BOOSTING PERFORMANCE WITH ENTERPRISE FLASH DRIVES IN DELL/EMC CX4 SERIES STORAGE

By Bharath Vasudevan
Annette Cormier
Eric Cannell

Enterprise flash drives represent a huge leap forward over traditional mechanical drives. When deployed in Dell/EMC CX4 Series storage arrays, these drives can dramatically boost storage performance for targeted data center workloads while helping increase resource utilization, enhance energy efficiency, and reduce total cost of ownership.

Tremendous performance advances over the years have enabled nearly exponential growth in many areas of technology—processors, memory, I/O buses, and more. Until recently, however, performance improvements in storage have lagged considerably behind other technologies, often resulting in application performance bottlenecks. In particular, the rotational speeds of traditional mechanical hard disk drives (HDDs) have increased only fourfold over the past quarter century or so—from 3,600 rpm to 15,000 rpm. To compensate for the performance gap, administrators often resort to work-arounds such as short stroking and striping. While these measures help achieve the desired performance, the trade-off can be a significant underutilization of capacity. These work-arounds cannot completely overcome the rotational and seek delays inherent in traditional HDDs, which can reduce latency to an acceptable threshold.

Today, enterprise flash drives (EFDs), also known as solid-state drives (SSDs), represent a huge leap forward in storage technology. Because EFDs have no moving parts, they incur none of the rotational and seek delays inherent in traditional HDDs. As a result, EFDs offer tremendous enhancements in performance and response time relative to traditional HDDs. Also, because they require no specialized techniques such as short stroking to enhance performance, they help increase utilization and reduce the number of disks needed to meet performance and response-time requirements. Additionally, because solid-state EFDs are smaller and more energy-efficient than traditional mechanical HDDs, they help reduce data center energy and space requirements.

Because EFDs offer exceptional performance in configurations consisting of far fewer drives than traditional HDDs, EFDs in Dell/EMC CX4 Series storage area network (SAN) arrays allow organizations not only to meet application performance and response-time requirements but also to increase resource utilization and enhance energy efficiency—while helping to reduce total cost of ownership. EFDs are one of several disk drive options available in Dell/EMC CX4 Series storage arrays, which also include Fibre Channel and Serial ATA (SATA) drives.

All drives are designed to leverage the same management and replication features available in Dell/EMC CX4 Series arrays, such as virtual provisioning; EMC® MirrorView™, EMC SAN Copy™, and EMC SnapView™ software; EMC Navisphere® Management Suite software; and the EMC CLARiiON® splitter built into the EMC FLARE® code for integration into EMC RecoverPoint™ appliances. All drives are tested and validated to be used with EMC's leading-edge data protection software: Replication Manager, a host-based tool for synchronizing application and OS processes with snapshots and clones of valued data, and RecoverPoint,

an appliance that connects to Dell/EMC CX4 Series arrays to offload concurrent local and remote any-point-in-time recovery of Dell/EMC CX4 volumes. Dell/EMC CX4 Series arrays can also integrate with Dell OpenManage™ systems management software for seamless management of end-to-end storage and servers.

> **"EFDs in Dell/EMC CX4 Series SAN arrays allow organizations not only to meet application performance and response-time requirements but also to increase resource utilization and enhance energy efficiency—while helping to reduce total cost of ownership."**

IDENTIFYING APPROPRIATE WORKLOADS FOR EFDs

Although EFDs deliver enormous performance advantages, their cost and capacity limitations make them most effective when used as a targeted solution for particular load profiles—especially for workloads with extremely high performance and low latency requirements. Examples of this type of workload include online transaction processing (OLTP) database applications such as customer relationship management (CRM), enterprise resource planning (ERP), and order-entry applications; applications requiring extremely fast retrieval and storage of data such as currency exchange and electronic trading systems; and highly read-intensive workloads such as search engine databases.

For example, workload profiles that may benefit from EFD technology are characterized by low cache read hit rates with random I/O patterns, small I/O requests of up to 16 KB, and extremely high transaction throughput requirements. Typical usage scenarios include hot database tables that require read-response times of less than 1 ms and/or high throughput, database temp areas, and high-write-bandwidth applications such as database loads and system backups.

Many other usage scenarios are well suited to EFDs, such as small, highly active file systems that have response-time service-level agreement (SLA) requirements and metadata control areas in cluster file systems that require minimum read-response time. In general, any situation in which performance and response-time requirements would necessitate HDD work-arounds such as short stroking or striping—or which could not achieve the requisite performance and response time even by short stroking traditional HDDs—is one that might benefit from EFD technology.

Enterprise flash drives in Dell/EMC CX4 Series storage arrays offer exceptional performance and energy efficiency for targeted enterprise applications

Not all workload situations are equally suited to EFDs, however. Usage scenarios that may not be cost-efficient for EFDs include workloads with high cache read hit rates already serviced at memory access speed, workloads with large-block sequential reads such as online analytical processing (OLAP) applications, and workloads that are not constrained by rigorous performance and response-time requirements.

BENCHMARKING EFD PERFORMANCE

To quantify the performance of EFDs in Dell/EMC CX4 Series SAN arrays, and to help administrators identify optimum usage scenarios for EFDs, EMC engineers conducted a series of benchmark tests in October 2008 comparing the performance and response time of EFD-based storage arrays with traditional HDD arrays. In particular, EMC compared the performance of EFDs relative to Fibre Channel drives in three configurations: a comparison of 150 Fibre Channel drives with 6 EFDs for an extremely read-intensive workload, a comparison of 75 Fibre Channel drives with 6 EFDs using a read/write OLTP workload, and a direct comparison of 6 Fibre Channel drives with 6 EFDs using a read/write OLTP workload.

The tests were conducted using a Dell™ PowerEdge™ R900 server with four

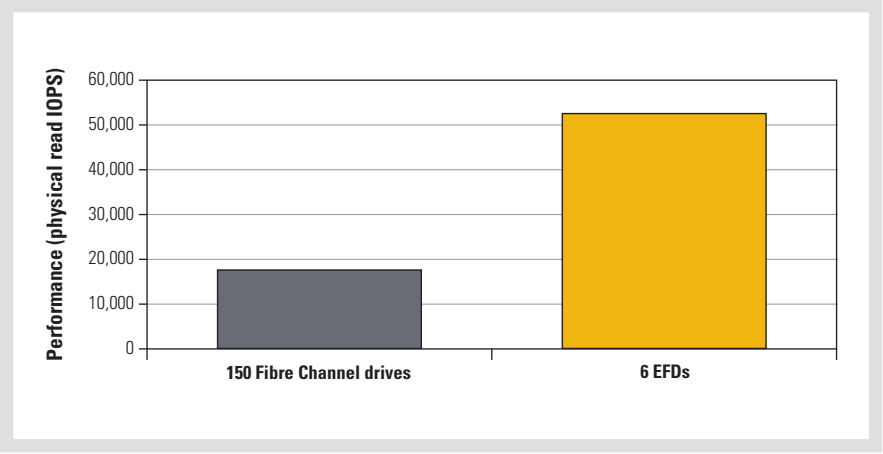

Figure 1. *Physical read I/Os per second for 150 Fibre Channel drives compared with 6 EFDs*

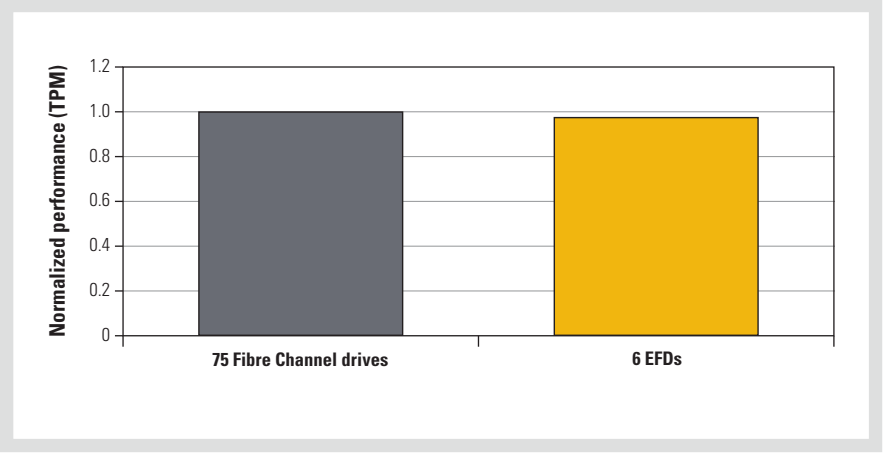

Figure 2. *Normalized transactions per minute for 75 Fibre Channel drives compared with 6 EFDs*

quad-core Intel® Xeon® processors, 32 GB of RAM, and two 146 GB, 15,000 rpm Serial Attached SCSI (SAS) drives supporting the Oracle® Enterprise Linux® 5.2 OS and Oracle database binaries. The database platform was Oracle Database 11.1.0.6 with Oracle Automatic Storage Management (ASM) for volume management and a 260 GB database with 600 warehouses and 100 concurrent users.

The storage configuration consisted of a Dell/EMC CX4-960 SAN array with the EMC FLARE 28 operating environment. The configuration utilized 300 GB, 15,000 rpm short-stroked Fibre Channel drives in a RAID-5 configuration and 73 GB EFDs in a RAID-5 configuration. Test loads were generated using a 64-bit Linux server, and the resulting system loads were measured in several ways: the storage tier load was measured by Navisphere Analyzer, the OS load was measured by vmstat and iostat, and the database usage was measured by Oracle Automatic Workload Repository (AWR) reports.

Read-intensive configuration: 150 Fibre Channel drives versus 6 EFDs

In the read-intensive configuration test, both drive configurations were subjected to a 99.9 percent read-intensive load, which would be typical of a workload such as heavy search indexing. The configuration of 150 short-stroked Fibre Channel drives subjected to this read-intensive load achieved a throughput of 17,459 physical read I/Os per second (IOPS) with an average latency of 4 ms. In contrast, the configuration of 6 EFDs subjected to the same load achieved a throughput of 53,055 physical read IOPS with an average latency of 1 ms.

In this read-intensive environment, EFDs contributed to an enormous enhancement in performance and latency—including over 3 times more read IOPS in total, and over 75 times more IOPS per drive (see Figure 1). Also, the EFD configuration delivered a 75 percent reduction in overall response time. Furthermore, these results were achieved using only 6 EFDs relative to 150 Fibre Channel drives—a 96 percent reduction in the number of drives required. These results illustrate that EFDs achieved more than 75 times more reads per drive than the Fibre Channel drives.

Read/write OLTP configuration: 75 Fibre Channel drives versus 6 EFDs

In the read/write OLTP configuration test, both drive configurations were subjected to a 60/40 read/write OLTP workload. To simulate a real-world environment, both drive configurations were also subjected to a background 40–45 percent load level with caches saturated. Redo logs were kept on the Fibre Channel drives.

In this environment, the configuration with 6 EFDs delivered almost the same transactions per minute (TPM) as the configuration with 75 Fibre Channel drives (see Figure 2). On a per-drive basis, the EFD configuration delivered 12.25 times more TPM per drive than the Fibre Channel configuration. This result was achieved using 92 percent fewer EFDs than Fibre Channel drives.

Direct OLTP comparison: 6 Fibre Channel drives versus 6 EFDs

To more directly measure the performance gains of EFDs relative to Fibre Channel drives, EMC also compared a configuration of 6 Fibre Channel drives directly against a configuration of 6 EFDs using an OLTP workload with a 60/40 read/write ratio. In this benchmark, the EFDs sustained an average of 19,000 TPM,

while the Fibre Channel drives sustained only 2,400 TPM (see Figure 3), delivering on average an eightfold enhancement over the Fibre Channel drives.[1]

BOOSTING PERFORMANCE COST-EFFICIENTLY

As shown by the benchmark results in this study, the EFDs in Dell/EMC CX4-960 Series SAN arrays can deliver dramatic performance gains relative to traditional HDDs in both read-intensive and high-transaction-rate OLTP environments. Because far fewer EFDs than traditional HDDs are required to achieve these results, EFD deployments in targeted applications can help increase utilization and energy efficiency while helping to reduce data center costs.

In particular, the significant reduction in drives that an EFD configuration enabled in this study indicates that EFDs can help reduce capital costs, energy requirements for power and cooling, and space requirements for data centers. And because EFDs are one of several disk drive options available in Dell/EMC CX4 Series storage arrays, EFD deployments can also benefit from the management and replication features available with Dell/EMC CX4 Series arrays, which contribute to reduced management costs as well.

> **"As shown by the benchmark results in this study, the EFDs in Dell/EMC CX4-960 Series SAN arrays can deliver dramatic performance gains relative to traditional HDDs in both read-intensive and high-transaction-rate OLTP environments."**

Bharath Vasudevan is the product manager for database programs at Dell. He previously led the solutions engineering team that delivered multiple commercial solutions including e-mail, unified communications, and flexible computing solutions. He has a master's degree in Electrical and Computer Engineering from Carnegie Mellon University.

Annette Cormier is a solutions marketing manager for Dell/EMC storage solutions. She has 20 years of experience in developing and bringing to market enterprise storage, network management, and security products for Dell, Hewlett-Packard, and SGI, and has previously been an IT SAS/Oracle database programmer at the Natural Resource Ecology Laboratory (NREL) at Colorado State and at Pacific Power and Light. Annette has a B.S. in Computer Science, Artificial Intelligence, from Colorado State University.

Eric Cannell is a product marketing manager for Dell/EMC storage systems. He has 17 years of experience developing enterprise technology products, before which he spent 6 years as a software engineer at the NASA Jet Propulsion Laboratory. Eric has B.S. degrees in Astronomy and Computer Science from the University of Illinois, an M.S. in Computer Science from the University of Southern California, and an M.B.A. from the Tuck School of Business at Dartmouth College.

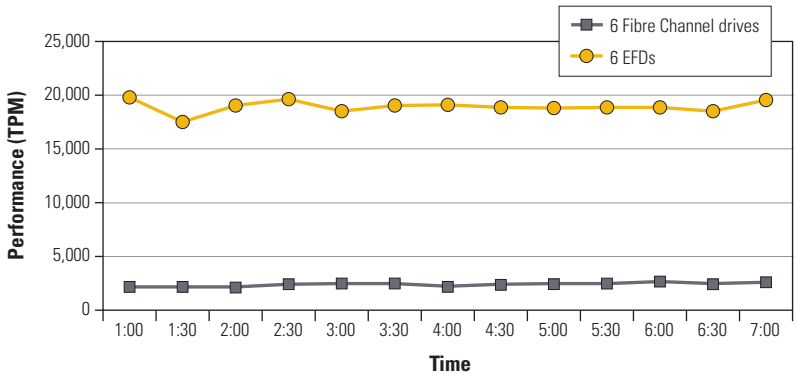

Figure 3. *Transactions per minute for 6 Fibre Channel drives compared with 6 EFDs*

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Dell/EMC storage:
DELL.COM/EMC

Register for a Dell/EMC CX4 event or Webcast:
www.dellenterprise.com/goto/emc/index.html

[1] For additional details and results from the tests described in this article, see "Leveraging Dell | EMC CX4 with Enterprise Flash Drives for Oracle Database Deployments," by Dell, April 2009, available at DELL.COM/EMC.