

Accelerating Oracle Data Warehousing with Emulex and Brocade Gen 6 Fibre Channel

32GFC provides performance boost for Oracle data warehousing workloads.



Executive Summary

Large enterprises choose Fibre Channel storage networking because of its proven performance, reliability and scalability for mission-critical workloads. Fibre Channel is seeing renewed interest from customers who have deployed other storage technologies. This comes as no surprise since Fibre Channel was purpose-built for storage – it delivers consistently high performance, does not drop packets, and has deep ecosystem support from all of the major operating system (OS) and hypervisor vendors.

Enterprise datacenters are increasingly turning to solid-state (flash) storage to improve storage performance, and they most frequently select Fibre Channel as the host interface to meet their flash storage requirements. In many cases, however, deploying solid-state storage moves the performance bottleneck to the storage network. This situation often drives users to push for higher performance from their storage network infrastructure.

The sixth generation of Fibre Channel is aimed at satisfying the needs of growing deployments of flash storage, hyper-scale virtualization, and new high-speed data center architectures such as NVMe.

Emulex commissioned Demartek to evaluate three generations of Emulex Fibre Channel adapters in an Oracle database data warehousing environment. We found that the new Emulex Gen 6 (32GFC) HBAs provided the performance needed to alleviate network bottlenecks that can be caused by all-flash arrays and dramatically improve data warehousing application

performance. New features boost reliability and deliver a suite of diagnostic, troubleshooting and deployment features.

Key Findings

- > Gen 6 products from Brocade and Emulex provide the powerful combination that can be used today to accelerate applications even with existing 8GFC and 16GFC storage arrays.
- > The Emulex LPe32002 (32GFC) HBA reduced Oracle 12c data warehouse workload query time by almost half compared to 16GFC and reduced it to a quarter compared to 8GFC connected servers.
- > Reduced latency for both target and initiator by half when compared to the LPe16002 (Gen5) HBA.
- > Advanced features developed in conjunction with Brocade accelerate deployment and simplify support that integrates with Brocade Fabric Vision™.
- > Easy-to-upgrade from previous generations of Fibre Channel (4GFC/8GFC) with forward/backward compatibility – using the same connectors, management, and cabling.

Emulex Gen 6 HBAs

Emulex Gen 6 HBAs deliver 2x greater bandwidth than the previous generation – 12,800MBps (2 ports, 32G, full duplex) – as well as less than half the latency and over 1.6 million IOPS on a single port [the highest FC per port performance seen thus far]. Improvements in the Emulex Gen 6 Fibre Channel adapter technology include operating at faster clock rates and an increased number of hardware offloads that reduce the number of firmware “touch points” required for Fibre Channel connection initialization functions.

One of the design goals of the new Emulex Gen 6 products was to lower latency in order to complement the growing all-flash storage array market. We visited Emulex labs and validated on specialized laboratory equipment that both the initiator and target latency was cut in half when compared to the previous generation products.

LPe32000-series is available in single and dual-port models with 32GFC optics. The Emulex LPe31000-series is available in single, dual and quad-port models with 16GFC optics. The single and dual-port are 32GFC-ready. Optics can be upgraded to 32GFC, enabling data centers to scale-up when needed. The quad-port model is a low-profile board with fixed optics. In order to accommodate a quad-port low-profile form-factor, the transceivers are not upgradeable.

NVMe over Fabrics Ready

NVM Express (NVMe) is a relatively new protocol for solid-state storage devices built with non-volatile memory technologies. NVMe provides substantially lower latency for storage I/O operations and significantly higher IOPS per device. NVMe will scale-up the number of devices it can address by adopting “NVMe over Fabrics” technology. Fibre Channel is one of the fabric technologies that will be supported by NVMe over Fabrics, and the Emulex Gen 6 Fibre Channel adapters are “NVMe over Fabrics Ready.”


Emulex Gen 6 Fibre Channel Adapter Models

Model	Ports	32GFC	16GFC	8GFC	4GFC
LPe32000-M2	1	✓	✓	✓	
LPe32002-M2	2	✓	✓	✓	
LPe31000-M6	1	*	✓	✓	✓
LPe31002-M6	2	*	✓	✓	✓
LPe31004-M6	4		✓	✓	✓

* Upgradeable to 32GFC with Emulex-certified transceiver kit

Emulex Dynamic Multi-Core Architecture

Emulex's unique Dynamic Multi-Core Architecture provides performance where it's needed. The architecture delivers full IOPS performance- 1.6 million IOPS - to a single-port, which is 2x more IOPS than other HBA designs can deliver. The advantages of this design are critical when using dual-port HBAs in an active-standby configuration. In fact, an estimated 80% of HBAs sold are dual-port and configured for active-standby fail-over mode.



ExpressLane™

In high-density virtual environments with mixed storage, scaling to meet business needs can be complex and often results in performance reduction. Emulex ExpressLane provides Quality of Service (QoS) and application performance between servers and across the fabric by tackling congestion in storage environments. Emulex ExpressLane assigns traffic priorities on the host application's storage path. This priority allows Fibre Channel frames to get to targets quickly, much like a highway express lane, allowing one LUN to be given higher priority than the other LUNs on

the same connection, improving performance and lowering latency for the prioritized traffic.

Secure Firmware Update

Secure Firmware Updates is a new feature whereby the Fibre Channel adapter firmware is digitally signed by Emulex. The signature is verified at the time of a firmware update, only allowing authentic firmware images to be loaded on the adapter. This eliminates the possibility of loading firmware that has been tampered with, making security attacks more difficult for hackers.

Forward Error Correction (FEC)

Forward Error Correction (FEC) is a Gen 6 Fibre Channel standard feature for HBAs and switches that provides enhanced data reliability and performance by automatically detecting and recovering from bit errors. It is especially useful in diverse and complex user environments such as blade system implementations. FEC is a digital signal processing technique that introduces redundant data, called an error correcting code, prior to data transmission or storage. FEC then provides the receiver with the ability to correct errors without a reverse channel to request the retransmission of data, thereby improving performance.

ClearLink™ Diagnostics Integration

Emulex integrates Brocade's ClearLink diagnostic capabilities for high performing SANs that demand nonstop availability for mission critical applications.

ClearLink diagnostics allow users to automate a battery of tests to measure and validate latency and distance across the switch links. These tests allow the IT administrator to detect physical cable issues, find faulty ports, estimate cable length and perform data integrity checks quickly and efficiently. Management of ClearLink is seamlessly integrated into the Emulex OneCommand Manager application.

Brocade Gen 6 Fibre Channel Switches

The LPe32000/LPe31000-series of adapters support Brocade I/O Insight for Gen 6 FC which proactively and non-intrusively monitors device and application-level I/O to gain insights into performance and availability, ensuring predictable performance and operational stability. SAN administrators have traditionally had poor visibility into the performance of storage I/O workloads between host and storage devices, and this inability to locate IOPS and latency issues can create storage related performance problems. The Brocade I/O Insight feature solves this problem by obtaining I/O latency and IOPS metrics for a storage device to diagnose I/O operational issues, including the ability to:

- > Identify I/Os that deviate from expected behaviors to facilitate fault isolation
- > Locate source of resource contention, congestion and errant storage devices with proactive I/O monitoring of the following:
 - > Total I/Os
 - > First response max/avg
 - > I/O latency (ECT) max/avg
 - > Outstanding I/O max/avg

The following advanced features and functionality are supported by Emulex Gen 6 HBAs, when connected to a Brocade Gen 5 or Gen 6 switch, providing a complete end-to-end solution:

- > ClearLink™ (D_Port) – automated end-to-end signal integrity checks help identify any cabling, optics or port issues in minutes versus hours.
- > Link Cable Beaconsing – locates a connection on either the Brocade switch port or the Emulex HBA port by making the LED port glow for easy end-to-end identification.
- > Host Name Registration – automated capture of information from Emulex HBA ports eliminates the need to manually associate worldwide port names with servers.

- > Read Diagnostic Parameters – Brocade switches and Emulex HBA ports will self-report diagnostic information, including port speed, link errors, and SFP information (temperature, Tx and Rx power, etc.).

Oracle Database

Oracle Database 12c delivers performance, scalability, high availability, data optimization, data security and ease of management to support the most demanding OLTP, Data Warehousing and Big Data requirements. Oracle Database 12c offers a variety of features including consolidating databases and enabling Database as a Service with Oracle Multitenant; real time analytics with database in-memory optimizations; automatically optimizing data storage and compression according to usage patterns; providing continuous access with Oracle's Maximum Availability Architecture; securing enterprise data with a comprehensive defense-in-depth strategy; simplifying the analysis and integration of Big Data and efficient database management using Enterprise Manager Cloud Control.

Red Enterprise Linux (RHEL)

Red Hat Enterprise Linux 6.6 has been tuned to optimize performance, system administration and virtualization. With support for higher processor counts and memory limits as well as kernel optimizations that allow for more efficient CPU utilization on large NUMA systems, Red Hat Enterprise Linux 6.6 fully accommodates dense single-server workloads. Other system performance enhancements include support for additional 40 GbE network adapters, reductions in network latency and jitter, and support for high performance, low latency applications.

Test Results

The primary workload used for these tests was a data warehousing application workload running on Oracle 12c. This read-intensive data warehousing workload is also known as a decision support application because it gives answers to critical business questions. It consists of a fixed set of 22 queries of relatively high complexity of a large database that examines large volumes of data.

Because this is a fixed set of work, improvements in infrastructure such as implementing faster Fibre Channel host bus adapters (HBAs) will result in completion of the work in a smaller amount of time.

We compared the performance of this data warehousing workload with three Emulex adapters.

- > Emulex LPe12002 – 8GFC
- > Emulex LPe16002 – 16GFC
- > Emulex LPe32002 – 32GFC

The same operating system (RedHat Enterprise Linux 6.6), the same database software (Oracle 12c) and the same application workload settings were used for all three tests.

Oracle Settings


The Oracle database used a 16GB cache for its data, with a cache-to-data ratio of 1:12.5. A “round robin” setting was used for the multi-path I/O for the eight paths to the all-flash storage array.

Throughput

When using real database workloads, the I/O rate varies as the workload progresses because the application is not only issuing I/O requests but is also consuming varying amounts of host CPU and memory resources. In

our tests, we achieved full line rate for each of the adapters, but the amount of time to complete the workload was much faster for the LPe32002 (32GFC) adapter than the other adapters, as shown in the chart.

HBA Throughput - Data Warehousing Workload Oracle 12c on RHEL 6.6




Query Response Time

In the following chart, the time to complete each of the individual 22 queries is shown for each of the three adapters. Using the LPe12002 (8GFC) completion time as a baseline, we observed that the LPe32002 (32GFC) adapter achieved a 72% reduction in the time to

complete the workload, from approximately 43 minutes to approximately 12 minutes. We believe that this same rate of reduction in time to complete the work would scale to much larger databases.

Completion Time of Decision Support Queries Oracle 12c on RHEL 6.6



Storage Array Speeds

These tests show that even with all-flash arrays that have 8GFC target ports, substantial improvements in performance can be achieved by upgrading the host adapters and switches to the 32GFC technology without changes on the target storage system.


Demartek has tested 16GFC all-flash arrays and found that in addition to higher available bandwidth, these arrays also provide lower latency. Even greater performance will be seen with the introduction of 32GFC arrays.

A Brief Commentary on Response Time

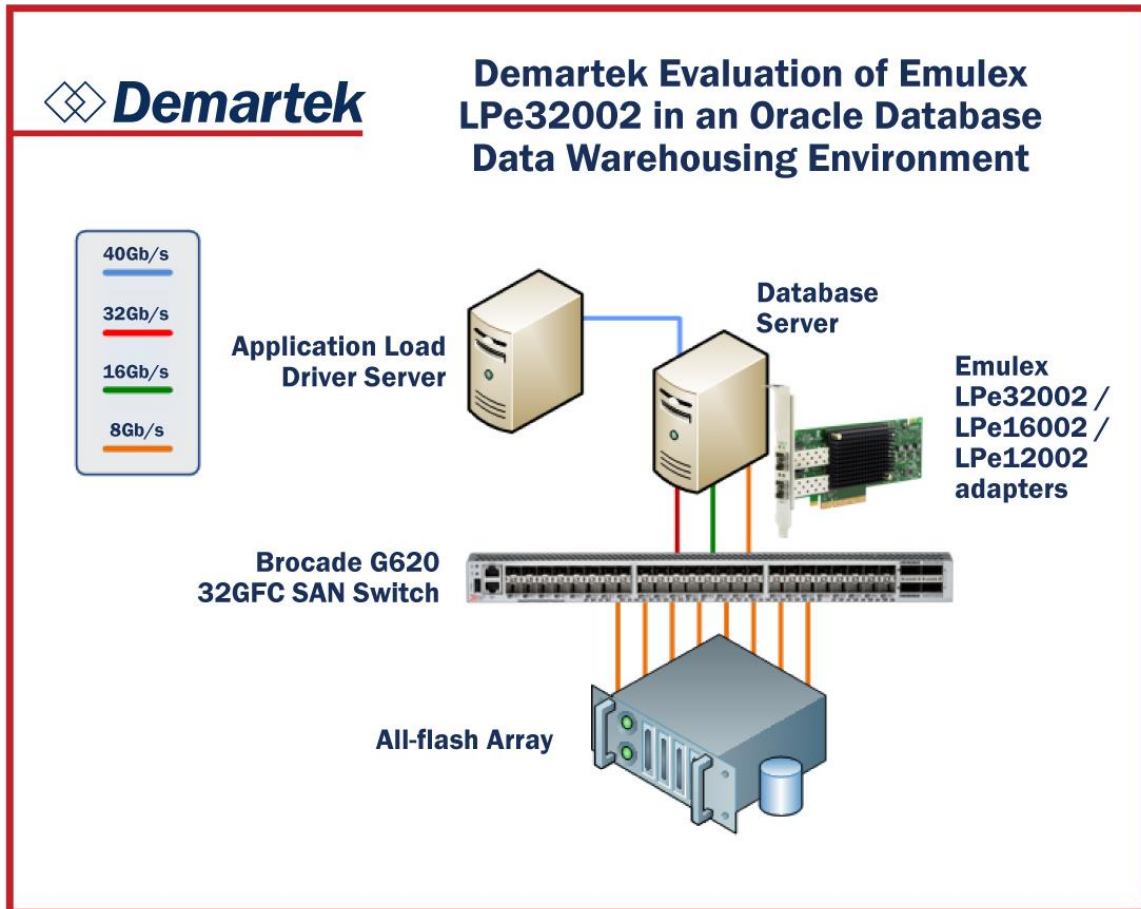
In our testing of all-flash arrays, we have found that application response time, or latency, is heavily dependent on the application workload. At one end of the spectrum, some workloads such as online transaction processing (OLTP) have very low latencies because they request small amounts of data. On the other hand, large data warehousing applications have much higher latencies because of the large volume of data that is usually requested. Other applications typically fall in between these two extremes.

In a separate test of the Emulex adapters, we observed substantial reductions in response time as measured by the host server for the various application workloads, based on their block sizes. The Emulex LPe32002 (32GFC) adapter provided approximately 70% lower application response time than the 8GFC adapter across four different types of workloads. The 32GFC adapter achieved approximately 50% lower response time for three of the workloads and 40% lower response times for the OLTP workload compared to the 16GFC adapter.

Average I/O Response Time by Blocksize (lower is better)



Test Environment



Servers (application and database servers)

- > 2x Intel Xeon E5-2643v3, 3.4GHz, 12 total cores, 24 total threads
- > 256 GB RAM
- > 40 Gb Ethernet NIC (Emulex OCe14401)
- > RedHat Enterprise Linux 6.6
- > Oracle 12c (database server only)

Fibre Channel Switch

- > Brocade G620 32GFC Switch

Storage System

- > All-flash array
- > 8x 8GFC target ports

Fibre Channel Adapters

- > Emulex LPe12002 (8GFC), firmware 2.02.A0
- > Emulex LPe16002 (16GFC), firmware 11.0.235.14
- > Emulex LPe32002 (32GFC), firmware 11.0.235.14

Summary and Conclusion

The Gen 6 Fibre Channel products from Emulex and Brocade provide new performance and feature capabilities ideally suited to connecting to high-speed flash-based storage systems. We found the following benefits and advantages to deploying these products together:

- > Accelerates applications even with existing 8GFC and 16GFC storage arrays
 - > Up to 4x faster for data warehousing workloads
 - > Cuts down data warehousing query time by 72%, from 43 minutes to 12 minutes in one example, enabling faster decision making and offering substantial business value.
- > Unique features such as Emulex ExpressLane more effectively use high performance tiered storage assets
- > Advanced features developed in conjunction with Brocade accelerate deployment and simplify support that integrates with Brocade Fabric Vision
- > Easy-to-upgrade from previous generations of Fibre Channel (4GFC/8GFC) with forward/backward compatibility – using the same connectors, management, and cabling

Appendix – Fibre Channel Technology Overview



History

Fibre Channel is an industry standard storage solution. Development began in 1988, receiving ANSI standard approval in 1994. Fibre Channel technology was introduced to the market in 1997 at 1 gigabit per second (Gbps) and its single-lane speed has doubled every three to five years. With Gen 6 Fibre Channel, there is a single-lane (32GFC) and quad-lane (128GFC) specification.

Backward Compatibility

Fibre Channel technology is backward compatible with the two previous generations. This means that Gen 6 (32GFC) Fibre Channel switches, host bus adapters (HBAs) and optics (transceivers) are backward compatible with 16GFC and 8GFC equipment. This provides a smooth upgrade path for enterprises.

Fibre Channel Speed History
Gbps



Fibre-Optic Cable Types and Lengths

Demartek recommends that customers purchase OM3 or OM4 multi-mode fiber-optic cables for Fibre Channel deployments within datacenters. The orange-colored OM1 and OM2 fiber-optic cables are considered obsolete for new deployments.

Fiber Optic Cable Lengths by Speed and Distance

	OM1	OM2	OM3	OM4
Cable Color	Orange	Orange	Aqua	Aqua
4GFC	70m	150m	380m	400m
8GFC	21m	50m	150m	190m
16GFC	15m	35m	100m	125m
32GFC	—	20m	70m	100m

Accelerating Oracle Data Warehousing with Emulex and Brocade Gen 6 FC

The most current version of this report is available at

http://www.demartek.com/Demartek_Emulex_32GFC_Accelerating_Oracle_performance_on_RedHat_Evaluation_2016-07.html on the Demartek website.

Broadcom®, the pulse logo, Emulex and LightPulse are among the trademarks of Broadcom Limited and/or its affiliates in the United States, certain other countries and/or the EU.

Brocade is registered trademark of Brocade Communications Systems, Inc.

Demartek is a registered trademark of Demartek, LLC.

All other trademarks are the property of their respective owners.