BROCADE

DATA CENTER

Accelerating Oracle Databases with EMC CX4 Flash Drives and Brocade Connectivity

Solution Benefits:

- Ability to implement a Tier 0 storage layer capable of delivering very high I/O performance at very low latency.
- Ability to dramatically improve Oracle OLTP throughput and maintain very low response times.
- Seamless interoperability and Brocade advanced networking features via Brocade 8 Gbps FC Connectivity

ORACLE CERTIFIED PARTNER

EMC² where information lives®

CONTENTS

THE CHALLENGE

In the drive to accomplish more with less time and fewer resources, achieving better performance, scalability, and response times have continued to be the challenge for companies seeking to optimize information access and delivery to their corporate databases.

With the availability of the EMC® Enterprise Flash Drives (EFD) for the EMC CLARiiON mid-range storage array and Brocade® 8 Gigabits per second (Gbps) Storage Area Network (SAN) products and advanced network capabilities, customers can now implement solutions with shared access to Oracle® databases running on flash drives to deliver millisecond response times and up to 20x greater I/O operations Per Second (IOPS)..

THE SOLUTION

As flash drives are accessed using the standard Fibre Channel (FC) interfaces, the drives look like a standard FC disk, allowing Brocade advanced networking features to be used and enabling the administrator to use all the standard management and administration tools available for an FC storage array. In addition, flash drives benefit from the advanced capabilities that EMC provides, including local and remote replication, cache partitioning, and priority controls.

The enterprise flash solution allows database and application managers to improve the performances of their applications, which can directly increase business revenue and productivity. For these applications, two important benefits can be realized:

- A single flash drive can replace many short-stroked drives with its ability to provide very high transaction rates (IOPS). This reduces the total number of drives needed for the application, increases power savings with a reduced number of spinning disks, and eventually means reduced floor space in the data center.
- Flash drives provide very low latency, so applications for which a predicable low response time is critical and not all data can be kept at the host or in the cache will greatly benefit from these drives. Because no rotating media exists in EFDs, their transfer rate is extremely high and data is served much faster than the best response time that can be achieved with short-stroked hard disks.

ENTERPRISE FLASH DRIVE STRATEGIES

In summary, the EMC/Brocade solution for Oracle databases provides a database and application manager with the following enterprise EFD strategies:

- Place the entire database on EFDs when database performance is tied directly to business revenue.
- Place a portion of the database on EFDs as a storage Tier 0. Carefully identify the busiest portion that could benefit the most from EFDs based on joint EMC and Oracle guidelines. Examples are tablespaces, materialized views, indices, and temp.
- Place the active database partitions on EFDs with an Information Lifecycle Management (ILM) strategy in mind. As they get less active, migrate them to Tier 1, 2, or 3 storage, making room for the new active partitions on the EFDs.

In general, EFDs benefit random-read workloads the most and can easily be positioned for OnLine Transaction Processing (OLTP) applications. However, many Decision Support System (DSS) workloads also benefit, as they tend to become random by the time they reach the storage (due to multiple layers of striping and high query concurrency). Pure sequential workloads benefit as well, although to a lesser degree than random workloads.

NOTE: By default Oracle ASM stripes the data everywhere and makes any workload random.

By improving the performance of the busiest database components, it is likely that the performance of the rest of the database will also improve as the HDDs are less busy and can perform better.

Remember that disks are not always the bottleneck and that overall performance benefits are in proportion to the time Oracle was waiting to serve I/O. However, almost any workload can benefit from EFDs. Ensure that there is enough CPU power, connectivity, and I/O concurrency to benefit from the new Tier 0 capabilities.

SOLUTION COMPONENTS

Oracle Database

There are no simple, definitive rules to readily identify applications that best suit EFDs, but a few guidelines are available. It is very important to understand the load profile of an application before putting it on the EFDs, since most databases have different workload profiles during different times of the day. EFDs are suitable for highly read-intensive and extremely latency-sensitive applications and using these drives against the wrong target may not yield the desired benefit.

Understanding he following terminology can help you decide whether EFDs are suitable for particular workloads.

- **Write cache.** Most storage systems have large write-side cache and all write IOPS from a host are generally written to cache and incur no delay due to physical disk access. CLARiiON storage arrays have write caches sized that match the disk count supported by the controller and support enabling and disabling write cache at the LUN level if needed.
- **Read hit.** A read request from a database host can be served by storage system immediately if it already exists in storage cache because of a recent read or write or due to pre-fetch. A read serviced from the storage cache without causing disk access is called a “read hit.” If the requested data is not available in storage cache, the CLARiiON must retrieve it from disk, called a “read miss.”
- **Short-stroked drives.** Some extremely latency-sensitive applications use this technique on regular Fibre Channel drives to obtain low latencies. This is a technique in which data is laid out on many partially populated disks to reduce the spindle head movement and provide high IOPS at a very low latency.

Workloads with high CLARiiON cache read-hit rates are already serviced at memory access speed, and deploying them on flash drive technology may not return a significant benefit. Workloads with low CLARiiON cache read-hit rates that exhibit random I/O patterns, with small I/O requests of up to 16K, and that require high transaction throughput will benefit most from the low latency of EFDs.

Database and application managers can easily point to mission-critical applications that directly improve business revenue and productivity when business transaction throughput is increased, along with reduced service latencies. Cognizant of these applications, storage administrators often resort to short stroking more drives to ensure the highest possible I/O service level supported for them. EFDs can provide two very important benefits for such applications namely the ability to provide very high transaction rate (IOPS) and very low latency.

EMC CLARiiON CX4

The CLARiiON CX4 series with UltraFlex™ technology is based on a new, breakthrough architecture and extensive technological innovation, providing a mid-range storage solution that is highly scalable, meeting the price points of most mid-range customers. The CX4 is the fourth-generation CX series and continues EMC's commitment to maximizing customer investments in CLARiiON technology by ensuring that existing resources and capital assets are optimally utilized as customers adopt new technology.

The new CLARiiON CX4 systems shown in Figure 1 are the next generation in the CX series that extends EMC's leadership in the mid-range enterprise storage market. The CX4 delivers immediate support for the latest generation of disk drive technologies, such as EFDs, 4 Gbps FC drives for high performance, and SATA II for high capacity. CLARiiON CX4 is the first mid-range storage system that can support all of these latest generations of disk drive technologies. CLARiiON CX4 with the latest release of FLARE® (R28) has been optimized for maximum performance and tiered storage functional flexibility. A complete introduction to the CX4 series is beyond the scope of this paper, however Figure 1 highlights major features of the series. EFDs are supported on all four models of CLARiiON CX4 listed here.

CX4-960

- Up to 960 drives; 32 GB memory
- Standard 8 Fibre Channel/4 iSCSI
- Maximum 32 front-end Fibre Channel and/or iSCSI

CX4-480

- Up to 480 drives; 16 GB memory
- Standard 8 Fibre Channel/4 iSCSI
- Maximum 24 front-end Fibre Channel and/or iSCSI

CX4-240

- Up to 240 drives; 8 GB memory
- Standard 4 Fibre Channel/4 iSCSI
- Maximum 20 front-end Fibre Channel and/or iSCSI

CX4-120

- Up to 120 drives; 6 GB memory
- Standard 4 Fibre Channel/4 iSCSI
- Maximum 16 front-end Fibre Channel and/or iSCSI

Figure 1. EMC CLARiiON CX4 models

Enterprise Flash Drives

The enterprise-class EMC Flash drives supported by CLARiiON CX4 are constructed with nonvolatile semiconductor NAND Flash memory and are packaged in a standard 3.5-inch disk drive form factor used in existing CLARiiON disk drive array enclosures. These drives are especially well suited for latency-sensitive applications that require consistently low read/write response times. EFDs also benefit from the advanced capabilities that CLARiiON provides, including local and remote replication, Virtual LUN Migration, Navisphere® Quality of Service Management, and "five nines" (99.999 percent) availability.

CLARiiON storage arrays currently support EFDs in two sizes, 73 GB and 400 GB. Both these drives have slightly different application performance characteristics depending on workload nature. They have almost similar read characteristics but somewhat different write characteristics. The 400 GB drives are specifically optimized for space and provide very good $/GB ratio whereas the 73 GB drives are optimized for performance and would offer better $/IOPS. The customer can choose either drive depending on their space and performance requirements.

Brocade 5100 and 5300 Switches

Brocade has the largest installed base of networked storage infrastructure supporting Oracle applications running in thousands of Brocade production environments worldwide. Today, there are more than 13 million Brocade SAN ports in data centers, representing approximately 75 percent of the overall networked storage market.

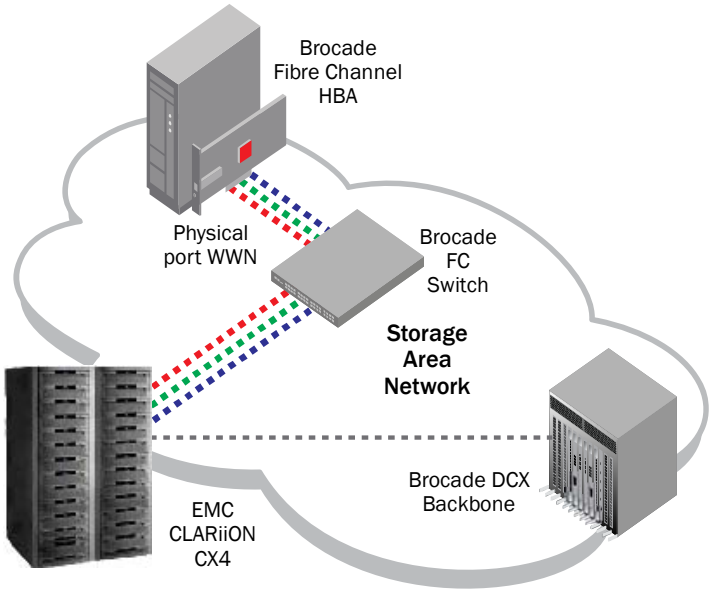

Figure 2. Brocade data center fabric with Adaptive Networking Quality of Service (QoS)

Both the Brocade 5100 and 5300 Switches support 1, 2, 4, and 8 Gbps technology in efficient 1U and 2U form factors respectively. The Brocade 5100 (shown in Figure 2) is available in of 24-, 32-, or 40-port configurations, and the Brocade 5300 in 48-, 64-, or 80-port configurations. The combination of port density, performance, and "pay-as-you-grow" scalability enables higher throughput and greater storage utilization, while reducing complexity for virtualized data centers. The evolutionary design makes it very efficient with regards to power consumption, cooling, and rack density—enabling medium- and large-scale server and storage consolidation for greater cost savings, better price/performance, and manageability. Compared to competitive SAN switches, Brocade provides greater power efficiency; more effective rack consolidation, and a significantly lower carbon footprint.

Brocade SAN switches feature a non-blocking architecture with ports concurrently active at 8 Gbps (full duplex) with no over-subscription. This level of performance and connectivity is ideal for expanding virtual server environments. In addition, enhanced Inter-Switch Link (ISL) Trunking can supply up to 64 Gbps of balanced data throughput in a single logical link.

Brocade 5100 and 5300 Switches offer Top Talkers (Advanced Performance Monitoring) and Adaptive Networking services, a suite of tools that includes Ingress Rate Limiting (IRL), Traffic Isolation (TI), and Quality of Service (QoS). These advanced capabilities help optimize fabric utilization and allocate ample bandwidth for mission-critical Oracle applications. Providing maximum flexibility, both products have Integrated Routing (IR) capabilities to connect switches across fabrics. Also, Virtual Fabrics enables the partitioning of a physical SAN into Logical Fabrics. This provides fabric isolation based on applications, business groups, customers, or traffic types without sacrificing performance, scalability, security, or reliability.

Brocade switches utilize cost-effective management solutions to optimize storage network resources, improve efficiency, and reduce Total Cost of Ownership (TCO). The Brocade EZSwitchSetup wizard is designed to simplify deployment. For multi-switch environments, Brocade Data Center Fabric Manager (DCFM™) Professional streamlines management and provides fabric-wide monitoring capabilities.

Brocade 815 and 825 Fibre Channel HBAs

The Brocade 815 (single-port) and Brocade 825 (dual-port) FC HBAs lay the foundation for extending fabric intelligence to servers, virtual machines, applications, and services—providing end-to-end storage network management. This approach enables tighter integration across the enterprise, including both physical and virtual networks.

By installing Brocade HBAs, organizations have a foundation to leverage N_Port ID Virtualization (NPIV) and maintain QoS from the virtual ports throughout the fabric. And to reduce administrative workload, they can automate the SAN boot process to quickly deploy diskless servers and operating systems.

SUMMARY

Incorporation of EFDs into the EMC CLARiiON CX4 series provides a new Tier 0 storage layer capable of delivering very high I/O performance at a very low latency, which can dramatically improve OLTP throughput and maintain very low response times. With comprehensive qualification and testing to ensure reliability and seamless interoperability, Tier 0 is supported by all the key CLARiiON software applications, such as data replication and remote protection.

Magnetic disk drive technology no longer defines the performance boundaries for mission-critical storage environments. The costly approach of spreading workloads over dozens or hundreds of underutilized disk drives is no longer necessary. EMC CLARiiON combines the performance and power efficiency of EFD technology with traditional disk drive technology in a single array managed with a single set of software tools to deliver advanced functionality, ultra-performance, and expanded storage tiering options.

LEARN MORE

Brocade partners with companies of all sizes to deliver innovative solutions that help organizations maximize the value of their most critical information. To learn more, visit www.brocade.com/alliance.

About Brocade

Brocade connects the world's most important information—delivering proven networking solutions for today's most data-intensive organizations. From the data center to high-performance Ethernet networks, Brocade is extending its near-15-year heritage as a leading partner for advanced storage and networking solutions.

The innovative Brocade approach to networking facilitates strategic partnerships by providing:

- Standards-based solutions that streamline qualification and help ensure interoperability
- Leading-edge technology that enables comprehensive, best-in-class solutions
- The best alternative choice in terms of overall business value

The world's largest enterprise networks, government entities, and global service providers rely on Brocade to maximize the business return on their data. It's no wonder 90 percent of the world's most critical business information flows through Brocade solutions. Quite simply, Brocade enables today's complex businesses to run. Where other vendors produce networking that's ordinary, Brocade is committed to delivering the extraordinary.

About EMC

EMC Corporation (NYSE: EMC) is the world's leading developer and provider of information infrastructure technology and solutions that enable organizations of all sizes to transform the way they compete and create value from their information. Information about EMC's products and services can be found at www.EMC.com.

About Oracle

Oracle (NASDAQ: ORCL) is the world's largest business software company. For more information about Oracle, please visit the Web site at http://www.oracle.com.