

Implementing EMC Symmetrix DMX-4 Flash Drives with Oracle Databases

Applied Technology



Abstract

This white paper examines the performance considerations of placing Oracle Databases on enterprise Flash drives versus conventional hard disk drives, and discusses the physical placement of database files.

April 2009

Copyright © 2008, 2009 EMC Corporation. All rights reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com

All other trademarks used herein are the property of their respective owners.

Part Number H5699.3

Table of Contents

Executive summary	4
Introduction	4
Audience	5
Technology overview	6
Symmetrix DMX-4 overview.....	6
Enterprise Flash drives	6
Which database workloads best fit enterprise Flash drives?	7
Random reads workloads	7
DSS and BI workloads	7
Writes workloads.....	7
Latency critical workloads	8
Summary of enterprise Flash drive strategies	8
Which database components to place on Flash drives?	9
Inspection of an Oracle AWR report	9
Load Profile	9
Top 5 Timed Events	10
Instance Efficiency	10
Tablespace I/O Stats and File I/O Stats	11
Use Case #1: Benefits of moving an entire OLTP database to EFDs	11
Part I – Database on one RAID 5 HDD vs. one RAID 5 EFD.....	11
Part II – Database on 80 RAID 1 HDDs vs. 12 RAID 5 EFDs	13
Use Case #2: A single tablespace migration to Flash	14
Flash drives and Information Lifecycle Management (ILM) strategy	16
Conclusion	17

Executive summary

EMC has enhanced the latest release of Symmetrix® Enginuity™ version 5773 to integrate enterprise-class Flash drives into the Symmetrix DMX-4 storage array. EMC® Symmetrix is the first and only enterprise array with support for this emerging generation of drive technology. With this capability, EMC creates a new “Tier 0” ultra-performance storage tier that removes the performance limitations of magnetic disk drives. By combining enterprise-class Flash drives optimized with EMC technology and advanced Symmetrix functionality, organizations now have new tiering options previously unavailable from any enterprise storage vendor.

Enterprise Flash drives (EFDs) dramatically increase performance for latency sensitive applications. Enterprise Flash drives, also known as solid state drives (SSD), contain no moving parts and appear as standard Fibre Channel drives to existing Symmetrix management tools, allowing administrators to manage Tier 0 without special processes or custom tools. Tier 0 Flash drives are ideally suited for applications with high transaction rates and those requiring the fastest possible retrieval and storage of data, such as currency exchange and electronic trading systems, or real-time data acquisition and processing. A Symmetrix DMX-4 with Flash drives can deliver millisecond application response times and up to 30 times more I/O operations per second (IOPS) than traditional Fibre Channel hard disk drives. Additionally, because there are no mechanical components, Flash drives consume significantly less energy per IOPS than traditional hard disk drives.

This white paper examines some of the use cases and best practices for using enterprise Flash drives with Oracle Database workloads. It shows that Flash drives deliver vastly increased performance to the database application when compared to traditional Fibre Channel drives, both in transaction rates per minute as well as transaction response time.

Introduction

Database performance has long been constrained by the I/O capability of hard disk drives (HDD) and the performance of the HDD has been limited by intrinsic mechanical delays of head seek and rotational latency. Flash drives, however, have no moving parts and therefore no seek or rotational latency delays, which dramatically improves their ability to sustain very high numbers of IOPS with very low overall response times.

Figure 1 shows the theoretical IOPS rates that can be sustained by traditional HDD based on average seek and latency times as compared to Flash drive technology. Over the past 25 years, the rotational speeds of HDDs have improved from 3,600 rpm to 15,000 rpm, yielding only four times the improvement in IOPS. Flash drive technology represents a significant leap in performance and can sustain over 30 times the IOPS of traditional HDD technology.

FLASH Drive vs. HDD IOPS




Figure 1. Relative IOPS of various drive technologies

The introduction of the industry’s first enterprise-class Flash drives in EMC Symmetrix DMX-4 disk arrays meets the growing demand for higher transaction rates and faster response times. Companies no longer have to purchase large numbers of the fastest Fibre Channel disk drives and only utilize a small portion of their capacity (known as short stroking) to satisfy the IOPS performance requirements of very demanding random workloads.

Relational databases are often at the core of business applications and increasing their performance, while keeping storage power consumption and footprint to a minimum, reducing total cost of ownership (TCO) and helping in growing data centers constraints. The deployment of Tier 0 Flash drives together with slower tiers such as Fibre Channel and SATA drives enables customers to structure the application data layout where each tier of storage meets the I/O demands of the application data it hosts.

Audience

This white paper is intended for Oracle Database administrators, storage architects, customers, and EMC field personnel who want to understand the implementation of enterprise Flash drives in Oracle Database environments to improve the performance of business applications.

Technology overview


Symmetrix DMX-4 overview

The new Symmetrix DMX-4 system, as shown in Figures 2 and 3, is the next generation in the DMX series and extends EMC's leadership in the high-end enterprise storage market. The DMX-4 delivers immediate support for the latest generation of disk drive technologies: Flash drives and 4 Gb/s FC drives for high performance and SATA II for high capacity. Symmetrix DMX-4 is the first and only high-end storage system that can support all of these latest generations of disk drive technologies. The DMX-4 with the latest Enginuity release 5773 has been optimized for maximum performance and tiered storage functional flexibility.



**Combinations may be limited or restricted based on configuration*

Figure 2. Symmetrix DMX-4: World's largest high-end storage array



**Combinations may be limited or restricted based on configuration*

Figure 3. Symmetrix DMX-4 950: Entry point for DMX-4 technology

Enterprise Flash drives

With the introduction of the Symmetrix DMX-4 running Enginuity 5773, EMC now supports enterprise-class Flash drives. The enterprise-class EMC Flash drives are constructed with nonvolatile semiconductor NAND Flash memory and are packaged in a standard 3.5-inch disk drive form factor used in existing Symmetrix DMX-4 drive array enclosures. These drives are especially well suited for low-latency applications that require consistently low read/write response times. These drives exhibit as much as 30 times improvement in IOPS over 15k rpm HDD technology.

Flash drives benefit from the advanced capabilities that Symmetrix provides, including local and remote replication, cache partitioning, and priority controls.

Which database workloads best fit enterprise Flash drives?

Random reads workloads

In the case of a read request, if the requested data is in the global cache either because of recent read or write, or due to sequential prefetch, the request is immediately serviced without additional disk I/O. A read serviced from cache without causing disk access is called a read hit. If the requested data is not in the global cache, the Symmetrix must retrieve it from disk; this is referred to as a read miss. A read miss will incur increased I/O response time due to the innate mechanical delays of hard disk drives.

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

DSS and BI workloads

When Oracle performs full table scan operations such as batch, reports, or during backups, it will issue sequential reads with large I/O (up to 1 MB in size). Such sequential reads are long operations and rely on best throughput (MB/s) rather than best latency (I/O response time). The better the throughput the quicker the report, query, or backup will complete. Improving backup performance using TimeFinder® and RMAN incremental backups is described in the following Symmetrix and Oracle joint-logo white paper, *Reducing Backup Window and Recovery Time with Oracle Database 11g RMAN and EMC TimeFinder/Clone - Applied Technology*. However improving the performance of Decision Support Systems (DSS) and Business Intelligence (BI) is an important task for many mission-critical data warehouse and BI systems. When Symmetrix identifies sequential read streams it will attempt to use prefetch to bring the requested data into cache ahead of time, thereby increasing the cache read-hit rate for the following sequential read I/O requests.

However there are two common phenomena in sequential read workloads. The first is that even when Oracle issues large sequential read requests, they commonly get broken down by too many layers of striping (EMC recommends not more than two) and become random by the time they get to the storage. The second is that with high user or query concurrency, the requests will seem as random I/Os to the storage.

EMC provides specific best practices for DSS and BI workloads to maximize Symmetrix prefetch effectiveness and throughput predictability.

When using spinning hard disk drives (HDD), if the I/O workload arriving to the storage is indeed sequential (by following our best practices and reducing layers of striping) even high-density slow drives like SATA can support high throughput. But as explained earlier, when the workload becomes more random the ability of HDDs to retain their throughput drops to less than 50 percent and in many cases 15k rpm drives are offered in BI systems to support the workload requirements. Flash drives are not susceptible to such fluctuations and also they can maintain their throughput as they have double the bandwidth over HDD in doing sequential reads, even as the read workload becomes random.

Writes workloads

It is important to understand that any I/O request from the host is serviced by the Symmetrix from its global cache. Under normal circumstances, a write request is always written to cache and incurs no delay due to physical disk access. If write activity experiences high response time it usually points to very large I/O size (and is therefore expected to take longer) or points to issues such as high I/O queuing due to not enough connectivity. Such connectivity imbalances can occur when not enough LUNs are presented to the host (host queuing), not enough HBA ports (HBA queuing), not enough Symmetrix front-end or back-end ports are utilized, or not enough disks are used. Therefore write performance issues are best dealt with by Root Cause Analysis and be directly fixing any connectivity imbalances, rather than assuming that Flash drives will immediately fix the problem.

Oracle online redo logs and archived log file activity is mostly sequential writes and because writes are serviced from Symmetrix cache, logs are not necessarily best candidates for placement on Flash drives and can be placed on HDD to leave more room for other files on the Flash drives. However if there is enough room, logs will still benefit from being on Flash drives as disk I/Os to Symmetrix cache can be serviced much faster.

Latency critical workloads

Database and application managers can easily point to mission-critical applications that, if made much faster, will directly effect an increase in business revenue and productivity. In a similar way the storage managers can point to these same applications, since for performance planning reasons they use a large number of *short stroked* drives, in other words, data is laid out on many partially populated disks in order to satisfy IOPS and latency requirements. When such applications are identified, Flash drives can provide two very important benefits. First, a single Flash drive can replace many short stroked drives by its ability to provide a very high transaction rate (IOPS). This reduces the total number of drives needed for the application, increases power saving by not having to keep many spinning disks, and eventually means reduced floor space in the data center. Second, Flash drives provide very low latency so applications where a predictable low response time is critical, and not all the data can be kept at the host or Symmetrix cache, will greatly benefit from using such drives. Because no rotating media exists in Flash drives, 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 drives.

Summary of enterprise Flash drive strategies

- Place the entire database on EFDs when database performance ties directly to business revenue.
- Place a portion of the database on EFDs as a storage Tier 0. Carefully identify the busiest portion that will benefit the most from EFDs based on joint EMC and Oracle guidelines. Examples are tablespaces, materialized views, indices, or temp.
- Place the active database partitions on EFDs with an ILM strategy in mind. As they get less active they will be migrated 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 OLTP applications. However, many 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 will benefit as well, although to a lesser degree than random.

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 the overall performance benefits will be in proportion to the time Oracle was waiting to serve I/O. However, almost any workload will benefit from EFDs. Ensure there is enough CPU power, connectivity, and I/O concurrency to benefit from the new Tier 0 capabilities.

Which database components to place on Flash drives?

Inspection of an Oracle AWR report

When possible, the entire database can be placed on Flash drives. In some cases however, it is more cost-effective to place only certain database components on Flash drives and leave the rest on hard disk drives. This section will focus on recognizing workloads and database components that best fit Flash drives by using Oracle statistics. The use of Flash drives as part of an ILM solution is discussed in more detail later.

It is possible to review Oracle statistics using the statspack report or Automatic Workload Repository (AWR) report, or by querying the v\$sysstat view directly, to identify whether the database will benefit from being on Flash drives. Furthermore, the statistics can also be used to determine the database components that will benefit the most. Since an AWR report provides data that is averaged over the report interval (two points in time, or snapshots), make sure to keep the interval reasonable and during peak activity. A one hour interval is usually appropriate. For example, looking at an Oracle AWR report, a number of areas can easily be focused on:

Load Profile

At the beginning of the AWR report, in the *Load Profile* section, inspect the logical reads, physical reads, and physical writes. If we recognize a read-mostly workload where a significant number of reads are physical (reads served from storage rather than Oracle cache) this indicates a workload that may greatly benefit from Flash drives.

Load Profile

	Per Second	Per Transaction	Per Exec	Per Call
DB Time(s):	64.8	0.2	0.02	0.14
DB CPU(s):	3.9	0.0	0.00	0.01
Redo size:	2,105,159.7	5,091.2		
Logical reads:	26,125.1	63.2		
Block changes:	10,572.2	25.6		
Physical reads:	9,627.9	23.3		
Physical writes:	3,782.3	9.2		
User calls:	478.6	1.2		
Parses:	2.3	0.0		
Hard parses:	0.0	0.0		
W/A MB processed:	29,543.3	71.5		
Logons:	0.2	0.0		
Executes:	2,670.0	6.5		
Rollbacks:	4.0	0.0		
Transactions:	413.5			

Figure 4 AWR report – Load Profile

Top 5 Timed Events

Next inspect the **Top 5 Timed Events** section. Look for the *db file sequential read* event, which indicates small data block reads and therefore random workload (the name is misleading). If you find it check the time ratio (*% of Total Call Time*) spent on this event relative to the others. If the ratio is high, then it provides a good indication that Flash drives can be of great benefit to this workload. Both response times for random reads can be significantly reduced, as well as the disk I/O retrieval bottleneck removed. If, however, high time ratio exists for wait events indicating sequential workload, such as *db file scattered read*, or *direct path read* (read to PGA memory usually by parallel query slaves), then refer to the earlier paragraph about Flash drive benefits for DSS workloads. As explained before, even if the workload seems sequential by Oracle, it may be random at the storage layer due to LVM striping (including Oracle ASM) or high user concurrency. In that case Flash drives may still be of great benefit (although not to the same extent as small block random reads). It is highly recommended to involve EMC professionals to provide storage statistics to accompany the AWR report analysis.

Event	Waits	Time(s)	Avg wait (ms)	% DB time	Wait Class
db file sequential read	7,286,335	47,332	6	81.05	User I/O
DB CPU		3,498		5.99	
db file parallel read		824	17	4.84	User I/O
gc cr grant 2-way	2,556,455	1,243	0	2.13	Cluster
log file sync	369,662	1,236	3	2.12	Commit

Figure 5 AWR report – Top 5 Timed Events

Instance Efficiency

Inspect *In-memory Sort %* in the **Instance Efficiency** section. Also look for a wait event called *direct path read temp* and see if it is high on the top 5 wait events list. If a high percentage of sorting is done on disk and causing delays it is likely that placing TEMP files on Flash drives will remove this bottleneck.

Buffer Nowait %:	99.99	Redo NoWait %:	100.00
Buffer Hit %:	63.15	In-memory Sort %:	100.00
Library Hit %:	100.07	Soft Parse %:	98.80
Execute to Parse %:	99.91	Latch Hit %:	99.76
Parse CPU to Parse Elapsed %:	0.02	% Non-Parse CPU:	100.00

Figure 6 AWR report – Instance Efficiency

Tablespace I/O Stats and File I/O Stats

There are two sections of interest at the end of the report. The first is called **Tablespace I/O Stats** and the second is **File I/O Stats**. First inspect the tablespaces that appear at the top of the list (highest I/O rate). If the average response time (*Av Rd (ms)*) numbers are high (remember that averages can be misleading as they hide bursts and reduce peaks over the sample time) these tablespaces can be candidates for moving to Flash drives. These can be data files, temp, undo, index, and so on. This provides an alternative for moving the entire database. Ensure that the writes/sec (*Av Writes/s*) are not very high relative to the reads/sec (*Av Reads/s*) since, as explained earlier, write workloads will already benefit from Symmetrix cache.

Tablespace	Reads	Av Reads/s	Av Rd(ms)	Av Blks/Rd	Writes	Av Writes/s	Buffer Waits	Av Buf Wt(ms)
DATA_0	5,268,726	5,846	9.00	1.00	2,891,970	3,209	100	4.20
INDEX_0	3,397,545	3,770	5.18	1.00	170,842	190	1,865	5.43
UNDOTBS1	103	0	1.26	1.00	25,331	28	35	0.00
SYSAUX	703	1	3.12	1.18	357	0	0	0.00
SYSTEM	28	0	8.57	1.00	35	0	0	0.00
UNDOTBS2	1	0	10.00	1.00	1	0	2	5.00
USERS	1	0	10.00	1.00	1	0	0	0.00

Figure 7 AWR report – Tablespace I/O Stats

A review of the data files' I/O stats provides finer granularity of reads and average response times.

Moving only certain data files and not entire tablespaces to Flash drives is not recommended since Oracle spreads new data extents across all data files belonging to a specific tablespace.

Using this simple inspection of Oracle statistics can provide information such as whether the workload would significantly benefit from Flash drives and then, if the entire database cannot be moved, which database components will benefit most by moving them to Flash drives.

Use Case #1: Benefits of moving an entire OLTP database to EFDs

Part I – Database on one RAID 5 HDD vs. one RAID 5 EFD

To compare the performance advantages of enterprise Flash drives relative to hard disk drives (HDD), two identical Oracle 11g databases on ASM were deployed on a two dual-core Linux servers, each on to a single RAID 5 (3+1) group, one on 146 GB Flash drives, the other on 300 GB¹ 15k rpm HDD. Both configurations had the RAID 5 group divided into eight logical volumes (LUNs) of 45 GB each that were used for the data files. Since database logs don't benefit greatly from Flash drives, the redo logs were placed on other HDD in both tests. An industry-standard OLTP benchmark with a 60/40 read/write ratio was used to generate identical workloads against both databases with a 64-bit Linux server driving the activity. The results are shown in Figure 8.

¹ Both 300 GB 15k rpm and 146 GB 15k rpm disks have the same speed. To avoid the capacity differences, the same size logical volumes were created on both Flash drives and HDD, and the same number of disks was compared.

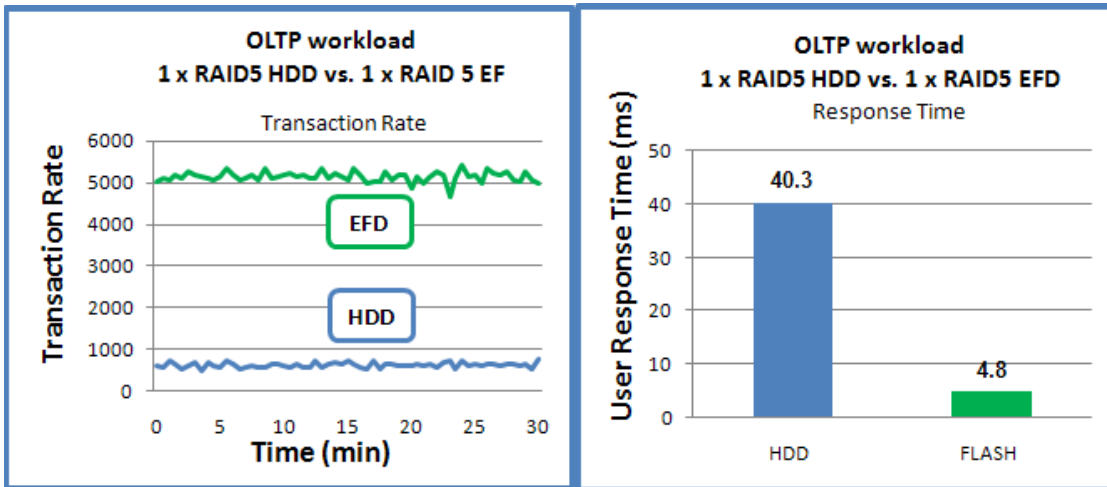


Figure 8. Transactions per minute rate and response time of Flash vs. HDD

As Figure 8 shows, in the configuration used for the test, the Flash drives sustained an average of 5,100 transactions per minute (TPM) and the HDD sustained around 620 TPM, roughly an eight-fold improvement in sustained TPM over HDD. In addition to the increased IOPS (I/O per second), the response time on the Flash drive was one-tenth of that observed on the HDD.

This use case points to an important fact, that changing only the disks in the configuration can result in a significant improvement in IOPS and response time. However, with any change, an analysis of the system should take place to make sure it is balanced for its new capabilities. For example, in this case, the bottleneck moved from back-end disk I/O to the host and resulted in a high level of queuing on the HBAs. The number of HBAs (two in this case) was not sufficient to achieve the most benefits out of the Flash drives. Although not shown here, by adding more HBAs, when the server architecture and CPU can support it, or adding more servers, such as by using Oracle Real Application Clusters (RAC), host I/O can be increased and an even higher transaction rate would have been realized.

During these tests all writes were serviced by Symmetrix cache so there was no difference in write performance.

Part II – Database on 80 RAID 1 HDDs vs. 12 RAID 5 EFDs

The next test carried out showed the advantage in moving an entire database from 80 146 GB 15k rpm RAID 1 disks to 12 146 GB EFDs in a RAID 5 configuration. The following environment outlines the software and hardware components used.

Table 1. Hardware and software components

		Description
1 x DMX-4 2500	Enginuity	5773.130.90
	Cache	64 GB Cache (4 x 32 Gb mirrored cache)
	HDD	80 x 146 GB 15k rpm (RAID 1)
	EFD	8 x 146 GB EFD (RAID 5) + 1 hot spare
2 x Dell R900	Operating System	RHEL 5.1
	CPU	4 x Quad core each
	HBA	2 x Dual port LP11002 each
	PowerPath®	5.1.1
Oracle	Enterprise Edition	Oracle Database 11g R1
	RAC, ASM	Two nodes RAC and Oracle ASM with separate disk groups for +DATA, +LOG, +TEMP, +FRA, and +FLASH
	SGA	11 GB on each node

The aim of the test was to demonstrate the significant increase in transactions and a corresponding decrease in response time. A 1 TB database was created across a two-node Oracle RAC, Oracle Database 11g and ASM. A similar OLTP workload as before was used in each test to generate transaction workload from both cluster nodes. The exact same workload and database configuration was used for the HDD and EFD tests.

The test results, in Figure 10, show a two-fold increase in transaction rate, with a 50 percent decrease in response time. It points to a vast consolidation achieved by moving from the 80 RAID 1 HDDs to only 12 EFDs in RAID 5 (12 EFDs were necessary to contain the entire 1 TB database in this use case). Such consolidation has the additional saving in data center cooling, energy consumption, and potentially floor space. It also shows the advantage of running the entire 1 TB database on EFDs by reducing response time (improving user experience) and more than doubling the transaction rate.

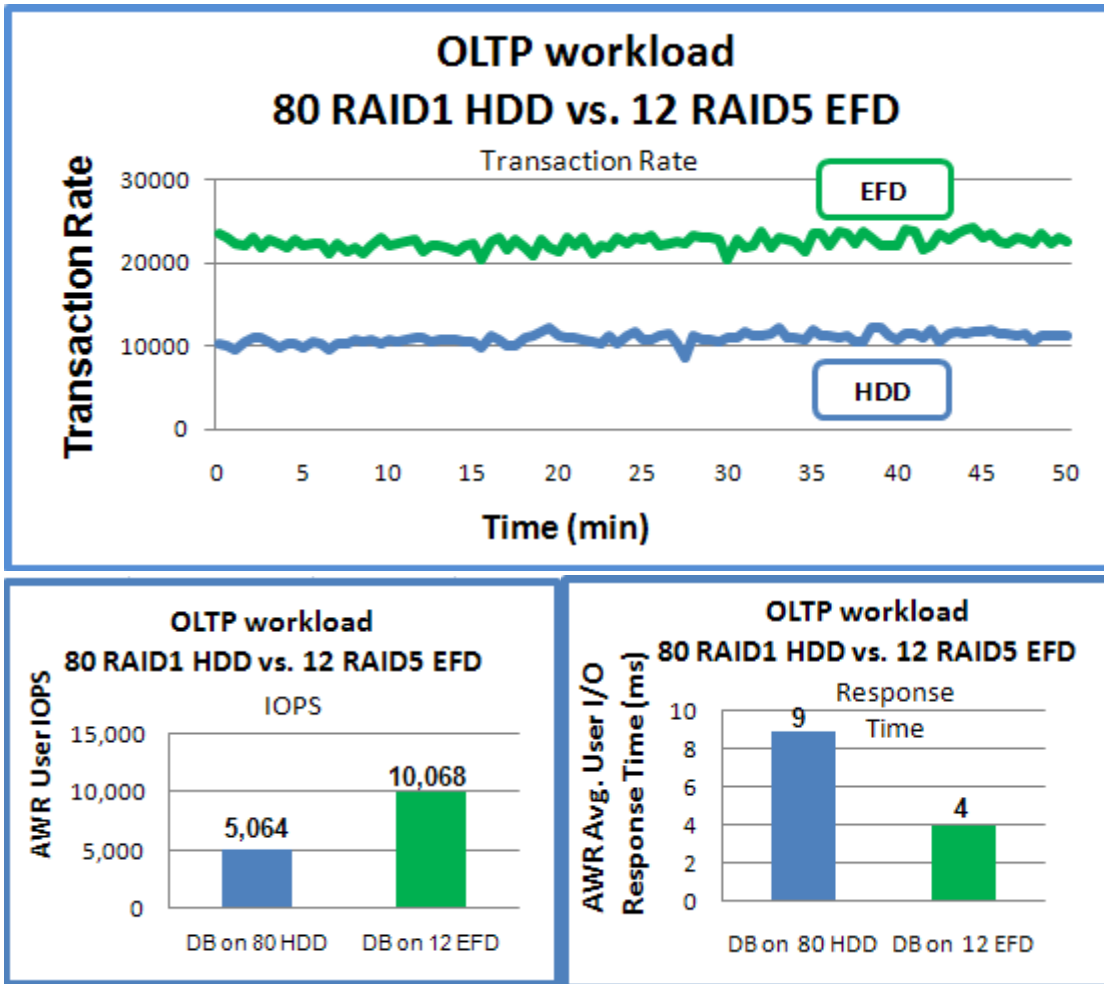


Figure 9. Oracle full database move from HDD to EFD

Use Case #2: A single tablespace migration to Flash

A 1 TB database was created across a two-node RAC distributed across 80 x 146 GB 15k rpm (RAID 1) devices.

The aim of this use case was to show the advantages in moving a portion of the database to enterprise Flash drives (EFDs). The use case showed an environment where only 25 percent of the I/Os were served from disk (the rest were from the 11 GB Oracle SGA). A single tablespace that was found to be the busiest (75 percent of these I/Os) was moved from 80 HDD to eight Flash drives.

Using a standard OLTP workload generator (as described earlier) a baseline was achieved where the heaviest utilized component was seen to be the disks, running at 80 percent utilization.

Oracle AWR reports from the run showed that the Oracle buffer cache had a 75 percent hit rate (25 percent of the I/Os were served from disk). It also showed that tablespace (STOK) was running 75 percent of the workload. Therefore it was decided to migrate the STOK tablespace to eight EFDs in a RAID 5 configuration. After the tablespace move the remainder of the database continued to run across the 80 HDDs, however they experienced a large decrease in both utilization and response times. Potentially many could have been removed.

The test results in Figure 10 show a two-fold increase in transaction rate and a 50 percent decrease in response time for the entire database. This was an excellent outcome for moving only a single tablespace from 80 HDD to eight EFDs, and where 75 percent of the database transactions' IOs were actually served from Oracle buffer cache. By adding EFDs and improving less than 25 percent of the workload (less since STOK was doing only 75 percent of the physical I/Os) the transaction rate doubled.




Figure 10. Oracle tablespace move from HDD to EFD

Flash drives and Information Lifecycle Management (ILM) strategy

In most enterprise applications, the most recent data has the most demanding latency requirements over older data. This type of data classification is the beginning of an Information Lifecycle Management strategy, or ILM. Data classification is important in order to provide applications with the most cost-effective storage tier to support their workload needs. It can be done by placing each application on the storage tier that fits it best, but it can also be achieved by using multiple storage tiers within the same application.

A common way for deploying a single database over multiple tiers is by file type. For example archive logs and backup images can use SATA drives while redo logs and data files can use Fibre Channel HDD. It is now possible to add a new storage tier using Flash drive technology and place latency-critical data files, indices, or temp files on them as discussed earlier.

However, when the database is large, in order to achieve optimum utilization of drive resources, it might be better to place on Flash drives only the data that is accessed most frequently and/or has the most demanding latency requirements. Many databases can achieve this by using table partitioning.

Using analysis techniques such as those presented in this paper, the customer can determine the most active tablespaces and files that Flash drives are uniquely able to help. Placing the LUNs for these tablespaces on Flash drives will provide significant benefit while not requiring the entire database to be on Flash.

Table partitioning also creates subsets of the table, usually by date range, which can be placed in different data files, each belonging to a specific storage tier. Table partitioning is commonly used in data warehouses where it enhances index and data scans. However, with an ILM strategy in mind, customers should consider the advantages of using table partitioning for OLTP applications. While partitioning allows the distribution of the data over multiple storage tiers, including Flash drives, it does not address the data movement between tiers. Data migration between storage tiers is out of the scope of this paper. Solutions in this space are available by using Symmetrix virtual LUN technology, the Oracle online redefinition feature, or by using host volume management features. An example of using table partitioning is shown in Figure 11.

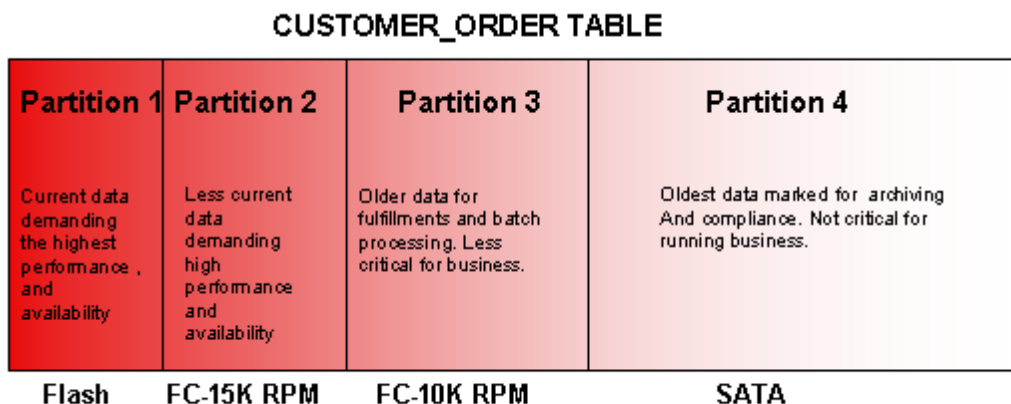


Figure 11. A partitioned table using tiered storage levels

Conclusion

Incorporation of Flash drives into Symmetrix provides a new Tier 0 storage layer that is 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 key Symmetrix software applications that enable advanced management tools.

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.

Symmetrix now combines the performance and power efficiency of Flash drive technology with traditional disk drive technology in a single array managed with a single set of software tools; this delivers advanced functionality, ultra-performance, and expanded storage tiering options.