

EMC Solutions for Enterprises
EMC Tiered Storage for Oracle
ILM Enabled by EMC Symmetrix V-Max

Reference Architecture

Copyright © 2009 EMC Corporation. All rights reserved.

Published April, 2009

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

Benchmark results are highly dependent upon workload, specific application requirements, and system design and implementation. Relative system performance will vary as a result of these and other factors. Therefore, this workload should not be used as a substitute for a specific customer application benchmark when critical capacity planning and/or product evaluation decisions are contemplated.

All performance data contained in this report was obtained in a rigorously controlled environment. Results obtained in other operating environments may vary significantly. EMC Corporation does not warrant or represent that a user can or will achieve similar performance expressed in transactions per minute.

No warranty of system performance or price/performance is expressed or implied in this document. 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: H6234.1

Contents

About this Solution

Purpose	4
The business challenge	4
The technology solution	5
Environment profile	12
Solution components	16
Solution validation	18
Conclusion	20

About this Solution

Purpose

The purpose of this document is to illustrate the techniques available to enable nondisruptive data movement in an Oracle infrastructure. With the introduction of the Symmetrix® V-Max™ Series these requirements can now be supported using EMC's Enhanced Virtual LUN Technology. Virtual LUN technology simplifies the use of multiple storage tiers within the array (such as SATA, HDD, and Flash drives). It also serves as a building block on the way to deploying an Information Lifecycle Management (ILM) strategy. By placing database components on storage tiers that match their performance, protection, and cost requirements, the overall database performance is increased, and total cost is reduced.

Specifically, these tests will illustrate nondisruptive data transfer between different RAID types and storage tiers in support of implementing Oracle ILM and/or the correct use of a storage tier strategy enabled by Symmetrix V-Max Virtual LUN technology.

The business challenge

Today's IT departments are being requested by the business to solve several challenges around optimizing the location and ensuring the availability of their Oracle infrastructure:

- **Availability:** It is critically important that all the database functions and data remain fully available to systems and applications that require that data, throughout the migration through the different tiers of data storage
- **Performance:** Minimizing the impact on the performance of the production database throughout the process of data movement

- Manageability: Change control is a significant challenge for customers, any changes to customer production environments can cause significant risk to the business

The technology solution

The following section describes EMC's Symmetrix V-Max Virtual LUN technology, which enables movement of LUNs within the Symmetrix V-Max array, nondisruptively. This capability enables the dynamic movement of data from enterprise Flash drives (EFD) to hard disk drives (HDD) to SATA drives. Data can be also moved from Symmetrix devices with one RAID type to a different RAID type to meet the performance and protection requirements of a particular application.

Symmetrix V-Max

The Symmetrix V-Max system is a high-end, scalable storage array comprising a system bay and separate storage bays. The system scales from a single high availability (HA) node configuration to eight node configuration and a maximum of 10 storage bays. Online system upgrades are achieved by adding HA nodes. Each HA node contains two integrated director boards with multi-core CPU processing power, cache memory, front-end ports, and back-end ports.

The Symmetrix V-Max presents a simplified, modular hardware design that is available in two models, the entry point Symmetrix V-Max SE Series with Enginuity™, and the high-end, scalable Symmetrix V-Max system. Both systems support host connectivity to IBM mainframes, System i, and open system hosts over Fibre Channel, FICON, iSCSI and Gigabit Ethernet. Symmetrix V-Max Enginuity software provides many enhancements to Symmetrix feature portfolio such as Remote Data Facility (SRDF®), TimeFinder® and others. For example, SRDF/Extended Distance Protection, or SRDF/EDP, enhances support to two-site DR solution over extended distances with zero or near zero data loss. In this configuration the storage cache alone is used on the intermediate site for a temporary pass-through data store of the modified tracks before copying them over to the tertiary site.

The Symmetrix V-Max system's back-end design includes expandable storage capacity (drive enclosures and bays), with support for 1 TB SATA drives, 10k rpm and 15k rpm Fibre Channel drives, and 200 / 400 GB enterprise Flash drives.

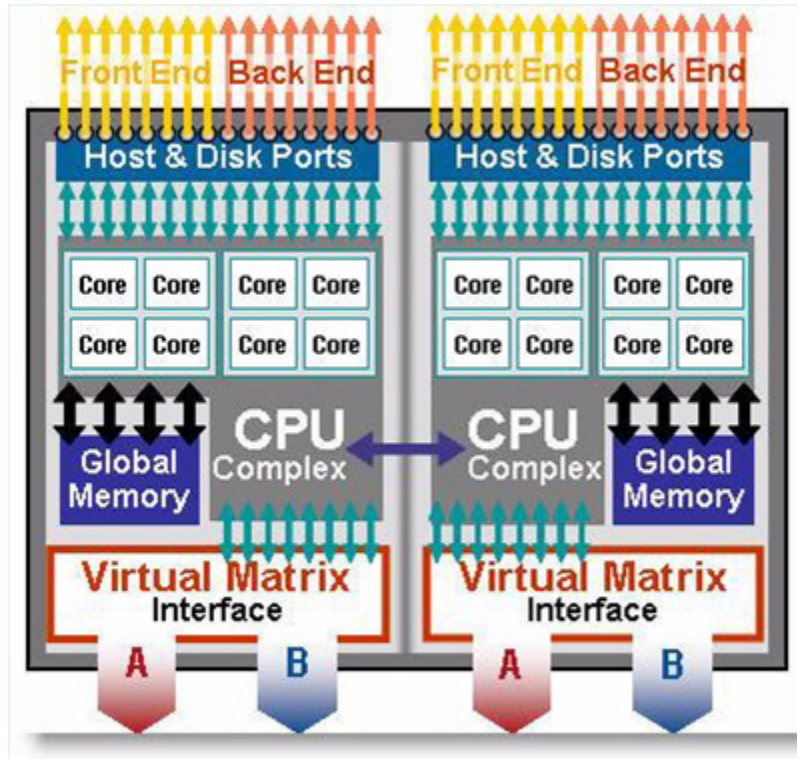


Figure 1 Symmetrix V-Max high availability node layout

Storage performance characteristics

The Symmetrix V-Max systems deliver performance, scalability, high-availability and protection to meet any requirements - from the small business with the SE series to the most demanding applications:

- High-performance integrated director board, including two quad-core processors per module, or 16 per node
- Enhanced, scalable memory, providing up to 1 TB of global cache (512 GB usable)
- Flash drive technology that delivers the ultimate performance with the lowest latencies
- Tools for tiered storage optimization, including Dynamic Cache Partitioning, Symmetrix Priority Controls, Symmetrix Optimizer, and Virtual LUN

Virtual LUN technology

EMC's Virtual LUN technology is a built-in feature on Symmetrix V-Max arrays that allows users to seamlessly migrate data between LUNs within the array without disruption to the applications. Until now most other storage-based methods for migrating data would require an outage to the applications before the migration would complete. Virtual LUN technology allows a user to migrate data to LUNs with better performance or other characteristics within the array, providing a greater level of control over the system and higher levels of service for business applications. The characteristics of a LUN that can be changed are dependent on RAID geometry and drive type.

All Virtual LUN migrations can be accomplished without the host ever knowing that anything has changed. [Figure 2](#) highlights a Virtual LUN movement example.




Figure 2 Virtual LUN movement

Virtual LUN migration

Virtual LUN migrations enable transparent, nondisruptive data mobility for standard Symmetrix volumes between storage tiers and between RAID protection schemes. Virtual LUN can be used to populate newly added drives or move devices between high performance and high capacity drives, thereby simplifying the use of tiered storage capabilities within a single Symmetrix array. Migrations are performed while providing constant data availability and protection.

RAID Virtual Architecture (RVA) allows, for the purposes of migration, two distinct RAID groups, of different types or on different storage tiers, to be associated with a logical volume. In this way, Virtual LUN allows for the migration of data from one protection scheme to another, for example RAID 1 to RAID 5, without interruption to the host or application accessing data on the Symmetrix device.

Data can be migrated to either configured or unconfigured space. In the case of the migration to unconfigured space, a target RAID group is created from the free pool in the array and migrated to. When the migration is completed, the original RAID group is deleted and the storage returned to the free pool.

Virtual LUN migration is fully interoperable with all other Symmetrix replication technologies - SRDF, TimeFinder/Clone, TimeFinder/Snap, and Open Replicator.

Virtual LUN migration can be managed via the SMC graphical interface, or the Solutions Enabler Command Line Interface (SYMCLI).

To perform a Virtual LUN migration, the following are required:

- Symmetrix V-Max array
- Enginuity operating environment for Symmetrix 5874
- Solutions Enabler 7.0
- Symmetrix Management Console 7.0 (optional)
- Symmetrix Optimizer license key

Note: All supported RAID protection changes can be found in the *EMC Solutions Enabler Version 7.0 Product Guide*.

Controlling Virtual LUN migration pace

A Virtual LUN migration initiates the synchronization of the source and target devices as a copy process at the back-end Disk Directors (DAs) on the array. Because of this, any other copy processes executed on the array may be impacted by this activity (i.e., TimeFinder/Clone and TimeFinder/Snap actions). As no Virtual LUN migration session can be aborted or canceled, the copy pace of the Virtual LUN migration can be adjusted using the Symmetrix Quality of Service (QoS) tools. Setting the QoS to either low or high can increase and decrease the speed of the migration while lessening the impact on other activities on the array. Customer application operations or processes running on the LUNs that are in a Virtual LUN migration session may be adversely affected by performance. This can be negated by setting the pace of the migration session to a slower speed.

The pace value can be set as integer values between 0 and 16 inclusively, with 0 being the fastest and 16 being the slowest. The default QoS value is 0.

Setting the pace value to 16 will allow the migration to synchronize approximately one track per minute, per device, effectively suspending the migration. Once the unrelated copy tasks have completed, the pace setting can be reset to 0. The QoS values can be set via SMC or Solutions Enabler.

ILM building block

Information Lifecycle Management (ILM) is a concept that has evolved over time and is driven by the need to improve and organize data storage. ILM is composed of three components, people/policy (process)/technology, an approach that is used to organize data as it progresses through its lifecycle; from the time data is created and stored until data is no longer needed or retired.

ILM is about right data/right location/right cost, thereby improving the business value of information (policy). To make ILM a reality, organizations need to control where information resides within the array. Using Virtual LUN technology as a building block component (technology) of ILM enables effective data migration. This is accomplished through the nondisruptive migration of data across storage tiers. This reduces an enterprise's TCO, while still providing the best performance and availability to the most valuable information in the system. Virtual LUNs accomplish this without any application downtime. Virtual LUNs ensure that business value can be maximized while still lowering costs, thus providing enterprise customers with a strategic advantage over competitors.

Autoprovisioning Groups

Solutions Enabler V7.0 introduces an easier, faster way to provision storage in the Symmetrix V-Max system. Initiator groups allow storage administrators to create groups of host initiators, front-end ports, and logical devices. These groups are then associated to a masking view, from which all controls can be managed.

Note: Initiator groups are not supported on Symmetrix DMX™ arrays running Enginuity 5773 and earlier. Storage administrators should continue to use the `symmask` and `symmaskdb` commands to mask devices in Symmetrix DMX arrays.

Most of the applications running on Symmetrix arrays require a fault-tolerant environment, with multiple paths to devices, and clusters of servers. Previous versions of Solutions Enabler required one command for each initiator/port combination through which the devices would be accessed. A new command, `symaccess`, provides all the storage provisioning requirements for Symmetrix V-Max arrays running Enginuity 5874.

Storage provisioning with `symaccess` allows users to create a group of devices, a group of director ports, a group of host initiators, and with one command, associates them in what is called a masking view. Once a masking view exists, devices, ports, and initiators can be easily added or removed from their respective groups.

This feature reduces the number of commands needed for masking devices, and allows for easy management of the masking view.

The steps for creating a masking view are:

1. Create a storage group (one or more devices).
2. Create a port group (one or more director/port combinations).
3. Create an initiator group (one or more host WWN or iSCSI).
4. Create a masking view containing the storage group, port group, and initiator group.

When a masking view is created, the devices are automatically masked and mapped.

Environment profile

The following test scenario infrastructure provided the baseline for the Oracle data movement scenarios. For more information, refer to “[Key results](#)” on page 18.

Hardware resources

[Table 1](#) lists the hardware resources used in the EMC Tiered Storage for Oracle ILM Enabled by EMC Symmetrix V-Max environment.

Table 1 Hardware resources

Equipment	Quantity	Configuration
Storage	1	EMC Symmetrix V-Max: 64 GB Cache (4 x 32 GB mirrored cache) HDD drives: 70 x 300 GB 15k rpm, RAID 1 52 x 300 GB 15k rpm, RAID 5 8 x EFDs
SAN	2	SAN switch, 64 port
Oracle RAC database server	2	11.1.0.6
ECC management server/SMC server	1	Dual-Core 3.4 GHz 16 GB RAM Two 4 Gb HBAs

Software resources

Table 2 lists the software resources used in the solution.

Table 2 Software resources

Title	Version	Comment
Red Hat	5.1 (64-bit)	OS for database server
Microsoft Windows	2003 SP2	32-bit
Oracle Database/Cluster/ASM	11.1.0.6	Database/cluster software/volume management
Enginuity	5874	Symmetrix storage array OS
Symmetrix Management Console (SMC)	7.0	Symmetrix storage management software
EMC PowerPath®	5.1 SP2	HBA load balancing and redundancy
Solutions Enabler	7.0	Symmetrix Command Line Interface host-based utility STP - statistics collection daemon. The storstpd daemon runs in conjunction with EMC Solutions Enabler

Physical architecture

Figure 3 illustrates the overall physical architecture of the solution.

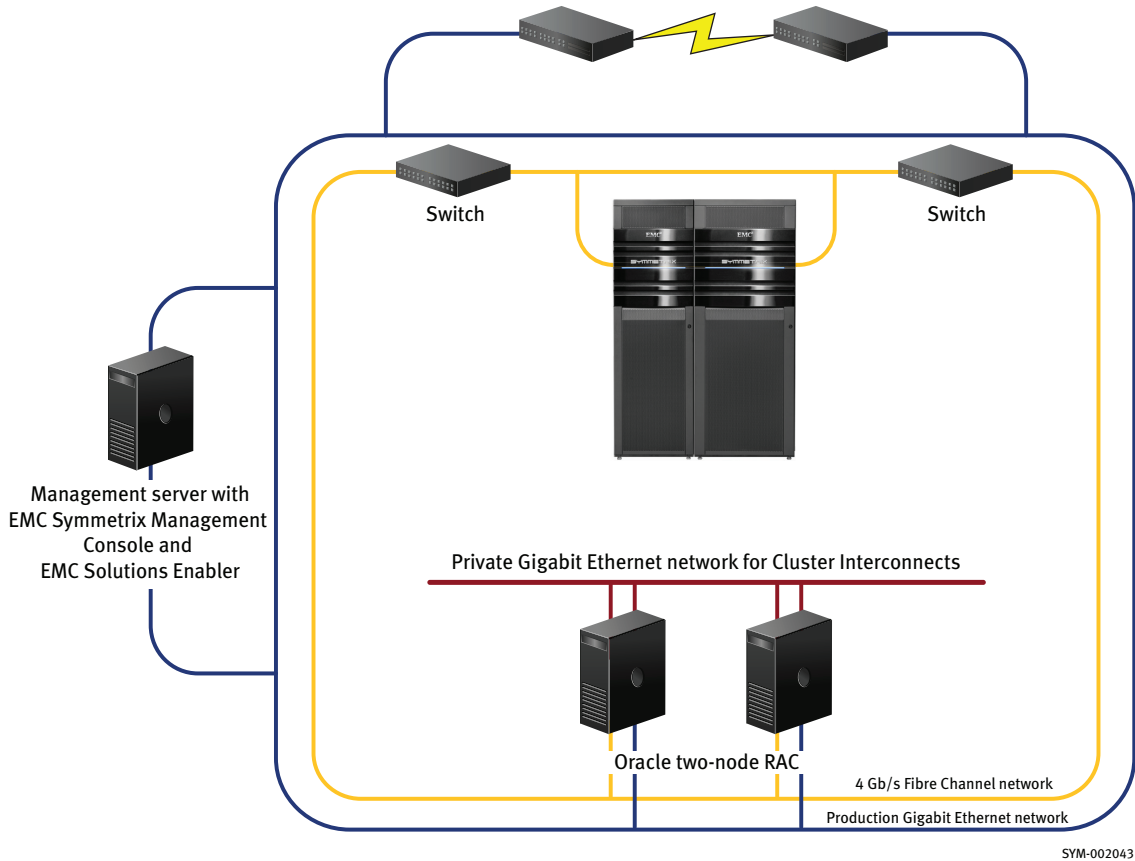


Figure 3 Overall physical architecture

The Symmetrix V-Max array used in this environment has one drive bay configured that can house a maximum of 240 drives.

For the 1 TB database, using Oracle ASM, hypervolumes were used. For this reference architecture, each physical spindle is divided into six hypervolumes on RAID 1 disks, see Figure 4.

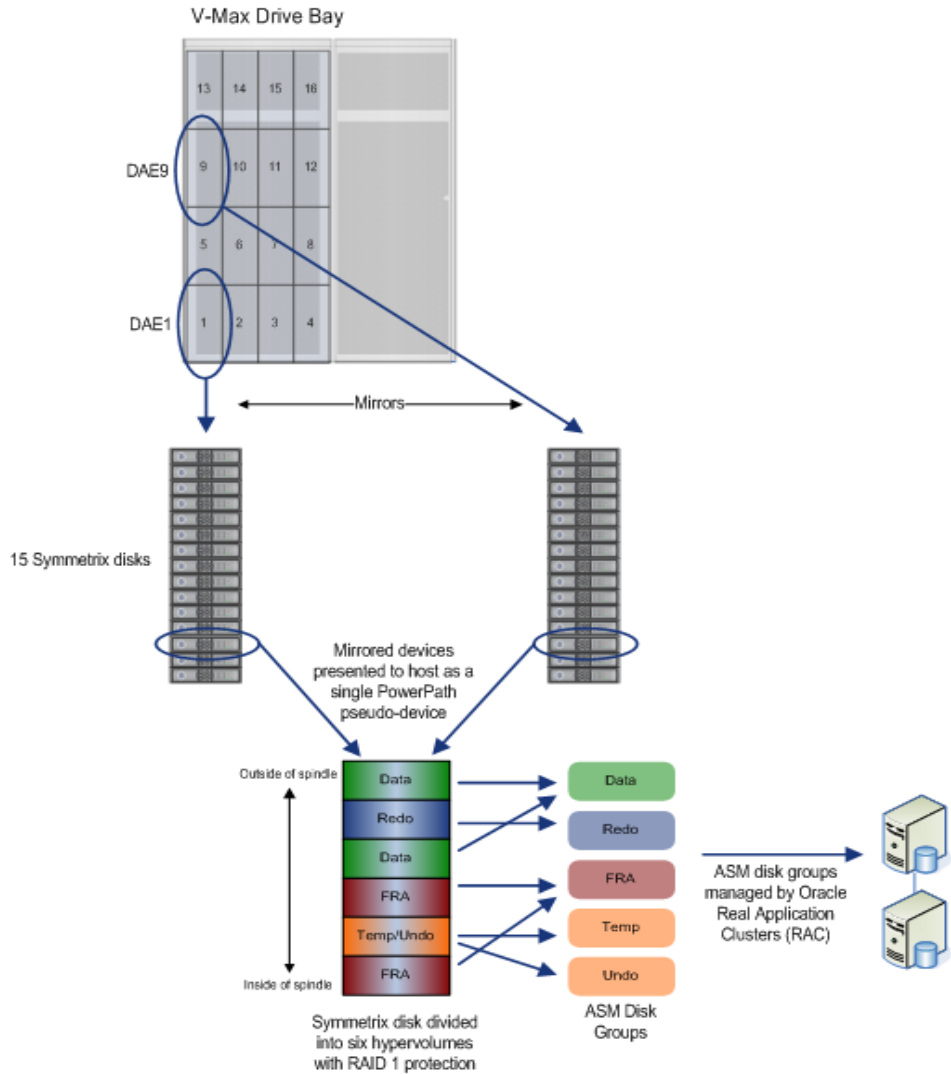


Figure 4 Symmetrix V-Max array disk layout

Solution components

In order to meet the business challenge, the following solution components were used.

EMC Symmetrix Management Console — SMC is a simple, intuitive, browser-based user interface for the configuration and management of Symmetrix arrays. SMC 7.0 was developed to concurrently support all the features of Enginuity 5874.

SMC presents the functionality of the Solutions Enabler SYMCLI (command line interface) in a browser interface.

Some of the new features available in SMC V7.0 are listed below:

- New wizards:
 - SRDF Replication wizard: guides you through the process of creating SRDF pairs and both concurrent and cascaded SRDF configurations.
 - LUN Migration wizard: guides you through the process of moving existing devices from one location (disk group) to another within a Symmetrix V-Max array running Enginuity 5874 or later. In addition, this wizard will also allow you to change the protection type of the devices.
 - Masking wizard: guides you through the process of creating masking records, or creating/modifying masking views.
- SMC on the Symmetrix V-Max service processor
Starting with Enginuity 5874, SMC is available on the Symmetrix service processor as part of EMC's new management integration. This feature is disabled by default. To enable it, contact EMC Customer Service.

EMC PowerPath — PowerPath works with the storage system intelligently managing I/O paths, and supporting multiple paths to a logical device. In this solution PowerPath manages four I/O paths and provides:

- Automatic failover in the event of a hardware failure. PowerPath automatically detects path failure and redirects I/O to another path.
- Dynamic multipath load balancing. PowerPath distributes I/O requests to a logical device across all available paths, thus improving I/O performance and reducing management time and downtime by eliminating the need to configure paths statically across logical devices.

Oracle infrastructure

The following Oracle components were used to create the Oracle infrastructure.

Oracle Database 11g Enterprise Edition — Oracle Database 11g Enterprise Edition delivers industry-leading performance, scalability, security and reliability on a choice of clustered or single servers running Windows, Linux, and UNIX. It provides comprehensive features easily managing the most demanding transaction processing, business intelligence, and content management applications.

Oracle Database 11g RAC — Oracle Real Application Clusters (RAC) is an optional feature of Oracle Database 11g Enterprise Edition. Oracle RAC supports the transparent deployment of a single database across a cluster of servers, providing fault tolerance from hardware failures or planned outages. Oracle RAC supports mainstream business applications of all kinds. This includes Online Transaction Processing (OLTP) and Decision Support System (DSS).

Oracle Automatic Storage Management (ASM) — Oracle ASM is an integrated database filesystem and disk manager. With ASM, filesystem and volume management capabilities are built into the Oracle database kernel. This reduces the complexity of managing the storage for the database.

In addition to providing performance and reliability benefits, ASM can also increase database availability because disks can be added or removed without shutting down the database. ASM automatically rebalances the files across the disk group after disks have been added or removed. It supports single-instance Oracle Databases and Oracle Real Application Clusters (Oracle RAC). ASM is Oracle's recommended volume manager solution.

Solution validation

This solution illustrates the various methods by which an EMC / Oracle customer can implement nondisruptive data transfer between different RAID types in support of an Oracle data management strategy enabled by Symmetrix V-Max VLUN technology. The following test scenarios were completed for the solution:

- Migration of an Oracle database to Symmetrix V-Max using EMC SRDF
- Nondisruptive movement of an entire Oracle database from RAID 1 to RAID 5 using EMC Virtual LUN
- Nondisruptive movement of an Oracle ASM disk group from HDD to EFDs using EMC Virtual LUN

Key results

Migration of an Oracle database to Symmetrix V-Max using EMC SRDF

Result — Seamless data transfer of the database to the remote array was achieved, with both storage systems' devices allocated to both nodes of the Oracle RAC database, with only one array's devices in use at any one time. Once data transfer completed the relationship was suspended between the SRDF R1 devices to the corresponding SRDF R2 devices, allowing switchover to the Symmetrix V-Max array in a matter of minutes. In a normal SRDF suspended state the R2 devices are in a write-disabled state. Once the relationship was suspended, the R2 devices were write-enabled, allowing the hosts access to the Symmetrix V-Max array instantly. This ensured minimal downtime to production while also providing a failback option if required.

Nondisruptive movement of an entire Oracle database from RAID 1 to RAID 5 using EMC Virtual LUN

Result — During this testing the Oracle database remained up and running, with no DBA administration actions required. The source devices' server LUN address, as well as the EMC PowerPath pseudo name, remained consistent from the server perspective. Once completed, the Virtual LUN migration returned the original source devices to the free pool, which could be used for another application or function.

Nondisruptive movement of an Oracle ASM disk group from HDD to EFDs using EMC Virtual LUN

Result — Using Symmetrix V-Max Virtual LUN technology an individual ASM disk group (DATA) can be moved nondisruptively between different storage tiers.

For this test scenario the entire DATA disk group of an Oracle OLTP database was moved from 300 GB 15k HDD drives to 400 GB EFDs.

Additionally, this test illustrated that the Oracle infrastructure was migrated from 70 HDD drives to only eight EFDs.

Performance increased and response time reduced at the database level. As in the previous test, once completed the Virtual LUN migration returned the original source LUNs to the free pool, which could be used for another application or function. With the reduced utilization on the existing HDDs, a number could have been returned for repurposing.

All of this data movement was completed without any outage to the application. This type of data relocation offers a significant advantage in 24x7 production level environments.

Conclusion

Enterprises maintain a competitive advantage by placing increasing value on both the location and availability of their Oracle infrastructure. More companies are realizing that this information is critical to the business and are investing in their surrounding infrastructure. The data needs to be highly available and the data needs to be in the right place, at the right time and at the right cost to the enterprise. Understanding the value of this information over time is critical to designing a complete ILM infrastructure.

This solution meets the business challenges in the following manner:

- **Availability:** The solution illustrated the ease with which data can be moved by using the Symmetrix V-Max Virtual LUN capabilities. Virtual LUN illustrates the capability of changing the characteristics of LUNs, such as drive type and RAID type, to ensure the highest service levels to the customer, nondisruptive to the Oracle database. Data remains fully available to systems and applications that require that information, throughout the migration across the different tiers of data storage
- **Performance:** The Symmetrix V-Max system delivers scalable performance to meet the most demanding access, protection, and distribution requirements through:
 - High-performance director engines
 - Enhanced, scalable memory
 - Improved algorithmic intelligence that reduces processing overhead
 - Flash drive technology
 - Dynamic Cache Partitioning
 - Symmetrix Priority Controls
 - Symmetrix Optimizer
 - Virtual LUN technology
- **Manageability:** By using Virtual LUN, new storage tiers or RAID protections can be deployed without any changes to production, for example, not requiring the introduction of new LUNs, or changes in Oracle file locations (which in turn affect backup and DR operations).

Since 1995, EMC and Oracle have invested to jointly engineer and integrate their technologies. The combination of Oracle and EMC software and best practices used in an integrated fashion can greatly reduce the cost of designing, implementing, and operating your IT infrastructure.

To learn more about this and other solutions contact an EMC representative or visit:

<http://www.emc.com/solutions/application-environment/oracle/index.htm>