EMULEX®

White Paper

8Gb/s HBA
Performance Advantages for Oracle Database

AT A GLANCE

This benchmark study demonstrates that even for smaller OLTP databases, the Emulex LPe12000 8Gb/s HBA provides increased performance relative to the LPe11000 4Gb/s HBA in both I/O transactions per second and database response time. The results also showed that the LPe12000 can nearly double the peak performance with Oracle data warehousing queries.

PRODUCTS

Emulex LightPulse® LPe12000 8Gb/s Fibre Channel host bus adapters (HBAs)

APPLICATIONS

- Oracle Enterprise Linux 5
- Oracle 11g Database
- Oracle Data Warehousing

Introduction

Emulex LPe12000 LightPulse® 8Gb/s HBAs double the link speed of previous-generation 4Gb/s HBAs. The LPe12000 also provides better performance with low I/O rates, and improves latency and efficiency for every I/O.

In most cases, however, the maximum performance capability of the LPe12000 far exceeds the actual I/O transaction rates of Oracle database servers. This leads to a number of questions with respect to the performance of 8Gb/s HBAs with Oracle Database servers:

1. **Will an LPe12000 8Gb/s HBA provide a performance advantage with Oracle databases relative to 4Gb/s HBAs?**
2. **Will an 8Gb/s HBA increase database performance with a storage array that only supports 4Gb/s port connections?**
3. **Will the increased link speed be beneficial for data warehousing applications?**

This white paper reports the results of application benchmark tests designed to answer these questions.

8Gb/s EXCELLENCE IN THE ENTERPRISE

FCoE EXCELLENCE IN THE ENTERPRISE

ORACLE®

Oracle Test Environment

Most HBA performance reports on Oracle databases have concentrated on the extrapolation of raw performance data as measured in the perfect world of IOmeter test tools and RAM disk storage devices. The assumption has been made that if an HBA has better maximum IOPS (I/Os per second) performance with sequential data, then the HBA will also have better performance with an Oracle database that uses the same data block size.

While this assumption might apply to the maximum performance capability of an HBA, it does not directly inform database administrators how an HBA will actually perform with a typical database and storage array. The following tests were intended to provide information that directly relates to real-world Oracle database environments.

Test scenarios

The benchmark of Oracle database performance was done in two stages:

- **Calibration**—The Oracle Orion I/O calibration tool was used to establish a baseline and correctly size the storage array configuration. This was done using the "simple" test profile recommended in the Orion documentation that measures random I/O over a range of load levels. The data warehousing script was also used to validate a high-throughput configuration for the storage array.
- **Performance**—The SwingBench load generator was used to profile an order entry workload and a data warehousing workload using the Oracle 11g database.

Hardware

The same hardware was used for both test scenarios to avoid introducing factors that could alter the performance outcomes.

- **Server**—IBM x3650 server that was configured with two quad-core 2.66Ghz Intel processors and 8GB of system memory.
- **Operating system**—Oracle Enterprise Linux 5.
- **HBAs**—Emulex LPe12000 (8Gb/s) and LPe11000 (4Gb/s) with the latest available firmware and the 8.2.0.29 driver kit downloaded from the Emulex website. All HBA driver parameters were at default settings.
- **Switch**—HP 8/20q 8Gb/s Fibre Channel switch.
- **Storage array**—HP EVA4400 4Gb/s Fibre Channel storage array. Each 4Gb/s controller port was connected to the switch. Two disk groups were created, each with eight 15K RPM 140GB capacity disk drives. Two Vdisks were created, each used VRAID0 for a capacity of 800GB. Each Vdisk had read cache enabled and write-back cache requested. Each Vdisk was managed by a different storage controller. Each Vdisk was individually presented to either an LPe11000 or LPe12000 HBA.

Figure 1—Oracle HBA performance test configuration.

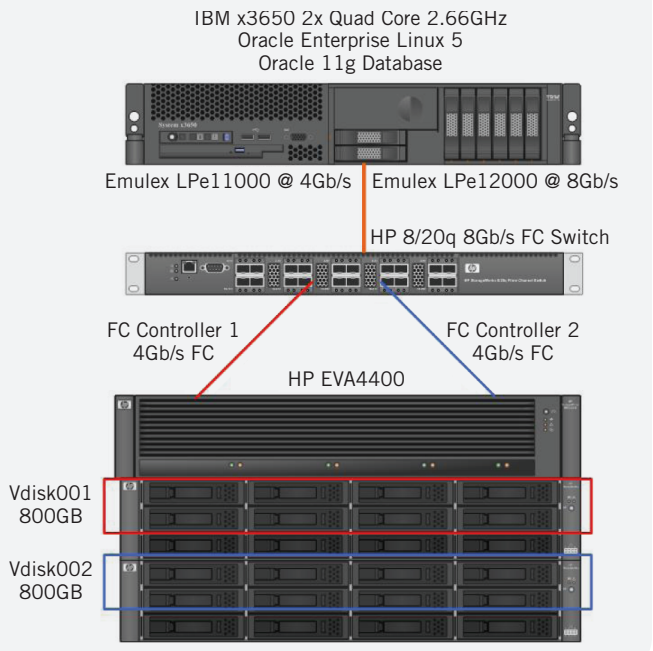

The Linux kernel was tuned according to the requirements stated by Oracle 11g installation guide. The following kernel tuning parameters were used by editing the /etc/sysctl.conf file:

```
fs.file-max = 512 * PROCESSES
kernel.shmall = 2097152
kernel.shmmax = 2147483648
kernel.shmmni = 4096
kernel.sem = 250 32000 100 128
net.ipv4.ip_local_port_range = 1024 65000
net.core.rmem_default = 4194304
net.core.rmem_max = 4194304
net.core.wmem_default = 262144
net.core.wmem_max = 262144
net.ipv4.tcp_wmem = 262144 262144 262144
net.ipv4.tcp_rmem = 4194304 4194304 4194304
```

Predicting Database I/O Performance

The Orion calibration tool was used to simulate Oracle I/O workloads that would predict real world performance. Oracle describes the tool as useful for understanding the performance capabilities of a storage system, either to uncover performance issues or to size a new database installation. Orion uses the Oracle database I/O libraries and can perform small I/Os to simulate an online transaction processing (OLTP) workload or large I/Os to simulate a data warehouse workload. The tool reduces benchmark complexity by eliminating the need to install and populate a live database environment.

One of the stated goals of this benchmark was to investigate if an LPe12000 8Gb/s HBA would improve the performance of a data warehouse application using an Oracle database. Data warehousing is characterized by very large data throughput due to a large sequential I/O profile, which should work well the increased bandwidth of the LPe12000.

The first requirement was to configure a storage array that could sustain throughput that exceeds a single 4Gb/s link. This was done by configuring two RAID groups distributed over two 4Gb/s ports of the EVA4400 controller.

The second requirement was to set up a test to simulate a data warehouse workload. The following script, which is similar to the data warehousing script in the Orion documentation, was used:

```
"./orion -run advanced -testname QLE2560_DW
-matrix point -num_large 4 -size_large 1024
-num_small 0 -num_disks 2 -type seq -num_
streamIO 16 -simulate concat -cache_size 2048
–verbose"
```

Figure 2 shows the simulated performance difference between a 4Gb/s LPe11000 and an 8Gb/s LPe12000 using the Orion test with the data warehousing I/O profile. As shown, the test confirmed that an 8Gb/s LPe12000 is capable of sustaining throughput levels that far exceed the 4Gb/s LPe11000. The performance was near the link rate limitation of 8Gb/s. Note that these tests were done using a 4Gb/s array as detailed previously.

Figure 2—The 8Gb/s LPe12000 data warehousing throughput performance is nearly twice that of the 4Gb/s LPe11000 HBA.

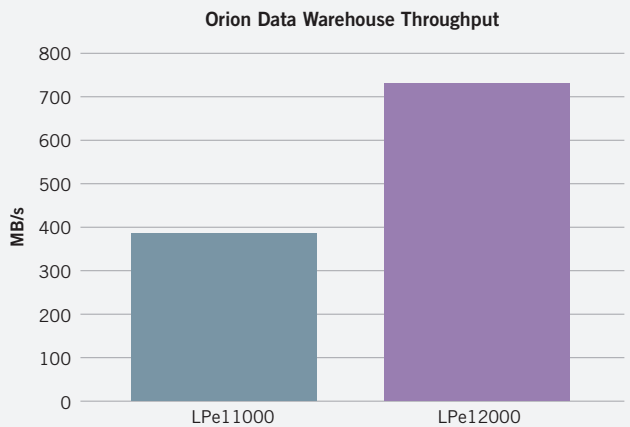

The next step was to do tests using the Orion OLTP I/O simulation, which is characterized by small random I/O. OLTP applications typically generate random reads and writes in data sizes that match the database block size. For Oracle, the default block size is 8KB.

Due to the extra processing required to move small data blocks, the performance metric changes from measuring link capacity in MB/s to the speed and response to which the I/O can be processed. The two typical measurements are IOPS (I/Os per second) and I/O latency. The "simple" script recommended in the Orion documentation was used to run small block random I/O. The storage configuration was the same as that used for the data warehousing test.

Figure 3 shows that the IOPS attained by the 8Gb/s HBA in the simple test exceeded that of the 4Gb/s HBA by nearly 15%. Figure 4 shows that I/O latency was also reduced by 16%.

Figure 3—Orion test results showing greater IOPS with the LPe12000.

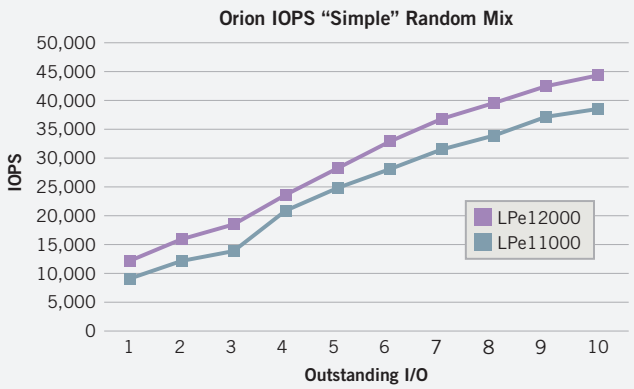

Figure 4—Orion test results showing lower latency with the LPe12000 over the entire load range.

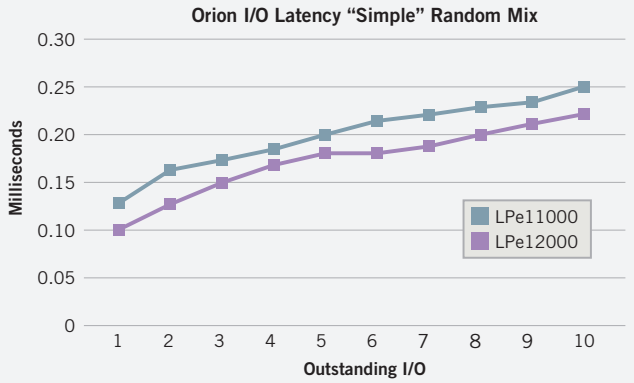

Oracle 11g Database Benchmark

With the performance capabilities of the I/O subsystem well understood from the results of the Orion simulation tool, the next step was to evaluate if the benefits of an 8Gb/s HBA would translate from simulation to the actual I/O dynamics of a live Oracle 11g database.

The Linux x86 version of Oracle 11g 1 (11.1) database was installed on the same server that was used for the Orion tests. SwingBench version 2.3 load generator and benchmark was used to stress test an Oracle 11g database. SwingBench includes two benchmarks - order entry and calling circle. Both require significantly large usage of CPU and memory resources from the server. Version 2.3 of SwingBench includes a DSS benchmark, which generates a significant I/O load that indicates if an 8Gb/s HBA would be beneficial for Oracle data warehousing.

In preparation for the DSS benchmark, the Oracle Database Configuration Assistant (./dbcl) was used to create a database across the same two storage LUNs used in the previous Orion benchmark. The Data Warehousing template suggested by Database Configuration Assistant was also used.

Figure 5—SwingBench user interface.

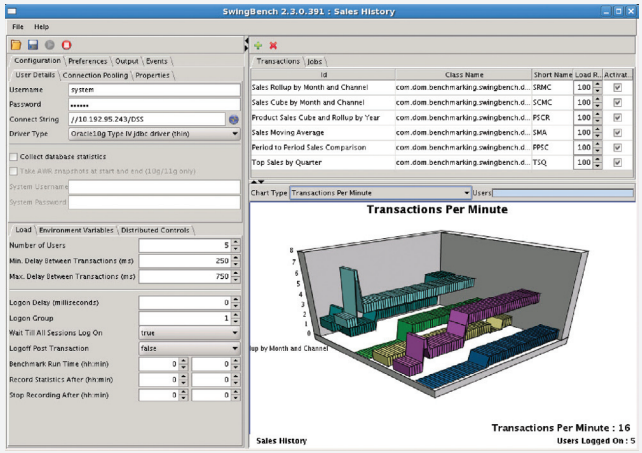

Figure 6—SwingBench test results showing higher peak throughput with the LPe12000.

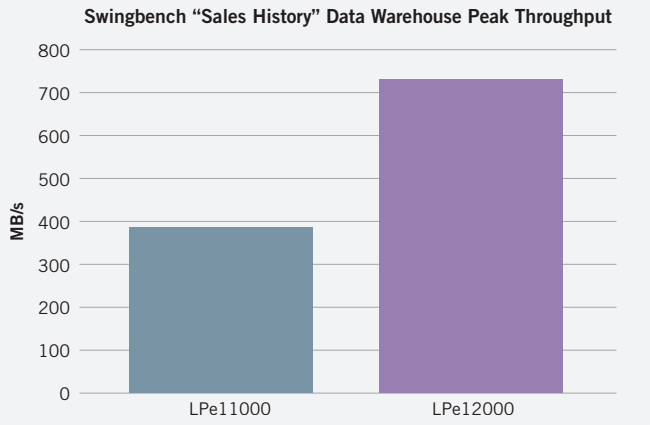

After installing the SwingBench code and data generator, data was loaded into the database using the SwingBench profile that is based on the "Sales History" sample schema that ships with Oracle 11g. The shconfig.xml profile was edited to add the necessary connection information.

A Fibre Channel analyzer was used to monitor the performance of the link between the HBA port and the switch port. As shown in Figure 6, the peak performance for data warehouse queries was indeed much higher for the 8Gb/s LPe12000 when compared to the 4Gb/s LPe11000. The peak throughput was nearly a match to the I/O simulation results that were measured with the Orion test.

Conclusion

The goal of this benchmark study was to investigate if the latest generation Emulex LPe12000 LightPulse® 8Gb/s HBAs would provide a performance benefit for Oracle database environments. The tests were done with a 4Gb/s storage array that would tend to minimize the improvement.

The results showed that even for smaller OLTP databases, the LPe12000 HBA provided increased performance relative to the LPe11000 4Gb/s HBA in both I/O transactions per second and database response time. The results also showed that the LPe12000 can nearly double the peak performance with Oracle data warehousing queries.

09-489 · 12/08

EMULEX®

www.emulex.com

World Headquarters 3333 Susan Street, Costa Mesa, CA 92626 +1 714 662 5600

Wokingham, UK +44 (0) 118 977 2929 | **Munich, Germany** +49 (0) 89 97007 177
Paris, France +33 (0) 158 580 022 | **Beijing, China** +86 10 68499547