TEST REPORT

Tolly.

#223133

September 2023

Commissioned by Broadcom

Emulex LPe36002 Host Bus Adapter & Intel Eagle Stream Servers

64G Fibre Channel Enables up to 4:1 Server Consolidation

EXECUTIVE SUMMARY

With each new generation of technology performance improvements are expected. There are times, though, where technological advancements on multiple fronts produce very dramatic benefits. This the case when the Broadcom Emulex LPe36002 64G Fibre Channel Host Bus Adapter (HBA) is paired with Intel's new Eagle Stream servers.

Broadcom commissioned Tolly to benchmark the performance of the Broadcom Emulex LPe36002 64G Fibre Channel dual-port host bus adapter (HBA) running in an Intel Eagle Stream-class server (hereafter Eagle Stream) and compare that to the same combined workload performance running in four separate, Intel Purley-class servers (hereafter Purley) each outfitted with a 16GFC HBA as was standard with that server generation.

Tests showed that the new Intel Eagle Stream-based platform's increased CPU power and improved memory performance/capacity improved multiple workloads and provide an environment where a database application can push the Emulex 64G FC HBA to full line rate performance of 64GFC thus matching the combined application throughput of four Purley platform servers using 16GFC HBAs. See Figure 1.

THE BOTTOM LINE

Intel Eagle Stream server and Emulex LPe36002 64G HBA benefits over older gen servers with 16G HBAs:

1. 1x Eagle Stream server with 64GFC HBA can achieve same TPROC-H query throughput compared to 4x Purley servers with 16GFC HBA
2. Effective 4:1 VMware server consolidation is demonstrated by 4 2-VM Purley servers with 16GFC HBA VM performing storage migration in the same time as a single 8-VM Eagle Stream server with 64GFC HBA
3. Consolidating Oracle DSS workloads from 4 Purley servers with 16GFC HBA to a single Eagle Stream server with 64GFC can significantly reduce I/O bound TPROC-H query time

Eagle Stream Server & Emulex LPe36002 64G HBA Oracle Database 19c
Query Time Improvement: Eagle Stream vs Four Purley Servers
(as reported by HammerDB TPROC-H v4.7)

Bars above the black line show the run times of four separate Purley-class servers with 16G FC running against a single data store.

Bar below the black line shows the Eagle Stream server with 64G FC running the same test against the same data store.

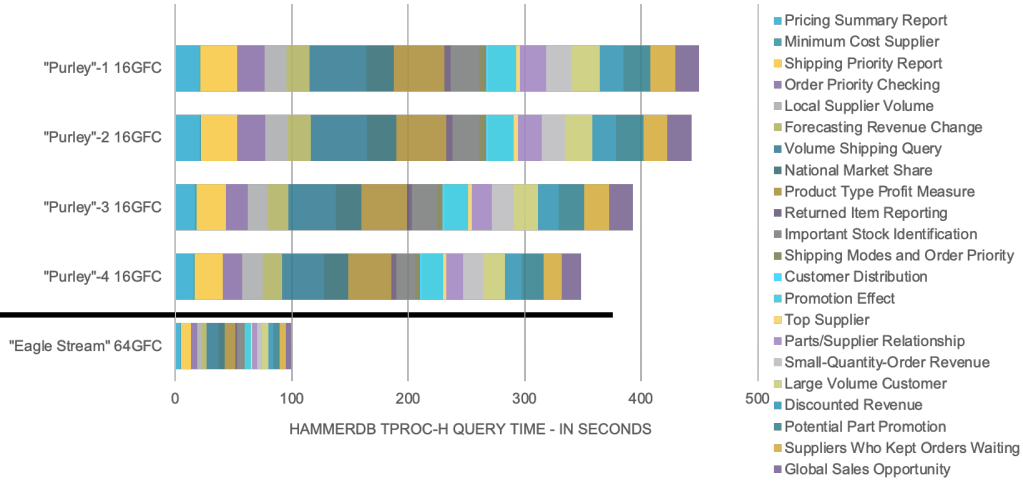

Notes: Systems running RHEL 8.8 on ESXi 8. "Purley-class" machines all ran the same generation Intel CPU and Broadcom Emulex 16G Fibre Channel adapters.
Source: Tolly, June 2023

Figure 1

Overview

The goal of these tests was to illustrate, simply, that a single Intel Eagle Stream server, using a single port of a PCIe 4-based, dual-port Emulex 64GFC can equal the I/O throughput of four individual, older generation, Purley servers each using a single port of a 16GFC HBA. Both database and VMware workloads were benchmarked. Additionally, the database benchmark was run on an Eagle Stream server using 16/32/64GFC network connections to illustrate the direct benefit of higher bandwidth in an Eagle Stream server.

The Purley servers use older, less powerful CPUs and use 16GFC HBAs that offer, at best, 25% of the 64G FC HBA's throughput. The HBAs are constrained by the bandwidth of the PCIe 3 bus architecture which would limit the benefits of using the higher FC speed HBAs in the older servers.

The broader point is that this significant performance improvement means that, for I/O-bound applications, a single Eagle Stream server can be used to replace and consolidate the workloads and operating expenses of up to four older servers thus improving the overall energy efficiency of the datacenter.

Query Time Improvement Results

The same test was run on all of the servers and consisted of running the TPROC-H analytics workload of HammerDB.[1] The tests were run using the Oracle 19c database environment but the results are generally applicable to any database or other input/output intensive workload.

The TPROC-H workload measures how long it takes to run a series of 22 different types of decision support queries. This type of workload is "read only" with no database updates taking place.

The test was run using two different scenarios. In the first scenario, four of the older servers ran the HammerDB benchmark simultaneously against the same NVMe-based, all-flash array. In the second scenario, the single Eagle Stream server ran the benchmark against the same data store.

Page 1, Figure 1, above the horizontal dividing line, summarizes results of the first scenario. Because those servers

Eagle Stream Server & Emulex LPe36002 64G HBA Performance Improvements Storage Migration: Eagle Stream vs 4 Purley Servers (as reported by VMware ESXi 8)

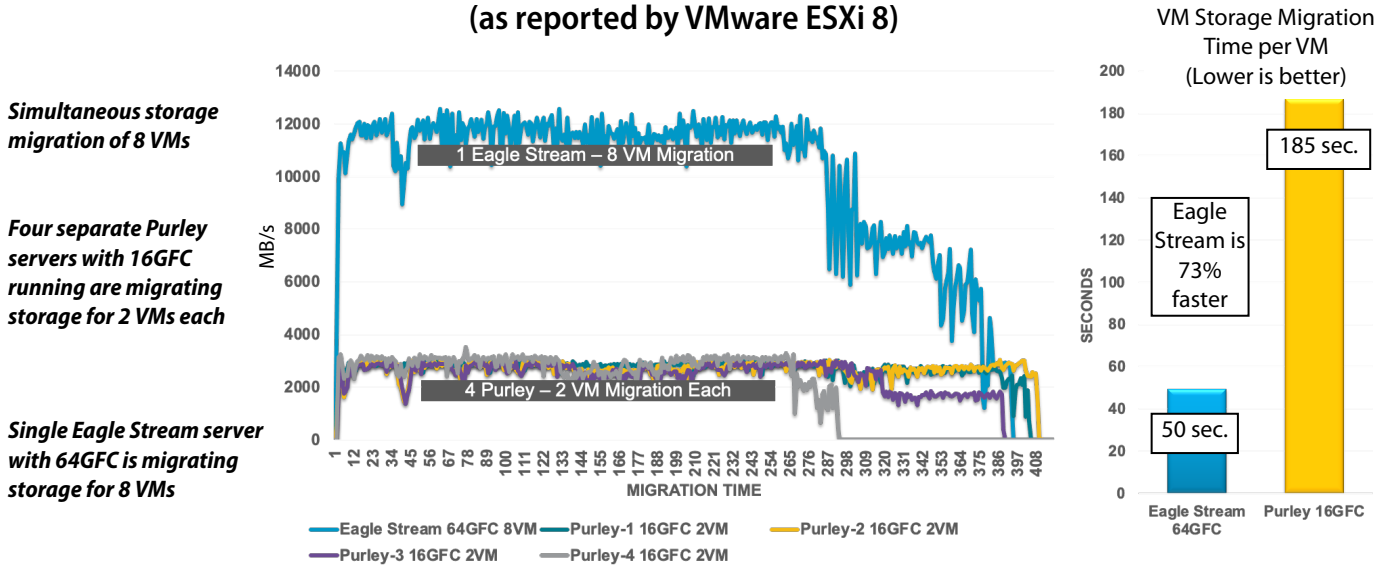

Source: Tolly, June 2023

Figure 2

[1] https://www.hammerdb.com/docs/ch11.html

were using 16GFC HBAs, 16G was the theoretical maximum for network I/O and, thus a potential bottleneck for each server. As each server finished the test, the reduced load on the target data store allowed subsequent server's tests to run more quickly. The fastest completion time was 335 seconds and the slowest was 448 seconds with the average being 405.5 seconds.

VM Storage Migration Results

Server virtualization is an important part of IT infrastructure for countless businesses and organizations worldwide. Efficient use of the underlying server hardware components is an important aspect of providing high quality end-user experience while controlling costs. Certain elements of server virtualization can place a tremendous load on I/O resources.

VMware documentation calls this "Storage vMotion" which refers to the storage assigned to a VM to be migrated from one data store to a different data store. VMware vCenter Server limits the number of simultaneous storage migrations to eight, thus the maximum for this test.

The test was run using two different scenarios. In the first scenario, four of the older servers each migrated the storage for two VMs simultaneously against the same physical data store. In the second scenario, the single Eagle Stream server migrated the storage of eight VMs simultaneously against the same physical data store.

Figure 2, on the previous page, summarizes the results of the storage migration tests in terms of storage I/O and migration time. Here, too, the I/O throughput difference between the single 64GFC server and the four 16GFC servers is readily apparent. Where each of the older servers delivers throughput of peaks around 3,000MB/s, the 64GFC server throughput was sustained at approximately 12,500MB/s. This increase in throughput on the 64GFC Eagle Stream server results in dramatically faster storage migration times for each of the VMs tested. As shown in the the figure, the average, per-VM storage migration time for VMs running on the Purley systems was 185s. The the average, per-VM storage time for VMs running on the Eagle Stream system was 50s.

Figure 3 illustrates the networking flow of the four older generation servers and the Eagle Stream across the Broadcom Brocade 64G Fibre Channel switch in both the database and virtualization tests.

Broadcom

Emulex LPe36002 HBA & Intel Eagle Stream Servers

64G FC Server Consolidation

Tested June 2023

Server Consolidation Test Bed Topology

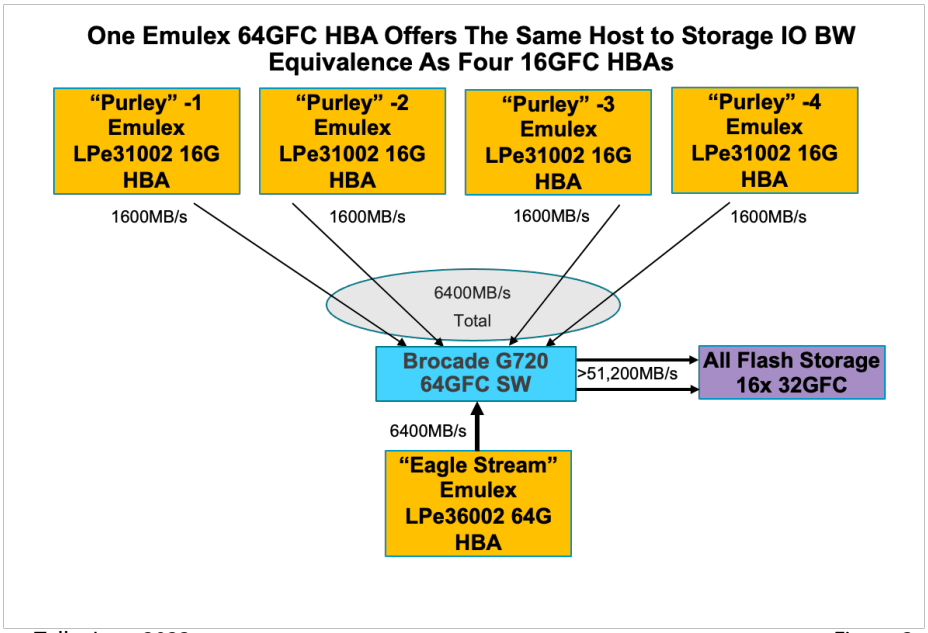

Source: Tolly, June 2023 Figure 3

Eagle Stream Server & Emulex LPe36002 64G HBA Oracle Database 19c

I/O Throughput & Transaction Improvement: Eagle Stream with 64GFC vs 16/32GFC

(as reported by HammerDB TPROC-H and Linux iostat)

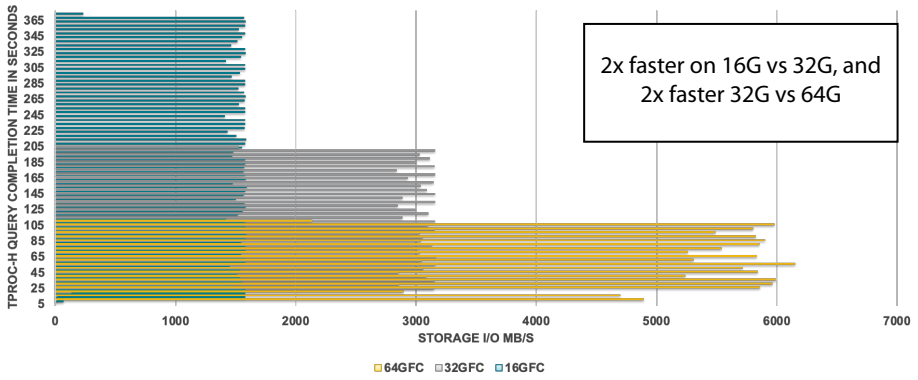

Source: Tolly, June 2023

Figure 4

Eagle Stream Server & 2-Port Emulex LPe36002 64G PCIe 4 HBA Performance Improvements

Eagle Stream (PCIe 5) vs Purley (PCIe 3) Servers

(as reported by HammerDB TPROC-H and Linux iostat)

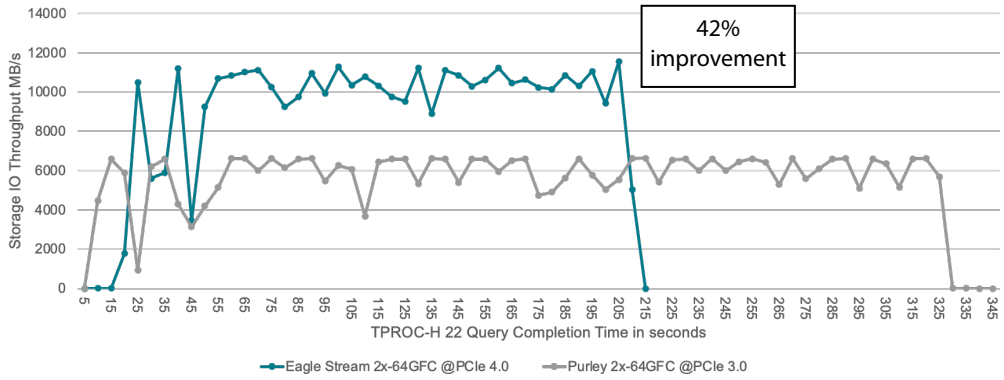

Source: Tolly, June 2023

Figure 5

64GFC vs 16/32GFC TPROC-H Results

This Oracle database analytics test used the TPROC-H database workload that measures how long it takes to run a series of 22 different types of decision support queries. This type of workload is "read only" with no database updates taking place.

This test was run three times with the only variable being the link speed between the server's FC HBA and the switch.

Figure 4 summarizes all three tests using two metrics: storage I/O throughput and query execution time as reported by the HammerDB database benchmark. What is important to note are the relative results across the three scenarios. The 16GFC HBA is clearly a bottleneck (brown bars) taking the longest to complete and delivering the lowest throughput.

Performance is improved, roughly by 2x, when the HBA is configured for 32GFC (gray bars) but, as will be seen, 32GFC still presented a transaction bottleneck.

When run using the 64GFC the database storage IO throughput is the highest and the query execution time is the shortest. Again, performance is

improved roughly by a factor of two over the 32GFC results.

64GFC Dual-Port HBA Performance

The Emulex LPe36002 64GFC HBA is a PCIe 4 interface card and is the recommended HBA for the Eagle Stream server. The card's total performance capacity is restricted by the bandwidth limitations of older generation servers that utilize PCIe 3.

As in the prior test, the TPROC-H database analytics benchmark was run on an Oracle 19c database multiple times using the same card but in servers that implement two different PCIe generation architectures.

Figure 5, on the previous page, illustrates the how the same dual-port 64GFC HBA delivers dramatically higher throughput and shorter database query times when deployed in a current generation server that implements PCIe 5 bus architecture.

Taking the same dual-port 64GFC HBA and deploying it in a PCIe 5 Eagle Stream server improved transaction time by 42% simply by removing the limitations imposed by the maximum bandwidth of the Purley PCIe 3 bus.

Figure 6, below, provides the consolidated test bed for the 16/32/64GFC tests.

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 1-10.

Server systems were all VMware ESXi 8 hosts running ESXi-8.0U1-21495797 (8U2 GA). Storage volumes mapped to each VM were configured as thick provisioned, eagerly zeroed. PVSCSI controller was used.

Each VM was assigned 100GB of memory and 40 vCPUs. Each VM was running RHEL 8.8

Details of the HammerDB tests, virtualization tests, and 16/32/64GFC tests are found in the test results section above for each test.

Consolidated Test Bed Topology: 64/32/16GFC & PCIe 5/PCIe 3 Tests

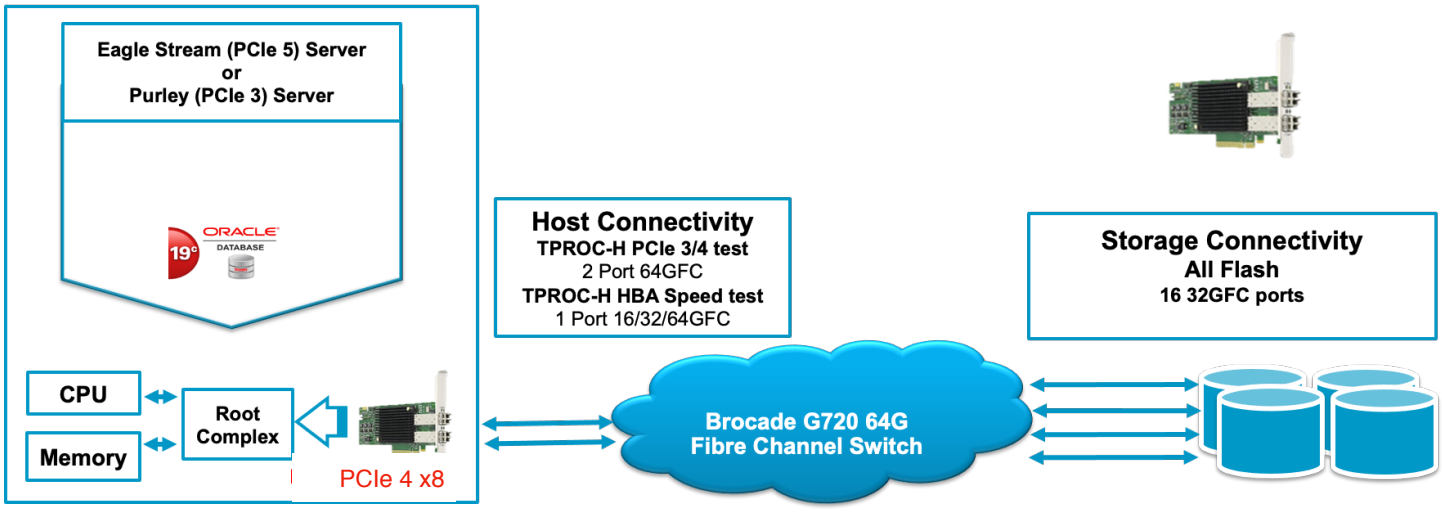

Source: Tolly, June 2023

Figure 6

Test Configuration Summary - 1 of 2

64G HBA Under Test

Vendor	Product Name	Bus Architecture	Firmware	Driver
Broadcom	Emulex LPe36002	PCIe 4	14.2.455.15	14.2.560.8

Table 1

Eagle Stream Server Configuration

Vendor/System	Intel Eagle Stream System
CPU	2 socket Intel(R) Xeon(R) Platinum 8468 @ 2.1 GHz
Number of CPUs	96
Memory (RAM)	512 GB
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-H settings	Degree of parallelism = 32 Scale factor = 30 Virtual users = 1 Ramp-up time: 2 minutes Run time: 5 minutes

Table 3

Oracle Database Configuration

Database	Oracle Database 19c (19.3)
Storage	Oracle Grid 19c, ASM disk group with external redundancy, 1 namespace for data
Dataset Size	40GB
Database Settings	SGA = 12000 MB PGA = 4000 MB Block size = 8 KB

Table 4

Storage Configuration

Data Store Description	NVMe-based, all-flash array
Ports	16 x 32G FC
Volumes	2 x NVMe: 200 GB and 1 TB
Performance Policy	High
Namespace/LUN	8 x 32G Target ports per Namespace
Network Fabric	Brocade G720 64G FC Switch v9.1.1

Table 5

Source: Tolly, June 2023

Test Configuration Summary - 2 of 2

16G HBA Under Test

Vendor	Product Name	Bus Architecture	Firmware	Driver
Broadcom	Emulex LPe31002	PCIe 3	14.2.455.11	14.2.560.8

Table 6

Purley Class Server Configuration	Host 1
CPU	2 socket Intel(R) Xeon(R) Gold 6146 @ 3.2GHz
Number of CPUs	24
Memory (RAM)	128 GB

Table 7

Purley Class Server Configuration	Host 2
CPU	2 socket Intel(R) Xeon(R) Platinum 8176 @ 2.10GHz
Number of CPUs	56
Memory (RAM)	128 GB

Table 8

Purley Class Server Configuration	Host 3
CPU	2 socket Intel(R) Xeon(R) Platinum 8176 @ 2.10GHz
Number of CPUs	56
Memory (RAM)	128 GB

Table 9

Purley Class Server Configuration	Host 4
CPU	2 socket Intel(R) Xeon(R) Gold 6148 @ 2.40GHz
Number of CPUs	40
Memory (RAM)	128 GB

Table 10

Source: Tolly, June 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.

Terms of Usage

223133 czse-9 -wt 2023-09-27- VerH