

Dell PowerEdge R760 Rack Server & Emulex LPe36002 Host Bus Adapter

64G Fibre Channel Oracle Database/Microsoft SQL Performance - NVMe/FC vs. SCSI/FC

EXECUTIVE SUMMARY

New generation servers can bring higher performance across a range of areas. This is certainly the case with Dell's 16th-generation server line. Similarly, newer protocols like NVM Express (NVMe) over Fibre Channel (FC) can provide greater throughput and efficiency than older SCSI over FC. Dell is unique in offering an end-to-end NVMe/FC connectivity solution in the mid-range storage marketplace with the PowerStore line.

Dell commissioned Tolly to benchmark the performance of the Broadcom Emulex LPe36002 64G Fibre Channel dual-port host bus adapter (HBA) running in the Dell PowerEdge R760 Rack Server by testing on actual database applications rather than simulated IO microbenchmarks. Testing focused on evaluating the database throughput, latency, and CPU efficiency of accessing both Oracle Database 19c and Microsoft SQL Server 2019 for Linux systems over older SCSI/FC and newer NVMe/FC. Databases were stored on a Dell PowerStore 9200T storage appliance.


Tests showed significant improvements in transaction throughput, latency reduction, and CPU efficiency. See Figure 1 for a summary of TPM improvements.

THE BOTTOM LINE

Dell PowerEdge R760 & Emulex LPe36002 64G HBA using NVMe/FC:

- 1 Improved database transactions by up to 53%
- 2 Reduced database stored procedure latency by up to 58%
- 3 Improved server CPU efficiency by up to 106%

Dell R760 & Emulex LPe36002 64G HBA Oracle Database 19c & Microsoft SQL Server 2019
Transactions Per Minute Improvement: NVMe/FC vs. SCSI/FC
 (as reported by HammerDB TPROC-C v4.7)



Notes: Test not designed as an Oracle/Microsoft comparison. Configurations differ and are not comparable. Goal is to show relative improvement.
 Source: Tolly, May 2023

Figure 1



Overview

The goal of this test was to illustrate the performance benefits of using the newer, more-efficient NVMe/FC protocol in lieu of the older, less-efficient SCSI/FC protocol in conjunction with Emulex 64G FC HBAs running under Linux in a Dell PowerEdge R760 Rack Server. (Dell sells the Emulex 64G FC HBA for the same price as the Emulex 32G FC HBA.)

The test was run using two popular database offerings to illustrate that the benefits were not confined to a particular database environment. All results should be viewed within a single database environment. While the Oracle and Microsoft test environments were similar, they were neither identical nor comparable and, thus, results should not be compared across database environments.

While low-level component benchmarks are instructive, ultimately system architects are rightly most interested in how network-level improvements can translate into application performance improvements. This benchmarking was done with HammerDB which generates actual user transactions against an actual database. The test was focused on TPROC-C which is the HammerDB, database-oriented implementation of the de facto standard TPC-C online transaction processing benchmark.

Tests showed 28% to 106% improvements in key benchmarks across both database environments. Details below.

SCSI/FC protocol was 541,751. When the test was re-run using the NVMe/FC protocol, the average TPM was 827,979 which was an improvement of 53%. See Figure 2 for all Oracle Database 19c results.

CPU Efficiency. This test metric showed the number of transactions processed for each one percent of CPU used. A higher number indicating more transactions per unit of CPU shows greater efficiency.

The average transactions per unit of CPU using the SCSI/FC protocol was 10,206. When the test was re-run using the NVMe/FC protocol, the average transactions per unit of CPU was 20,983 which was an improvement of 106%.

P95 Stored Procedure Latency. Delay is a synonym for latency. Thus, lower latency (less delay) is always better. "P95" is a more rigorous measure of latency and measures the time (in ms) for 95% of all

Test Results

Oracle Database 19c

Transaction Processing. The average transactions per minute (TPM) using the

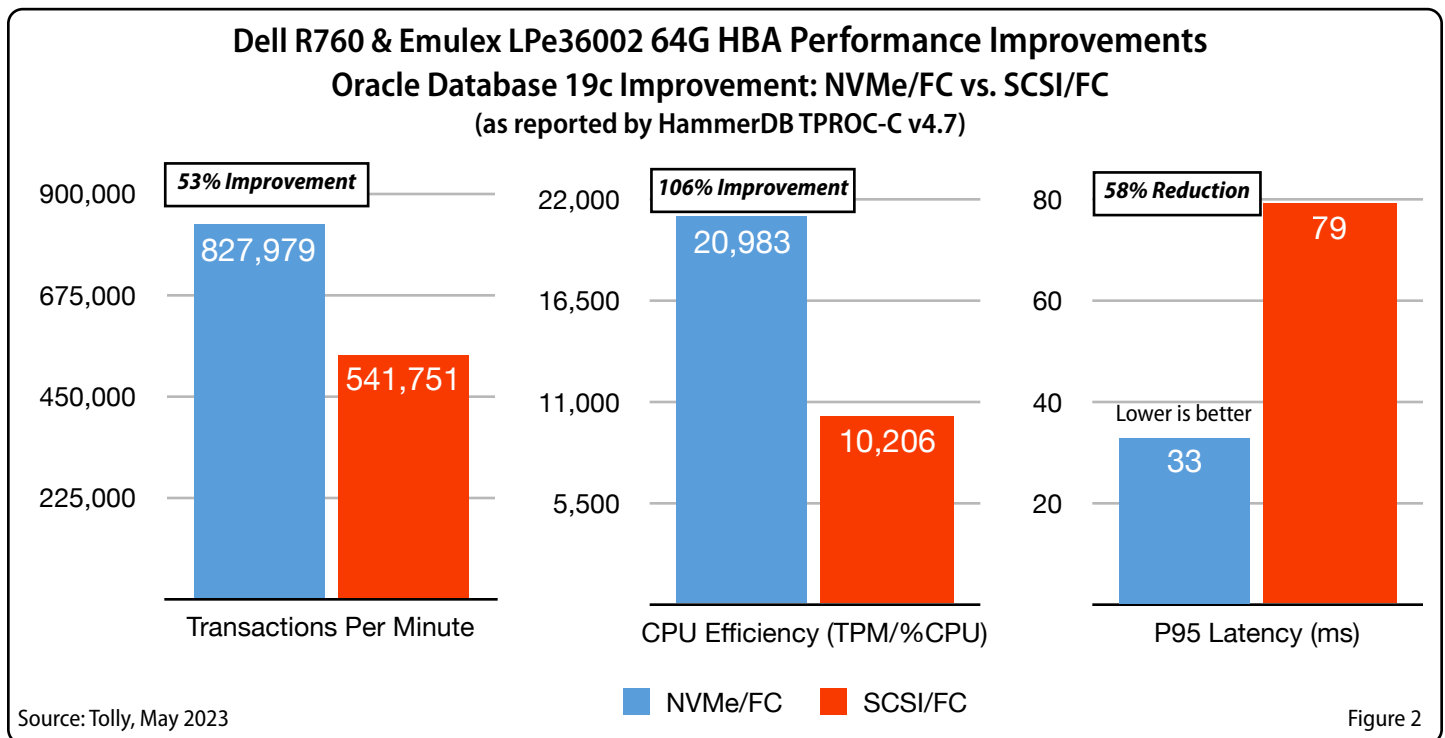


Figure 2



results. Thus, it is a more rigorous measure of latency than “average” latency.

The P95 latency using the SCSI/FC protocol was 79ms. When the test was re-run using the NVMe/FC protocol, the P95 latency was 33ms which was an improvement (i.e., reduction) of 58%.

Microsoft SQL Server 2019 for Linux

Transaction Processing. The average transactions per minute (TPM) using the SCSI/FC protocol was 625,648. When the test was re-run using the NVMe/FC protocol, the average TPM was 826,736 which was an improvement of 32%. See Figure 3 for all Microsoft SQL Server 2019 for Linux results.

CPU Efficiency. The average transactions per unit of CPU using the SCSI/FC protocol was 22,449. When the test was re-run using the NVMe/FC protocol, the average transactions per unit of CPU was 32,044 which was an improvement of 43%.

P95 Stored Procedure Latency. The P95 latency using the SCSI/FC protocol was 48ms. When the test was re-run using the NVMe/FC protocol, the P95 latency was 37ms which was an improvement of 28%.


Test Setup & Methodology

The HBA under test used current production drivers that are publicly available. Default settings were used. Details of the test environment and

Dell, Inc.

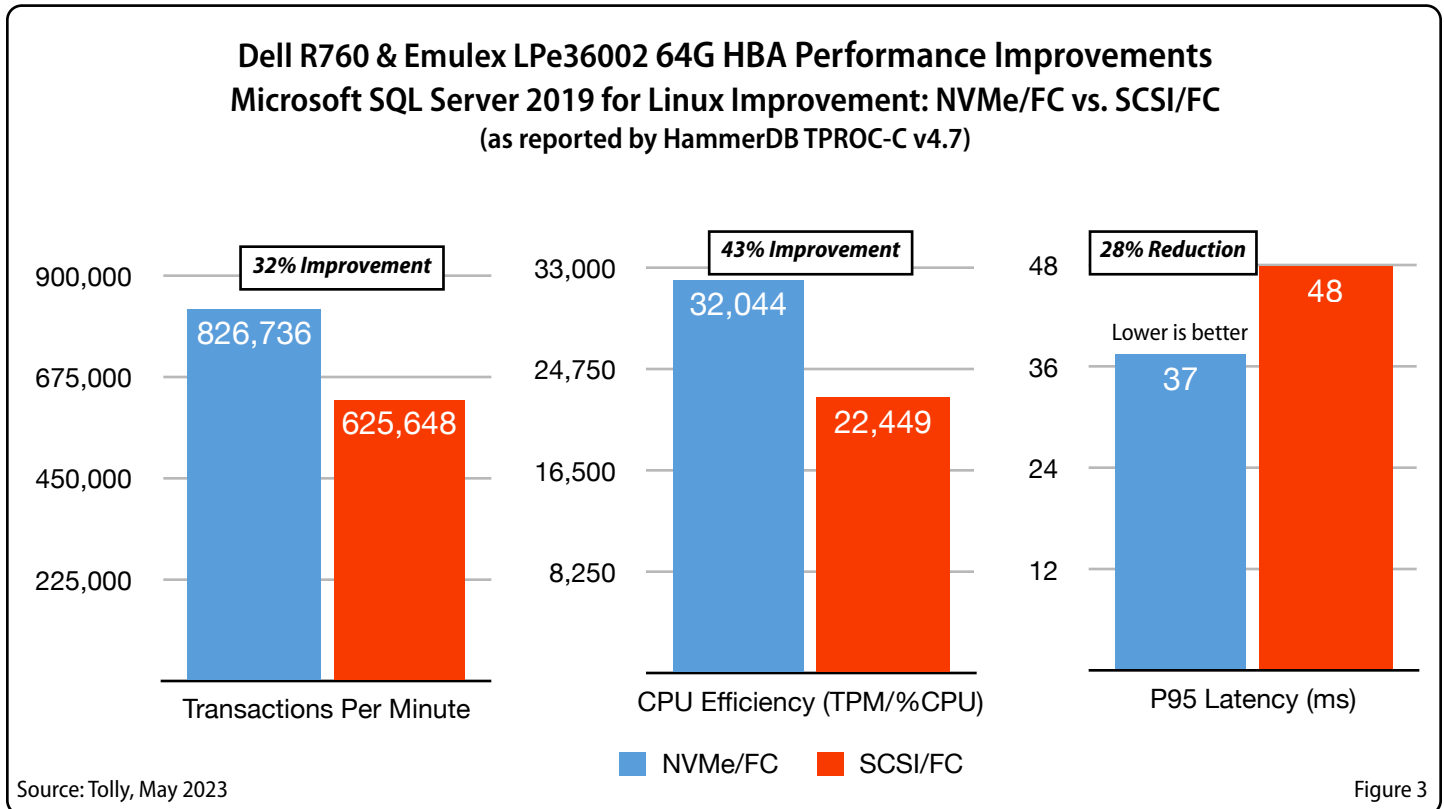
Dell PowerEdge R760 & Emulex LPe36002 HBA

64G Fibre Channel Database Performance

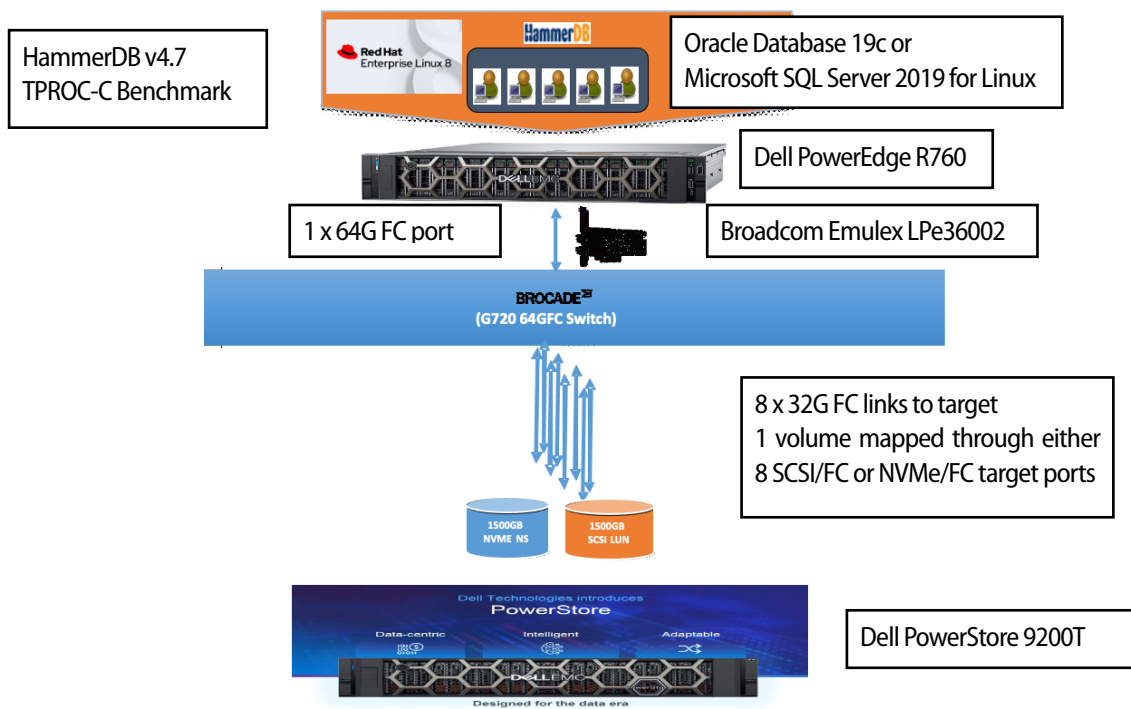


Tested May 2023

systems under test are found in Tables 1-5. Figure 4 shows a composite test environment.



Composite Test Bed Topology



Source: Tolly, May 2023

Figure 4

Database Test

The goal of this test was to benchmark the database transaction performance of each HBA running the HammerDB “TPROC-C” workload which, as noted earlier, is the HammerDB, database version of the Transaction Processing Council’s TPC-C OLTP benchmarked

A Dell PowerEdge R760 server was configured with the HBA under test. The Broadcom Emulex LPe36002 64G HBA connected to a Dell PowerStore 9200T via a Broadcom Brocade 64G Fibre Channel

switch. The test utilized a single 64G FC port of the Emulex HBA.

The server ran RHEL 8.7. SCSI Device Mapper and NVMe native multipath were enabled for the respective devices. NUMA was set to off and “transparent huge pages” was disabled.

For storage, path selection policy for NVMe native multipath was set to “round-robin”.¹ For SCSI Device mapper multipath was set to “queue-length 0”.²

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

The open source HammerDB test tool was used to populate the database schema and run the workload.

There were some differences in the database and test configuration between the Microsoft and Oracle database tests and those details are noted in the aforementioned tables. Because of those differences in configuration the reader should not compare results across database environments.

¹ Configuration followed the Dell the recommended multipath configuration for NVme/FC according to the latest version of the FC-NVMe support matrix. https://elabnavigator.dell.com/vault/pdf/NVMe_Connectivity.pdf

² Configuration followed the Dell recommended multipath configurations for SCSI according to PowerStore documentation <https://www.dell.com/support/manuals/en-us/powerstore-5000/pwrstr-hcg/configuration-with-device-mapper-multipathing-for-scsi?guid=guid-f78d5743-d77b-458f-b883-6f794b269231&lang=en-us>



Test Configuration Summary

HBA Under Test

Vendor	Product Name	Firmware	Driver
Broadcom	Emulex LPe36002 (64G) (PCIe 4.0)	12.8.497.10	14.0.0.15

Table 1

Server Configuration

Vendor/System	Dell PowerEdge R760
CPU	2 socket Intel(R) Xeon(R) Platinum 8468 @ 2.1 GHz 192 logical processors
Logical Processors	Enabled
Memory (RAM)	256 GB
Power Mode	Performance
OS	Red Hat Ent. Linux 8.7 (RHEL8)
Kernel	4.18.0-425.3.1

Table 2

Database Test Tool

Vendor	Open Source
Application	HammerDB 4.7
TPROC-C settings	Total # of Warehouses = 1,000 Transactions per user = 1 million Number of virtual users (Oracle) = 250 Number of virtual users (Microsoft) = 200 Ramp-up time: 2 minutes Run time: 5 minutes

Table 5

Storage Configuration

Vendor/Device	Dell PowerStore 9200T v3.5
Ports	8 x 32G FC
Volume Size	1,024GB volume each for NVMe/FC and SCSI/FC
Namespace/LUN (Microsoft SQL)	8 x 32G Target ports per Namespace
Namespace/LUN (Oracle 19c)	8 x 32G Target ports per Namespace
Network Fabric	Brocade G720 v9.0.1a

Table 6

Microsoft Database Configuration

Database	Microsoft SQL Server 2019 for Linux
Storage	Single volume, XFS
Dataset Size	100 GB
DB Memory Allocation	10G

Table 3

Oracle Database Configuration

Database	Oracle Database 19c
Storage	Oracle Grid 19c, ASM disk group with external redundancy, Oracleasm-lib-2.0.17-1
Dataset Size	100 GB, Temp table = 100GB, TPROC-C table space = 400GB
Database Settings	SGA = 10GB PGA = 6GB Block size = 8KB Set db for performance, disable background processes and change snapshot policies

Table 4

Source: Tolly, May 2023



About Tolly

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

Broadcom Emulex LPe36002

The Broadcom Emulex LPe36000-series Gen 7 Fibre Channel HBAs 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.

Gen 7 64G provides seamless backward compatibility to 32G and 16G networks.

Dell sells the LPe36002 64G HBA for the same price as the 32G model.

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.