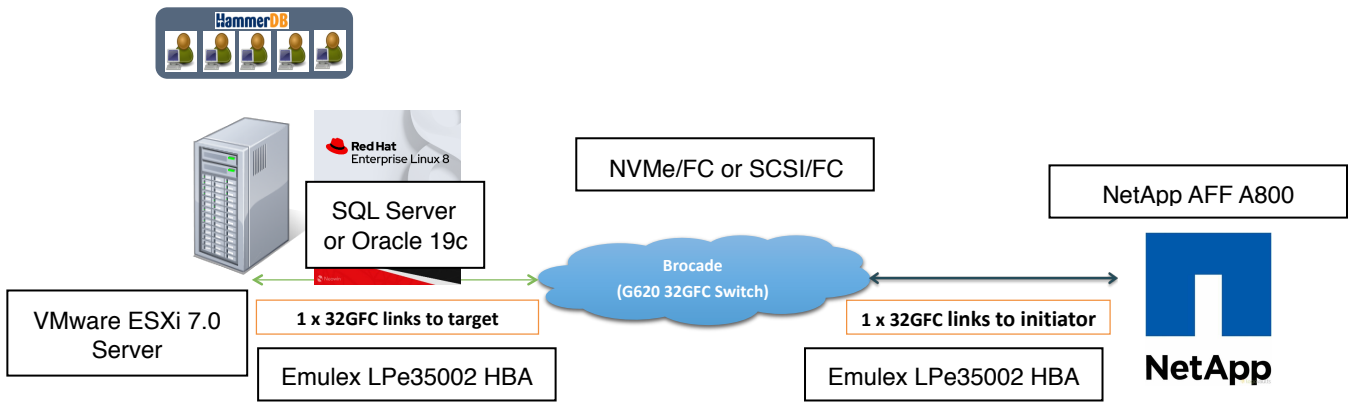

Source: Tolly, April 2020

Figure 2

Test Setup & Methodology

The HBAs under test used current production drivers that are publicly available. Default settings were used. Details of the test environment and systems under test are found in Tables 1-6. Figure 2 shows a composite test environment and indicates the main test variables of protocol (SCSI or NVMe/FC) and database application (Microsoft SQL Server or Oracle 19c).

Server Environment

Tests were run with the database applications running on a RHEL 8.0 virtual machine within a single VMware ESXi 7.0 server.

OLTP Database Test

The goal of this test was to benchmark the database transaction performance of each HBA. An initiator server was configured with the HBA under test. This HBA connected to a NetApp AFF A800 storage target system via a Brocade 32G FC switch. The test utilized a single port.

This test was run using two different database solutions: 1) Microsoft SQL Server 2017, 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 can be found in Tables 1-6. 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 HBA Under Test

Vendor	Product Name	Firmware	Transport Driver	FC Driver
Broadcom	Emulex LPe35002 (32G)	12.4.282.0	12.4.293.2-3vmw.700.1.0.15843807 (brcmnvmefc)	12.4.293.3-5vmw. 700.1.0.15843807 (lpfc)

Table 1

Server Configuration

CPU	2 socket - Intel(R) Xeon(R) Gold 6146 CPU @ 3.20GHz (48 CPUs)
Hyperthreading	Enabled
Memory (RAM)	64 GB
Power Mode	Performance
Hypervisor	VMware ESXi 7.0.0
Version	build-15843807
Guest OS	Red Hat Ent. Linux 8 (RHEL8)
Kernel	4.18.0-80.11.2.el8_0.x86_64

Table 2

Test Tool

Vendor	Open Source
System	HammerDB 3.2
TPC-C Load settings	Warehouses: 1,000 Virtual Users: 48 for SQL, 80 for Oracle Ramp up time: 5 min. Run time: 5 min.

Table 4

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 5

Microsoft Database Configuration

Database	SQL Server 2017
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
VM Configuration	vCPUs: 48, Memory: 50GB

Table 3

Oracle Database Configuration

Database	Oracle Database 19c
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 Oracle DB
VM Configuration	vCPUs: 48, Memory: 50GB

Table 6

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/FC deployments. Gen 7 32GFC provides seamless backward compatibility to 16GFC and 8GFC networks.

Source: Broadcom, April 2020

Terms of Usage

220117 czse2 -wt 2020-05-01 - VerE