TEST REPORT

Tolly.

#221123

May 2021

Commissioned by Broadcom, Inc.

Emulex® Gen 7 LPe36002 Host Bus Adapter

64G Fibre Channel Performance vs. 32G and 16G HBA Modes

EXECUTIVE SUMMARY

Data storage and retrieval is growing at an extraordinary pace. Analysts note that data grew world-wide by 4x in recent years and that database data warehousing will double in the next five years. Business Intelligence (BI) and Artificial Intelligence (AI) can dramatically increase data retrieval rates.

Next-generation Fibre Channel alone and in conjunction with next-generation PCIe 4.0 technology can provide dramatic improvements in application data throughput. Broadcom's Emulex Gen 7 LPe36002 64G Fibre Channel (64G) dual-port host bus adapter (HBA) leverages both 64G and PCIe 4.0.

Broadcom commissioned Tolly to benchmark the performance of its new adapter in terms of raw performance and compared to existing 32G Fibre Channel (32G) HBAs. The adapter demonstrated dual-port performance in excess of 11 million IOPS. Running relational database transaction benchmarks, the adapter effectively doubled performance by dramatically cutting run times. Similarly, the adapter leveraged PCIe 4.0 technology increasing application throughput to ~12,000MBPS compared to ~6,600MBPS maximum achieved on PCIe 3.0 servers.

THE BOTTOM LINE

The Emulex LPe36002 Dual-port 64G HBA:

1. Delivered greater than 11 million IOPS aggregate, dual-port performance, the highest performing Fibre Channel HBA Tolly has tested to date
2. Effectively doubled database application performance for Oracle Database 19c compared to 32G
3. Demonstrated ~12,000MBPS PCIe 4.0 throughput compared with ~6,600MBPS on PCIe 3.0 systems
4. Completed Oracle Database 19c queries 67% faster in a PCIe 4.0 server versus a PCIe 3.0 server

LPe36002 64G HBA Dual-Port Aggregate IOPS & Throughput

(as reported by Medusa Labs Test Tool Suite v7.4)

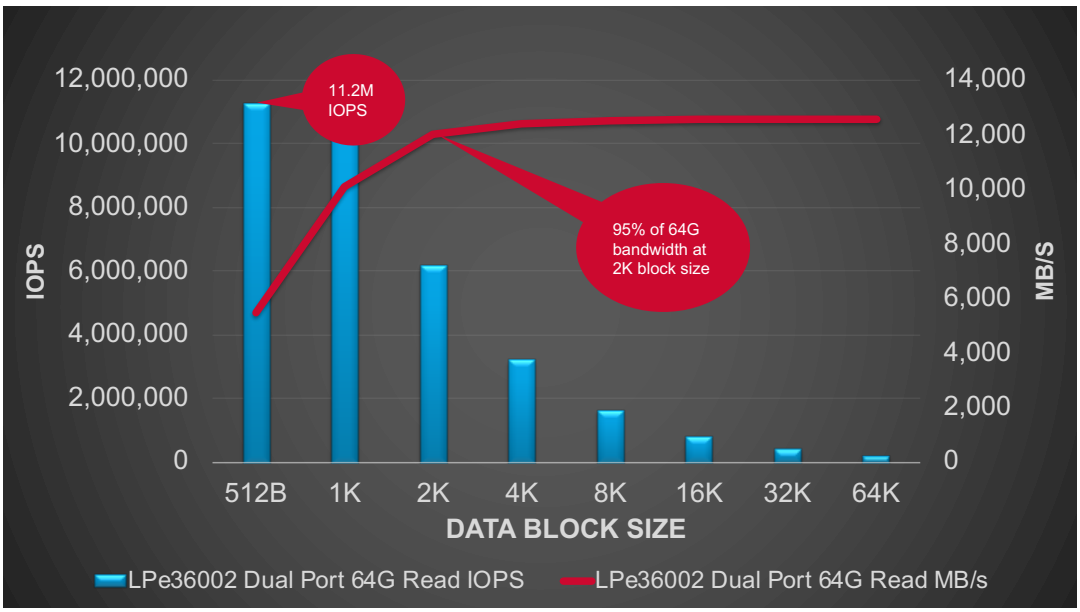

Notes: Server with PCIe 4.0 running Red Hat Linux Server 8.3 to 24 open source NVMe-FC SPDK targets.

Source: Tolly, April 2021

Figure 1

Overview

Testing profiled three different aspects of performance: 1) Raw performance vs. 32G, 2) Database application performance improvement vs 32G and 16G Fibre Channel (16G), and 3) Leveraging state-of-the-art PCIe 4.0 server architecture to increase throughput vs. PCIe 3.0.

As will be shown, the performance benefits in each of the aforementioned scenarios are dramatic with throughput roughly doubling and application run times roughly cut in half.

Of note is that the benefits of using the 64G adapter in the server were seen in an environment where the "back end" storage was connected via a 32G fabric. The 64G server connections were able to aggregate the data from multiple storage connections to achieve new heights in performance benchmarking.

The final test illustrated how the Emulex HBA can leverage the bandwidth-doubling benefits of PCIe 4.0 to double the throughput of existing applications.

Test Results

See the Test Methodology and Setup section for additional details of each test.

Dual-Port IOPS/ Throughput

This test used the Medusa Labs Test Suite to establish the IOPS and throughput performance of the HBA across a range of data block sizes from 512 bytes through 64KB. Lower sizes stress the IOPS capabilities where larger sizes test data throughput limits.

With a block size of 512 bytes, each port achieved performance in excess of 5,620,000 IOPS for an aggregate of 11,240,975 IOPS for the HBA. Broadcom claims 10 million IOPS for the 64G HBA and testing illustrates that claim to be conservative.

With 2KB data block sizes, the aggregate throughput reached ~95% of line rate. The two ports delivered a combined throughput of 11,998 MBPS. At 4KB data block sizes and higher, the HBA achieved line rate with aggregate throughput ranging from 12,378 MPBS at 4KB to 12,586 MBPS at 64KB. See Figure 1.

Database Transaction Performance

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.

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

This test benchmarked the run time of the HammerDB TPC-H Decision Support series of transactions. This consists of running 22 separate query sequences against the test database[1].

The test was run with the appropriate Emulex HBA configured to one of three different FC options: 64G, 32G and 16G to illustrate the benefits of higher speed Fibre Channel.

Emulex

LPe36002 HBA

64G Fibre Channel Host Bus Adapter Performance

Tolly. Certified

Tested April 2021

The test was run using a single port of the HBA and tested separately with the Oracle and Microsoft database solutions.

In the test of Oracle Database 19c, the test suite completed in 1,201 seconds with 16G, 608 seconds with 32G and only 315 seconds with 64G.

Over 64G the test ran 1.93x faster than over 32G and 3.81x faster than over 16G. See the upper chart of Figure 2.

Similarly, in the test of Microsoft SQL Server 2019, the test suite completed in 809 seconds with 16G, 499 seconds with 32G and only 286 seconds with 64G.

Over 64G the test ran 1.74x faster than over 32G and 2.83x faster than over 16G. See the lower chart of Figure 2.

[1] Because these are actual queries against the database, there are pauses in the the network traffic as data is being retrieved from the database in between the test queries.

LPe36002 64G HBA DSS Warehousing Performance: 64G vs. 32G vs. 16G

(as reported by HammerDB v3.3)

(Shorter runtime and higher throughput better)

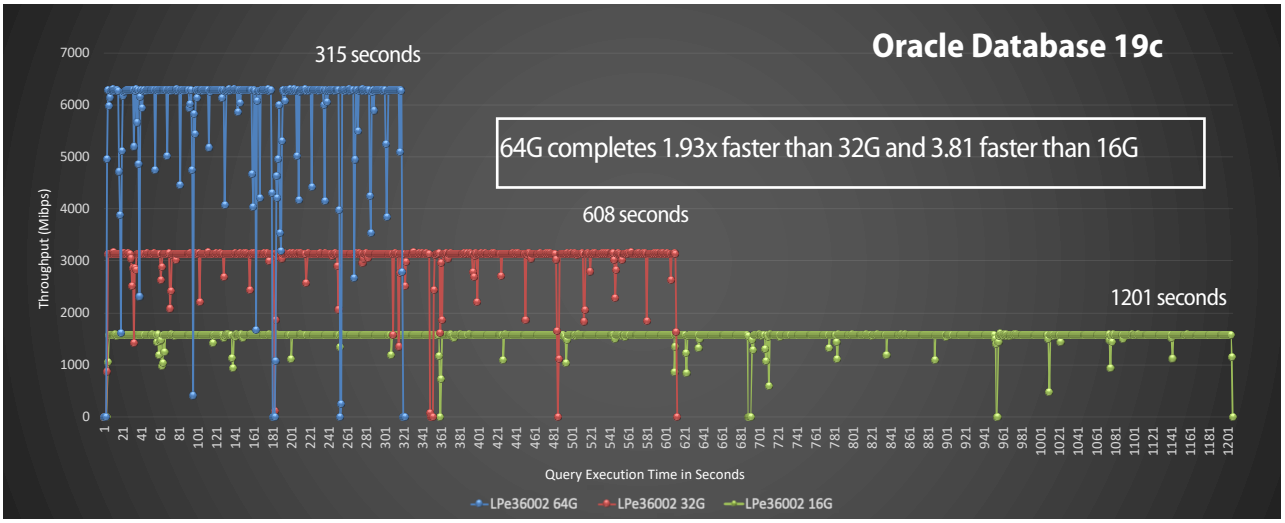

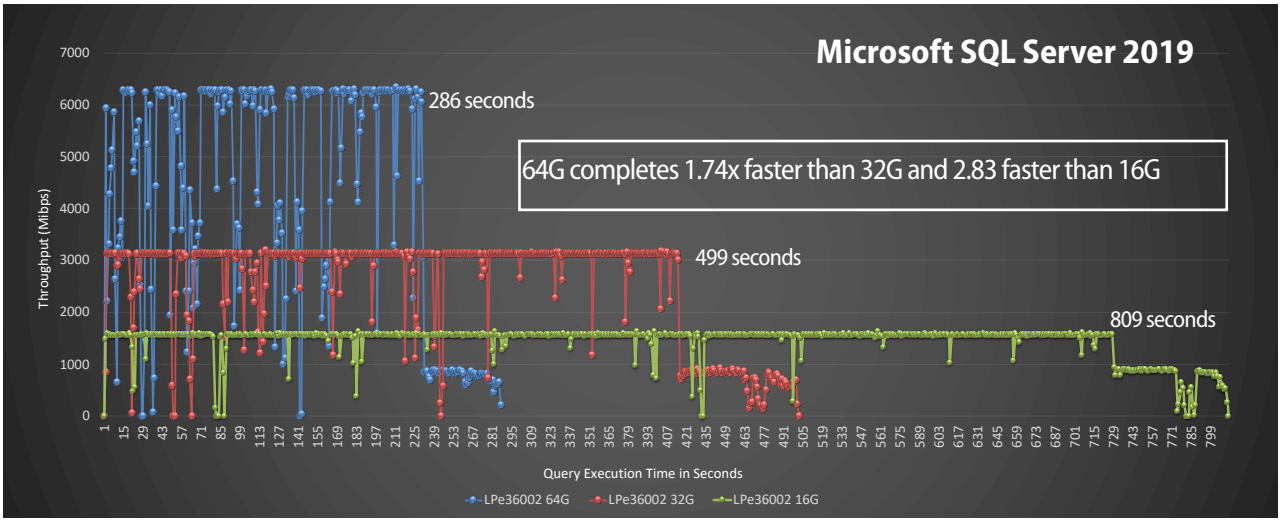

Source: Tolly, April 2021

Figure 2

LPe36002 64G HBA Dual-Port DSS Warehousing Performance: PCIe 4.0 vs. PCIe 3.0

(as reported by HammerDB v3.3)

(Shorter runtime and higher throughput better)

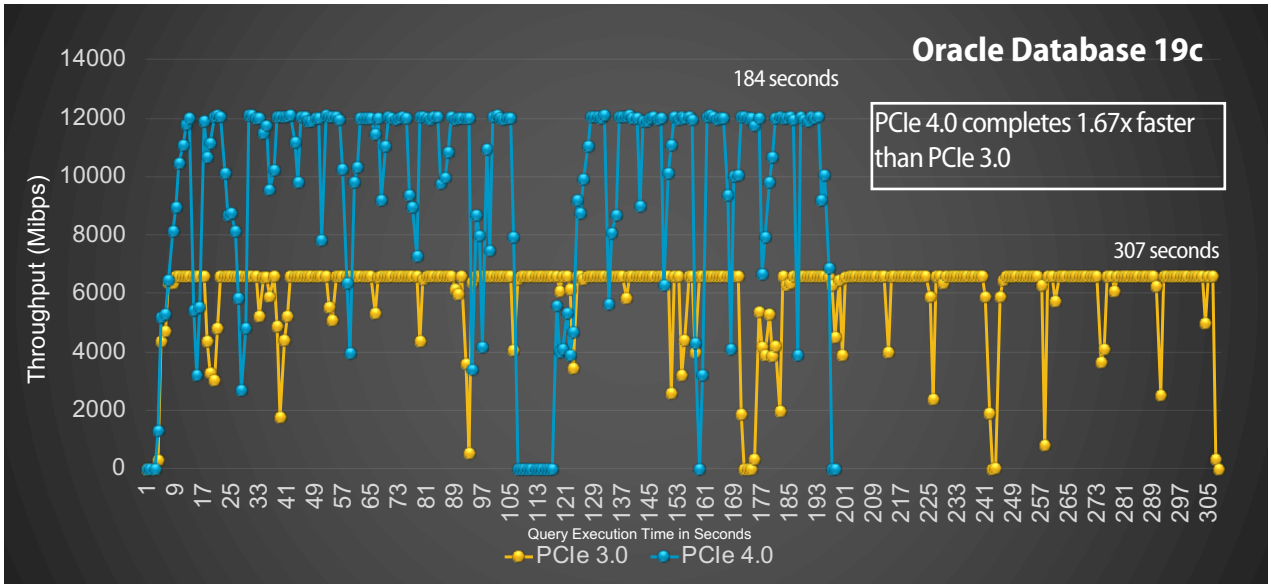

Source: Tolly, April 2021

Figure 3

PCIe 4.0 vs. PCIe 3.0 Performance

The LPe36002 is designed as an 8-lane PCIe 4.0 host bus adapter but can run in a PCIe 3.0 server slot although throughput levels when running both ports at 64G are best suited for PCIe 4.0 servers.

In this test, the DSS Oracle database test was run using a 64G dual-port adapter in two different servers: 1) PCIe 3.0, and 2) PCIe 4.0.

In the test of PCIe 3.0, the test suite completed in 307 seconds. The data throughput during the test hovered around 6,600 MBPS - just slightly higher than the bandwidth of two ports of 32G or one port of 64G.

In the test of PCIe 4.0, the test suite completed in only 184 seconds. The data throughput during the test hovered around 12,000 MBPS.

On the PCIe 4.0 server the test ran 1.67x faster than on the PCIe 3.0. See Figure 3.

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 systems under test are found in Tables 1-6. Figure 4 shows a composite test environment.

Dual-Port IOPS/ Throughput Test

The goal of this test was to benchmark the maximum input/output operations per second (IOPS) of the adapter under test using both ports simultaneously.

The test was run at all of the following block sizes: 512B, 1K, 2K, 4K, 8K, 16K, 32K, and 64K. The test measured the aggregate IOPS and throughput in Megabytes per second at each block size in separate read and write operations.

Viavi Solutions, Inc. Medusa Labs Test Tool Suite v7.4 was used to run the test. Option "nomerge=2" was used for all testing to disable merging of I/O operations to assure accurate results.

The adapter under test was run in a server with 72 cores running Red Hat Enterprise Linux 8.3. For the back end storage we used the open-sourced Storage Performance Development Kit (SPDK) running on servers using Emulex LPe35002 32G HBAs. Four SPDK storage servers were used for a total of 24 32G Target Ports serving 24 NVMe-FC Namespaces.

Test Configuration Summary

HBA Under Test

Vendor	Product Name	Firmware	Driver
Broadcom	Emulex LPe36002 (64G)	12.8.351.37	12.8.351.29

Table 1

Server Configuration (All except PCIe 3.0 Test)

CPU	2 socket - Intel(R) Xeon(R) Platinum 8360Y CPU @ 2.4GHz (72 CPUs)
Hyperthreading	Enabled
Memory (RAM)	448 GB
Power Mode	Performance
OS	Red Hat Ent. Linux 8.3 (RHEL8)
Kernel	4.18.0-240.el8.x86_64

Table 2

Microsoft Database Configuration

Database	Microsoft SQL Server 2019
Storage	4 disk LVM, 128K stripe, XFS
Dataset Size	100 GB
DB Memory Allocation	60G

Table 3

Oracle Database Configuration

Database	Oracle Database 19c
Storage	4 disk LVM, 128K stripe, XFS
Dataset Size	100 GB
Database Settings	SGA = 30GB PGA = 20GB Block size = 8KB FILESYSTEMIO_OPTIONS=SETALL

Table 4

Database Test Tool

Vendor	Open Source
System	HammerDB 3.3
TPC-H settings	TPC-H - MAXDOP = 24 TPC-H - Scale Factor = 100 TPC-H - Virtual User = 1

Table 5

Storage Configuration

Vendor/Device	Generic 2 Nodes/Clusters
Release	N/A
Namespace/LUN (Microsoft SQL)	4 400GB FC-NVMe Namespaces, 2 32G Target ports per Namespace
Namespace/LUN (Oracle 19c)	4 400GB FC-NVMe Namespaces, 2 32G Target ports per Namespace
Network Switch	Brocade G720 Gen 7 64G FC switch

Table 6

IOPS Test Configuration

Test Tool	Viavi MLTT version 7.4
I/O Type	Random 100% Read
Queue Depth	8
Threads	8
Storage	4 servers using SPDK 21.01 FC-NVMe inbox BRCMFC driver, 24 32LPe35002 32G ports/Namespaces

Table 7

Source: Tolly, April 2021

DSS Database Test

The goal of this test was to benchmark the database transaction performance of each HBA running the HammerDB "TPC-H" DSS Data Warehousing workload. An initiator server was configured with the HBA under test. The Emulex LPe36002 64G HBA connected to an enterprise all flash storage array via a Brocade 64G Fibre Channel switch. The test utilized a single port.

Composite Test Bed Topology

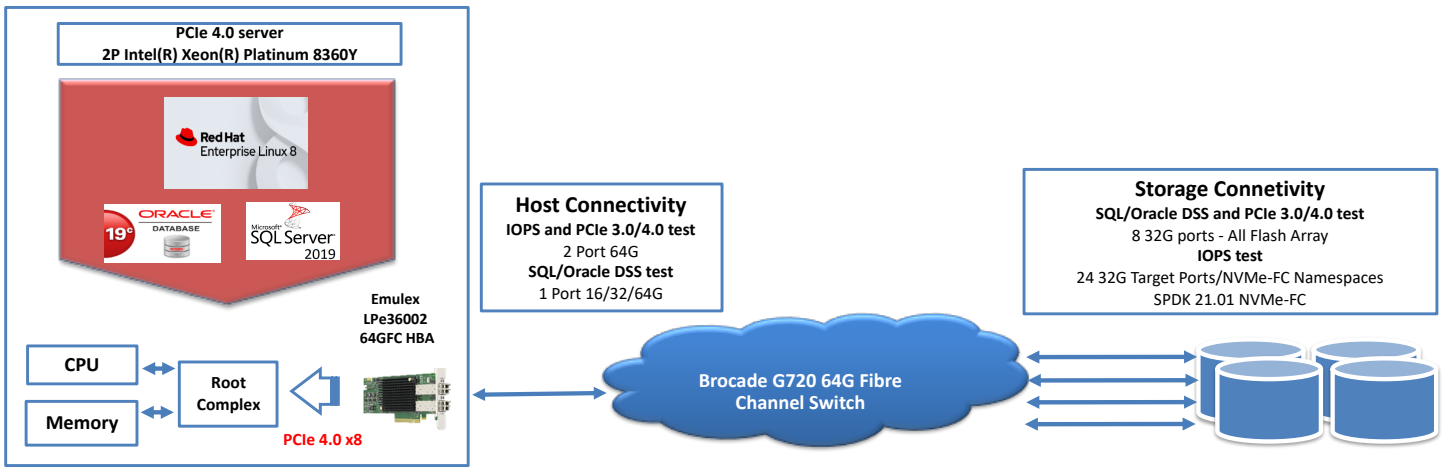

Source: Tolly, April 2021

Figure 4

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

The open source HammerDB 3.3 test tool was used to populate the database schema and run the workload using its TPC-H load test option. Additional details of the test environment can be found in Tables 1-7. 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.

PCIe 4.0 vs. PCI 3.0 Comparison Test

This test was run to illustrate the differences in performance when running the Emulex 64G HBA in a PCIe 3.0 vs. PCIe 4.0 server.

The test reproduces the database test for Oracle using dual-ports. The test was run in a server supporting PCIe 4.0 and again in a server supporting only PCIe 3.0 slots.

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

Emulex LPe36000-series

The Emulex® LPe36000-series Gen 7 Fibre Channel 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 LPe36000-series HBAs are available with single or dual 64G optics to tackle the toughest workloads and NVMe/FC deployments. Gen 7 64G provides seamless backward compatibility to 32G and 16G networks.

Terms of Usage

221123 czse2 -wt 2021-05-30 - Verl