

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

Applied Technology



Abstract

This white paper illustrates the methods by which an EMC[®] Symmetrix[®] V-Max[™]/ Oracle customer can implement nondisruptive data movement between different RAID types and storage tiers in support of an Information Life Management (ILM) strategy and improved ROI by utilizing storage tiers that match Oracle database data cost and performance requirements.

April 2009

Copyright © 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 H6233

Table of Contents

Executive summary	4
Introduction	4
Audience.....	4
Information Lifecycle Management	5
ILM building block.....	5
Tiered storage	6
Technology overview.....	7
Symmetrix V-Max.....	7
Storage performance characteristics.....	8
Virtual LUN technology.....	8
Oracle infrastructure	11
Use cases for Oracle data relocation	12
Use case environment profile	12
Use case 1: Migration to the Symmetrix V-Max platform using EMC SRDF.....	13
Use case 1 infrastructure	13
Migration method.....	13
Migration advantages	15
Use case 2: Nondisruptive movement of an entire database from RAID 1 to RAID 5 using EMC Virtual LUN.....	16
EMC Virtual LUN method	16
EMC Virtual LUN advantage	17
Use case 3: Nondisruptive movement of an Oracle ASM disk group from HDD to EFDs using EMC Virtual LUN.....	18
Oracle ASM disk group movement via EMC Virtual LUN method.....	18
EMC Virtual LUN advantage	20
Conclusion.....	21

Executive summary

As data continues to grow rapidly, and applications become more integrated to derive better business knowledge, companies increasingly realize the importance of their Oracle database infrastructure. Oracle data needs to be highly available and protected on one hand, and the right data needs to be in the right place at the right time.

EMC's Enhanced Virtual LUN Technology enables movement of storage devices' data (host LUNs) within the Symmetrix® V-Max™ array, nondisruptive to the host and the Oracle infrastructure. Oracle data can be migrated to less performant or more performant disk technologies, or different tiers within the array while users have full access to the Oracle database and without changing the host LUN IDs and therefore not requiring any host change-control procedures. This capability enables the dynamic movement of data between enterprise Flash drives (EFD), hard disk drives (HDD), and SATA drives. Data can be also moved from one RAID type to another to meet the performance and protection requirements of a particular application. Additionally, Virtual LUN technology can also be leveraged to improve capacity utilization by moving data to underutilized spindles in the array.

Introduction

This solution provides customers with different methods in which an EMC® Symmetrix V-Max / Oracle customer can implement Oracle data movement. Today's IT departments are being requested by the business to solve several challenges around optimizing their Oracle infrastructure:

1. **Availability:** It is increasingly important that all mission-critical database and applications remain fully available even as data, through its lifecycle, changes performance, protection, and cost requirements and therefore needs to be moved to different storage tiers and RAID protections.
2. **Performance:** Often database performance requirements are not known during initial design and implementation. Other times with the growth of data, user community, and other infrastructure changes, data performance requirements have changed. Being able to perform nondisruptive migration within the array between different disk technologies, EFD, HDD, and SATA makes a solution to these situations viable.
3. **Manageability:** Change control is a significant challenge for customers, and any changes to customer production environments can cause considerable risk to the business. There is a need to provide a data migration solution that doesn't require host LUN changes such as introducing new LUNs when new storage tiers are added, or when data movement between storage tiers and RAID protection is required.

As part of an overall Oracle data movement strategy, a customer may wish to carry out a number of data migration and nondisruptive data movement options within their Oracle infrastructure. These include the use of virtual LUN migration technology, Open Replicator, TimeFinder®, SRDF® and other EMC data movement solutions.

The paper illustrates a nondisruptive data movement in an Oracle infrastructure using the Symmetrix V-Max Enhanced Virtual LUN migration technology (Virtual LUN). This use case provides a building block component for an Information Lifecycle Management (ILM) strategy, allows migration of data based on LUN granularity between storage tiers and RAID types, and increases the availability of the enterprise's Oracle infrastructure.

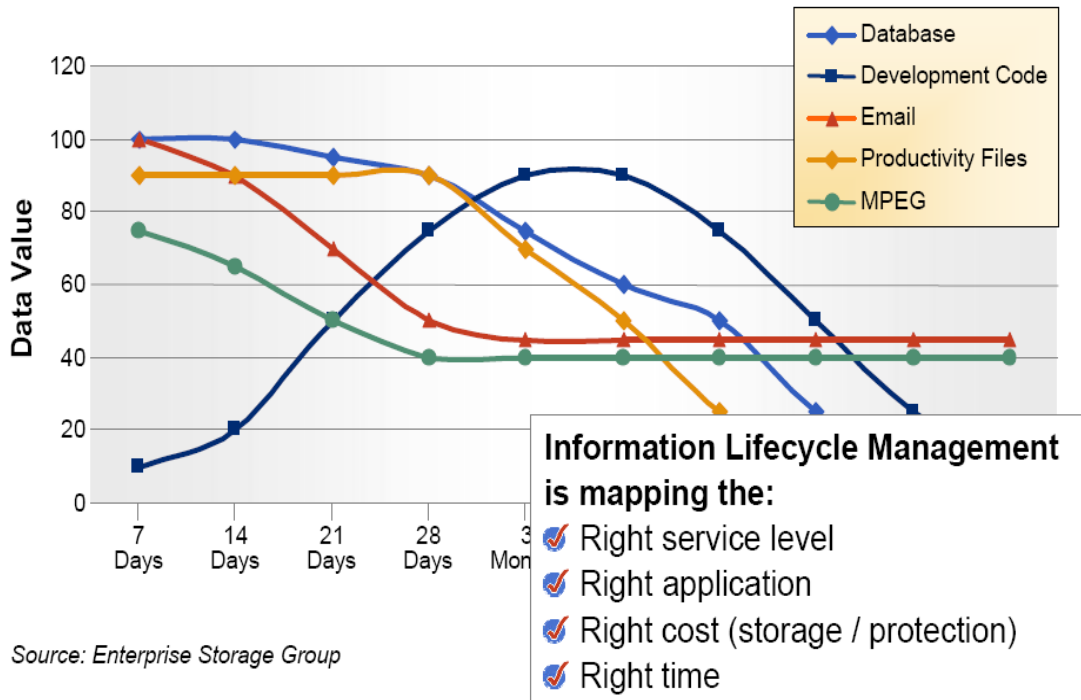
Audience

This white paper is intended for Oracle database administrators, storage architects, customers, and EMC field personnel who want to understand how EMC products and solutions can support the nondisruptive migration of Oracle databases.

Information Lifecycle Management

The following section illustrates the impact of lifecycle management of Oracle data from creation to retirement.

Information Lifecycle Management (ILM) is the practice of aligning the business value of information with the most appropriate and cost-effective infrastructure throughout its existence. The effective use of ILM allows users to lower the total cost of ownership (TCO) while still maintaining high service levels for the data and the applications that are accessed the most. Costs can be reduced by moving data with less business value to cheaper infrastructure, while maintaining higher valued data on better performing infrastructure. This is illustrated in Figure 1.



Source: Enterprise Storage Group

Figure 1 Information Lifecycle Management

Reducing the TCO can be accomplished by placing highly accessed and performance critical data on the fastest (and usually more expensive) storage tiers, like 15k rpm disks, or enterprise Flash drives, and placing less accessed data on lower cost large capacity drives, such as 10k rpm and SATA drives. Companies must examine their ILM and service level strategy to determine which application's data requires the highest performance. This analysis should be monitored and revised throughout the lifecycle of the application as requirements continue to change. Symmetrix V-Max provides an extensive set of solutions for lowering overall TCO while improving application availability, protection, and service levels.

ILM building block

ILM has to do with the placement of data in the right location (based on performance, availability, and cost requirements) throughout its lifecycle, thereby improving the business value of information (policy). To make ILM a reality, organizations need to control where information resides within their application and storage array. Using Virtual LUN technology enables an effective LUN base data migration within a Symmetrix V-Max array, which provides a building block in developing an ILM strategy. This is accomplished through the nondisruptive migration of data across storage tiers and therefore reduces an enterprise's TCO, while still providing the performance and availability as required by the Oracle database and applications. Virtual LUNs accomplish this without any Oracle application downtime.

Tiered storage

Tiered storage is the process of maintaining storage of varying performance characteristics within the same array or across multiple arrays. Tiered storage gives administrators the flexibility to utilize their resources effectively by aligning high-end storage technology to the appropriate information value. For example, highly active data can be placed on high performing storage tiers, like enterprise Flash drives, or 15k rpm drives, providing improved performance and user experience. An important enabler for the correct use of storage tiers that match business needs is the ability to move data without disrupting the users using these applications. Symmetrix V-Max Virtual LUN technology enables migration of data between storage tiers within the same array. Virtual LUN operates at a LUN-level granularity and therefore can be utilized in migration of data based on filesystems, ASM disk groups, Oracle tablespaces, and any other data that is fully contained in a group of host devices.

Symmetrix V-Max available tiers:

- EFDs - Enterprise Flash drives provide maximum performance drives and are available in 200 GB and 400 GB. Ideal for Tier 0 applications. EFD can be utilized to consolidate data that is thinly spread across many hard disk drives (“short-stroking” technique) for performance reasons, providing latency-critical and high-I/O rate applications with a performance advantage, and benefits almost any workload
- HDDs - Hard disk drives are equipped to cope with all Tier 1 applications once layout is optimal. They are available in various sizes and rotational speed: 146 GB, 300 GB, and 450 GB 15k rpm, and 400 GB 10k rpm.
- SATA drives - Serial ATA drives are usually associated with archiving of data on Tier 2 or Tier 3.



Figure 2 Symmetrix V-Max tiered storage

Technology overview

The following technology descriptions provide a Symmetrix V-Max overview and explain how it can perform Oracle data movement, maximizing the correct use of storage tiers and establishing a core functionality in the development of an ILM strategy for the Oracle databases.

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 an 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 HDD, FICON, iSCSI and Gigabit Ethernet. Symmetrix V-Max Enginuity software provides many enhancements to the Symmetrix feature portfolio such as Symmetrix Remote Data Facility (SRDF), TimeFinder, and others. For example, SRDF enhances support to two-site DR solutions 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 HDD drives, and 200 / 400 GB enterprise Flash drives.




Figure 3 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 director engines, including two quad-core processors per module, or 16 per node
- Enhanced, scalable memory, providing a maximum 128 GB of global memory on each logic enclosure
- 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. Prior to Virtual LUN technology, most other LUN migration methods required an outage to the applications before or after the migration, commonly to add the new target LUNs to the host or to point the application to the newly migrated LUNs. Virtual LUN technology allows a user to migrate data based on LUN granularity within the array, providing a greater level of control over the system and higher levels of service for business applications, as no host-related changes are required and the migration is performed while the application is accessing the LUNs. Virtual LUN technology allows you to change storage device (host LUN) RAID protection and drive type.

Virtual LUN migrations take place within the storage array and therefore do not use host or SAN resources and are transparent to the host. [Figure 4](#) illustrates a Virtual LUN movement example.




Figure 4 Virtual LUN movement

Virtual LUN migration

RAID Virtual Architecture (RVA) is a new and optimized design that provides Enginuity with a single code path to manage all Symmetrix RAID protections. This new design allows, for the purposes of migration, two distinct RAID groups, of different types or on different storage tiers, to be associated with a Symmetrix device. 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 migration to unconfigured space, a target RAID group is created from the free pool in the array and migrated. When the migration is completed, the original RAID group is deleted and the storage returned to the free pool. Migration can also take place to a configured space. In this case the target devices have to match in size to the migrated devices, and not be used by any host (unmasked). At the end of the migration the source and target devices will have swapped their RAID protection and disk type. In addition, the storage where the data was migrated from will be formatted to prevent unwanted exposure of the original data.

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 resources 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, other back-end operations 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.

Autoprovisioning Groups

Solutions Enabler V7.0 introduces an easier, faster way to provision storage in the Symmetrix V-Max system. Autoprovisioning 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: Autoprovisioning 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.

EMC Symmetrix Management Console (SMC)

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.

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) databases.

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 Oracle 11g 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 (RAC). ASM is Oracle's recommended volume manager solution.

ASM disks are the storage devices provisioned to ASM disk groups. In this environment an ASM disk corresponds to a LUN in the Symmetrix storage array. The devices are seen by the operating system as block devices. Oracle recommends using Oracle ASMLib on Linux. ASMLib enables the Linux kernel to discover and access block devices used for the Oracle database. It provides device naming and permission persistency.

An ASM disk group is a collection of disks that ASM manages as a unit. The contents of the files stored in a disk group are evenly striped across all the disks in the disk group

ASM files are allocated from disk groups and behave as any other type of Oracle database file.

An ASM file is completely contained within a single disk group. ASM automatically generates ASM filenames when creating parameter files, controlfiles, redo logs, tempfiles, and datafiles.

Use cases for Oracle data relocation

Customers may have many different reasons for wanting to relocate their Oracle application data. Some EMC Oracle customers may want to upgrade to a newer Symmetrix storage platform, while others may want to move their data to better performing drives, which in turn increases transactions and lowers response times. Whether the enterprise is planning to move data to slower or faster drives within the array, they have the ability to do it nondisruptively with the Symmetrix V-Max platform using Virtual LUN technology. The following sections describe the use case environment and three use cases that illustrate the Symmetrix V-Max Oracle data movement. The use cases demonstrate methods of data movement using EMC technologies including SRDF and Virtual LUN. Additional data movement such as TimeFinder, Open Replicator, Open Migrator, and others are not covered in this use case.

Use case environment profile

This use case environment consisted of two EMC storage arrays, a source Symmetrix DMX and a target Symmetrix V-Max. A 1 TB two-node RAC database was created on the source Symmetrix DMX, distributed across 70 x 300 GB 15k rpm (RAID 1) devices, and migrated to the target Symmetrix V-Max. In the following use cases an OLTP TPC-C-like toolkit was used to generate load against the Oracle RAC database. While this load was running against the array, Virtual LUN technology was used to migrate the database components nondisruptively.

Use case 1 shows how both SMC and Solutions Enabler were used to carry out the SRDF migration. All Oracle database devices were placed under SRDF control throughout the migration process and some post-configuration changes were required once migrated to the Symmetrix V-Max.

In use case 2 Virtual LUN technology was used, while running load on the Oracle 11g RAC database, to migrate the underlying storage from RAID 1 to RAID 5. The Oracle database was fully available performing transactions while the migration was in progress.

The third use case describes how an Oracle ASM disk group was moved nondisruptively from HDDs to EFDs, again using Virtual LUN technology.

[Table 1](#) outlines the software and hardware components used.

Components	Description
1 x Symmetrix V-Max	Engenuity 5874 Cache: 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
2 x Dell R900	RHEL 5.1 CPU: 4 x Quad-core each HBA: 2 x Dual-port LP11002 each
EMC PowerPath®	5.1. SP2
Solutions Enabler Symmetrix Management Console (SMC)	7.0
Oracle	Enterprise Edition 11g Real Application Clusters (RAC) Automatic Storage Management (ASM)

Use case 1: Migration to the Symmetrix V-Max platform using EMC SRDF

Use case 1 infrastructure

For customers that are planning to upgrade to the Symmetrix V-Max array it is important to choose the right migration process thus minimizing application downtime by using proven EMC technologies. Leveraging the remote SRDF replica copy of the Oracle database can accelerate time to deployment.

SRDF significantly enhances remote replication and disaster recovery operations by using fast and reliable replication technology. EMC SRDF can guarantee online consistency through the use of SRDF Consistency Group technology. Because the Oracle database used devices on an existing Symmetrix platform, EMC SRDF was the ideal technology choice for migration.

For this use case a 1 TB Oracle two-node RAC database was migrated to the Symmetrix V-Max array using EMC SRDF. The SRDF setup was configured by dedicating SRDF directors on both arrays to control and synchronize the data between the local and remote arrays. All Oracle database devices were placed under SRDF control, ensuring Oracle database consistency through the migration process.

Migration method

The following steps were used to perform the database migration:

1. Symmetrix devices being used by the Oracle database on the source Symmetrix DMX and associated SRDF R2 devices on the target Symmetrix V-Max were configured, enabling them to take part in the SRDF migration process. A device file was created to include the SRDF devices. This device file contained local devices on the left, with their corresponding target devices on the right for all Oracle components:

```
devfile.txt
```

```
00A1 0103
```

```
00A2 0104
```

2. Source and target SRDF devices were allocated to both nodes in the Oracle RAC environment, with only one of the arrays accessible by the database at any one time.
3. The following `symrdf` command was issued to initiate synchronization between the source and target devices specified in the `devfile.txt`:

```
symrdf -f devfile.txt establish -full
```

Note: The database was fully operational and all ongoing transactions were serviced from the R1 devices throughout the synchronization process

4. Once all data was fully synchronized the RAC database resources were stopped to guarantee consistency thus facilitating database data integrity across all R2 devices.
5. Once the database was shut down an SRDF query with the following command was issued to ensure all data successfully migrated:

```
symrdf -f devfile.txt query
```

6. The SRDF relationship on the `devfile.txt` was suspended and all R2 devices were converted from a write/disabled state to a read/write state using Solutions Enabler:

```
symrdf -f devfile.txt suspend
```

-
7. At this stage both the OCR and voting disks need the existing PowerPath pseudo name to be matched to the R2 devices. This provides the correct device location for the Oracle cluster registry and voting disk configuration. The OCR and voting disks are the ONLY devices that need their PowerPath pseudo names to be renamed to match the original production instance.

Note: In some cases customers may need to re-create the OCR and voting disks. In the case where customers are migrating both array and servers then the OCR and voting disks should be re-created. Refer to Oracle documentation for steps on how to re-create the OCR and voting disks.

The steps to rename these PowerPath pseudo name devices are:

- Identify the PowerPath pseudo name associated with the voting disk:

```
[root@tce-flash01 bin]# crsctl query css votedisk
```

```
0. 0 /dev/emcpower2
```

```
Located 1 voting disk(s).
```

- Identify the status and pseudo name of the PowerPath device associated with the OCR device:

```
[root@tce-flash01 bin]# ./ocrcheck
```

```
Status of Oracle Cluster Registry is as follows :
```

```
Version          :          2
Total space (kbytes) :      293592
Used space (kbytes)  :         3812
Available space (kbytes) :    289780
ID                : 1807044143
Device/File Name   : /dev/emcpower1
                   Device/File integrity check succeeded
                   Cluster registry integrity check succeeded
```

- List the used pseudo devices names. Note that emcpower20 has replaced emcpower7 in the output:

```
emcpadm getusedpseudos
```

- View the powermt display for the emcpower20 device:

```
powermt display dev=20
```

- Update the powermt.custom file with the new pseudo device mappings:

```
powermt save
```

8. Before starting Oracle Services check the following on both RAC nodes:

- Ensure that all R2 devices are split from their R1 devices, ensuring database consistency across all devices.(symrdf -f vmax.txt query).
- The ASM **init.ora** file parameter ASM_DISKSTRING includes the path to the R2 devices (asm_diskstring='ORCL:*').
- The ASM **init.ora** file parameter ASM_DISKGROUPS contains the names of all the disk groups (asm_diskgroups='DATA','FRA','REDO','TEMP','UNDO').

9. Start up the Oracle Clusterware, ASM, and database processes on both RAC nodes:

```
[root@tce-flash01 ~]# cd /u01/crs/oracle/11g/bin
[root@tce-flash01 bin]# ./crsctl start crs
Attempting to start Oracle Clusterware stack
CRS stack will be started shortly
[root@tce-flash02 bin]# ./crs_stat -t
```

ora.flash.db	application	ONLINE	ONLINE	tce-flash02
ora....h1.inst	application	ONLINE	ONLINE	tce-flash01
ora....h2.inst	application	ONLINE	ONLINE	tce-flash02
ora....SM1.asm	application	ONLINE	ONLINE	tce-flash01
ora....01.lsnr	application	ONLINE	ONLINE	tce-flash01
ora....h01.gsd	application	ONLINE	ONLINE	tce-flash01
ora....h01.ons	application	ONLINE	ONLINE	tce-flash01
ora....h01.vip	application	ONLINE	ONLINE	tce-flash01
ora....SM2.asm	application	ONLINE	ONLINE	tce-flash02
ora....02.lsnr	application	ONLINE	ONLINE	tce-flash02
ora....h02.gsd	application	ONLINE	ONLINE	tce-flash02
ora....h02.ons	application	ONLINE	ONLINE	tce-flash02
ora....h02.vip	application	ONLINE	ONLINE	tce-flash02

Migration advantages

Both DMX and Symmetrix V-Max storage arrays were presented to both nodes of the Oracle RAC but only one array's LUNs were accessible and usable at any one time. Once the data transfer between the arrays over SRDF completed it was necessary to suspend the SRDF link between the arrays. The devices originally presented from the DMX to the two-node RAC were removed to allow visibility of the new Symmetrix V-Max devices. These R2 devices were write-enabled, allowing the hosts access to the newly presented devices instantly.

Use case 2: Nondisruptive movement of an entire database from RAID 1 to RAID 5 using EMC Virtual LUN

As more and more information becomes mission-critical, users demand easier ways to migrate information without disrupting their systems. Some customers are reluctant to move their application data in fear of disturbing their system or creating application downtime. EMC's Virtual LUN technology, as described below, allows customers to move Oracle data nondisruptively within the Symmetrix V-Max array.

EMC Virtual LUN method

This use case details the process using Virtual LUN to migrate a 1 TB database created across a two-node-RAC from RAID 1 storage to RAID 5 within the same array. During this testing the Oracle database remained up and running, with no DBA or system administration actions required. The source devices' server LUN address, as well as the EMC PowerPath pseudo name, remained consistent from the server perspective.

As already stated the Oracle database continued to operate while the data migration was happening within the storage array. The following procedure provides the steps required to perform the migration from RAID 1 LUNs to RAID 5 LUNs within the Symmetrix V-Max array

1. To set up, validate, and establish the migration, the new LUNs and target LUNs were identified and added to a device file. The device file is a basic file with the source on the left and the target on the right, such as in the example below. The specific device file used in this use case consisted of 210 source and 210 target LUNs.

```
095 185
096 186
.....
```

2. Once the device file has been created then the migration of the devices can begin using the commands listed below.

The **`symmmigrate validate`** command is an optional command, that can be used in two different ways. The command can be used as a simple positive/negative test to determine if a migration is viable. This will be denoted by the return value of the command. The command can also be used when migrating to configured space to output a file that contains the exact pairs of devices that will be used for the migration. This output can be used to verify that the command will produce the exact migration that is desired by the user and to modify the resulting file if it is not.

Example:

```
symmmigrate validate -sid 209 -f devfile.txt -name oratce7
```

3. The **`symmmigrate establish`** command is used to create the migration session and begin the migration of the source devices to the new target location. The target protection type and disk group number must be provided by the user (unless using a device pair file). This is the point in the migration process that the mirrors are moved to the source device from the target, if using configured space, or are created if using unconfigured space.

Example:

```
symmmigrate establish -sid 209 -f devfile.txt -name oratce7
```

- Use the **symmigrate query** command to display the status of a specific migration session, or all migration sessions. The query command can be used with the session name only, or you can additionally filter the request by adding the specific filename, device group name, or storage group name that was used with that session. If the additional (file/storage group/device group) name is added to the session name, the Symmetrix ID is also required.

```
symmigrate query -name oratce7 -sid 209 -detail
```

EXAMPLE OUTPUT:

Src	Tgt	Invalid Tracks	Status SRC => TGT	Done (%)	TGT Dsk Grp	TGT Protection	Session Name
0095	0185	0	CreateInProg	N/A	01 RAID-5 (3+1)	oratce7	
0096	0186	0	CreateInProg	N/A	01 RAID-5 (3+1)	oratce7	

- Once the CreateInProg step completes it will follow on and enter the SyncInProg state as shown below.

Src	Tgt	Invalid Tracks	Status SRC => TGT	Done (%)	TGT Dsk Grp	TGT Protection	Session Name
0095	0185	516600	SyncInProg	1	01 RAID-5 (3+1)	oratce7	
0096	0186	518865	SyncInProg	1	01 RAID-5 (3+1)	oratce7	

- Once fully sync'd up a migrated status will appear against the device pairs. You will also be able to see a 100% completed on this query command. Here is an example output from a **query** command that shows that the migration has completed and all data is now running on the target devices.

Src	Tgt	Invalid Tracks	Status SRC => TGT	Done (%)	TGT Dsk Grp	TGT Protection	Session Name
0095	0185	0	Migrated	100	01 RAID-5 (3+1)	oratce7	
0096	0186	0	Migrated	100	01 RAID-5 (3+1)	oratce7	

```
symmigrate verify -name oratce7 -sid 209
All of the devices are in the 'Migrated' state.
```

Once the verify returns with a message, see above, a terminate operation can be carried out to break the relationship and complete the Virtual LUN migration process. The commands above were issued via Solutions Enabler but the same functionality is also available using SMC (Symmetrix Management Console).

EMC Virtual LUN advantage

During this testing the Oracle database remained up and running, with no DBA or system administration actions required. Once completed, the Virtual LUN migration returned the original source devices to the free pool, which could be used for another application or function. As more and more customers' information becomes mission-critical, users demand easier ways to migrate information without disrupting their systems. EMC's Virtual LUN technology, as described above, allows customers to move Oracle data nondisruptively within the Symmetrix V-Max array.

Use case 3: Nondisruptive movement of an Oracle ASM disk group from HDD to EFDs using EMC Virtual LUN

When customers are deploying an ILM strategy with their large Oracle databases multiple ASM disk groups are recommended. Each disk group can represent a different tier of storage. Over time, however, the type of disk that is supporting the workload of a particular disk group may not be sufficient to achieve the appropriate performance results necessary for the Oracle application. Using Symmetrix V-Max Virtual LUN technology an individual ASM disk group can be moved nondisruptively from one type of performance drive to a much faster drive type. Once the disk group has been moved then the application performance can be significantly improved. In the example below the entire DATA disk group of an Oracle OLTP database was moved from 300 GB 15k HDD drives to 400 GB EFDs.

Additionally we were able to migrate from 70 HDD drives to eight EFDs. Not only was performance increased at the database level but the 70 HDD drives were removed and returned for use with other applications. 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.

Before moving the data in an ASM disk group utilizing Virtual LUN, several tasks need to be done before the movement occurs. The process can be broken down into five steps. Listed below are the steps required to move a disk group using Symmetrix V-Max Virtual LUN technology. A critical step in this use case is having the ability to identify the devices associated with the ASM disk group that will be moved. Below are the SQL statements and scripts that were used to identify the disks in the DATA disk group for the database.

Oracle ASM disk group movement via EMC Virtual LUN method

1. Determine an ASM disk group to move and determine disks within the disk group.

Obtain all the ASM disks in the disk group:

```
SQL> select a.name gname
       ,b.name dname
       from v$asm_diskgroup a
       ,v$asm_disk b
       where a.name='DATA'
       and a.group_number=b.group_number;
```

```
ASM Disk Group ASM Disk
-----
DATA           DATA_001
DATA           DATA_002
DATA           DATA_003
DATA           DATA_004
DATA           DATA_005
DATA           DATA_006
DATA           DATA_007
DATA           DATA_008
DATA           DATA_009
DATA           DATA_010
```

2. Create a text file containing the ASM disks and run the following script to determine the Symmetrix device names.

asm.sh - script used for obtaining symmetrix device names

```
asm.sh
#!/bin/bash

for ASMDISK in `cat asmdisk.lst`
do
    major=`ls -l /dev/oracleasm/disks/${ASMDISK} | awk '{print $5}' | tr -d ', '`
    minor=`ls -l /dev/oracleasm/disks/${ASMDISK} | awk '{print $6}' | tr -s " "`
    partition=`cat /proc/partitions | tr -s " " | grep -we "${major} ${minor}" | awk
' {print $4}'`
    disk=`echo $partition | tr -d [:digit:]`
```

```
symmid=`/sbin/powermt display dev=$disk | grep Logical | cut -f2 -d="`
echo $ASMDISK /dev/$partition $symmid
```

done

execute script asm.sh

```
# asm.sh

DATA_001 /dev/emcpowertb1 0095
DATA_002 /dev/emcpowerta1 0096
DATA_003 /dev/emcpowersz1 0097
DATA_004 /dev/emcpowersy1 0098
DATA_005 /dev/emcpowersx1 0099
DATA_006 /dev/emcpowersw1 009A
DATA_007 /dev/emcpowersv1 009B
DATA_008 /dev/emcpowersul 009C
DATA_009 /dev/emcpowerst1 009D
DATA_010 /dev/emcpowerss1 009E
.....
```

After running the above script all of the 70 devices have been identified.

Once the devices have been identified then the process of moving the entire disk group can be done using the Virtual LUN process.

- Once the EMC devices in the DATA disk group have been identified then a device file can be created.

The device file is a very basic file with just the source on the left and target on the right, such as in the example below:

```
095      395
096      396
.....
```

- Once the device file has been created then the migration of the devices can begin using the commands listed below:

```
symmmigrate validate -sid 209 -f datatoEFD.txt -name oratce8
```

This command validates the migration you are attempting to carry out and ensures the device file created has the correct syntax.

- Once the validation commands return to the command prompt then an **establish** command is used to start the migration from the source LUN to the target LUN:

symmmigrate establish -sid 209 -f datatoEFD.txt -name oratce8

Different status are returned during the establish process as seen below. Each time the **query** command is run you can view the different status options that are shown:

symmmigrate query -name oratce8 -sid 209 -detail

Src	Tgt	Invalid Tracks	Status SRC => TGT	Done (%)	TGT Dsk Grp	TGT Protection	Session Name
0095	0395	0	CreateInProg	N/A	04 RAID-5 (7+1)	oratce8	
0096	0396	0	CreateInProg	N/A	04 RAID-5 (7+1)	oratce8	

```
.....
```

Src	Tgt	Invalid Tracks	Status SRC => TGT	Done (%)	TGT Dsk Grp	TGT Protection	Session Name
0095	0395	491658	SyncInProg	6	04 RAID-5 (7+1)	oratce8	
0096	0396	486864	SyncInProg	7	04 RAID-5 (7+1)	oratce8	

```
.....
```

```

Invalid   Status   Done   TGT
Src  Tgt   Tracks SRC => TGT  (%) Dsk Grp TGT Protection  Session Name
-----
0095 0395         0 Synchronized 100      04 RAID-5 (7+1)  oratce8
0096 0396         0 Synchronized 100      04 RAID-5 (7+1)  oratce8

```

```

Invalid   Status   Done   TGT
Src  Tgt   Tracks SRC => TGT  (%) Dsk Grp TGT Protection  Session Name
-----
0095 0395         0 MigrInProg 100      04 RAID-5 (7+1)  oratce8
0096 0396         0 MigrInProg 100      04 RAID-5 (7+1)  oratce8

```

```

Invalid   Status   Done   TGT
Src  Tgt   Tracks SRC => TGT  (%) Dsk Grp TGT Protection  Session Name
-----
0095 0395         0 Migrated 100      04 RAID-5 (7+1)  oratce8
0096 0396         0 Migrated 100      04 RAID-5 (7+1)  oratce8

```

Once the status of "migrated" is received, as above, then a verify is necessary to ensure the migration process has completed on the back-end in the array itself.

```
symmmigrate verify -sid 209 -name oratce
```

All session(s) with name 'oratce' are in 'Migrated' state.

EMC Virtual LUN advantage

The migration is now fully completed and no interruption or manual operation from the customer's application perspective was necessary. The above commands were issued via Solutions Enabler but the same functionality is also available using SMC (Symmetrix Management Console).

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

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 today need to create a competitive advantage for their IT infrastructure by placing increasing value on both the location and availability of their Oracle 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.

Therefore, investing in V-Max technology enables an enterprise's IT/Oracle infrastructure to succeed in this demanding 24x7 world with the following success factors:

- **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, across different storage tiers to ensure the highest availability service levels to the customer, without disruption 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/Scalability:** 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:** Virtual LUN technology enables ease-of-use management of new storage tiers or RAID protections as a building block in an enterprise's ILM strategy. With this nondisruptive method of data movement using SMC and Solutions Enabler, both DBAs and storage administrators can monitor the data movement progress visually and, once completed, it is transparent where their Oracle data resides on the storage array. This is a far more effective solution for data movement online and avoids any extra cycles necessary to complete an offline migration. This creates operational efficiency within the enterprise's IT/Oracle infrastructure. Those enterprise people resources (storage and Oracle administrators) can now be deployed on other pressing enterprise storage and Oracle issues.