

EMC Virtual Infrastructure for Oracle

Enabled by EMC Symmetrix V-Max and
VMware vSphere 4

Reference Architecture

EMC Global Solutions



Copyright © 2009 EMC Corporation. All rights reserved.

Published November, 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: H6504

Table of Contents

Reference architecture overview	4
Key components	7
Symmetrix V-Max.....	7
EMC ControlCenter.....	9
EMC Symmetrix Management Console	9
EMC Storage Viewer for vSphere Client.....	10
PowerPath/VE.....	11
VMware vSphere.....	12
Oracle Database 11g	14
Storage design and layout	15
Storage connectivity.....	16
Physical architecture.....	17
Validated environment profile	17
Hardware and software resources.....	18
Testing and validation.....	19
Conclusion.....	25

Reference architecture overview

Document purpose This document describes the reference architecture of the EMC Virtual Infrastructure for Oracle enabled by EMC Symmetrix V-Max™ and VMware vSphere 4 solution.

Solution purpose Virtualization of servers and storage can provide reduced operating and capital expense for customers who are trying to control data center expansion and reduce their environmental footprint. This document illustrates that, by using VMware's virtualization and EMC® Symmetrix® V-Max storage technologies, applications can be decoupled from the underlying hardware and infrastructure. The total amount of hardware can be consolidated and fully utilized to provide a more cost-effective and energy-efficient solution. This solution demonstrates the capabilities of vSphere in conjunction with the Symmetrix V-Max platform to show how consolidation can be fully realized.

EMC Proven™ Solutions are designed to build a replica of some aspects of a customer environment. Information about the creation of the environment and its operation is provided to field personnel and development organizations in the related Proven Solution Guides.

The purpose of this solution is to show how to build a customer environment consisting of Enterprise Oracle 11g databases on a Symmetrix V-Max and VMware vSphere.

The business challenge It is critical that today's data centers utilize resources more efficiently to support the dynamic requirements of running multiple business-critical applications. Applications must now be accessible 24 x 7 to meet the growing global access needs of organizations. At the same time, there are significant pressures to drive down IT costs and simplify management of the data center. A growing number of organizations are turning to server and storage virtualization technologies to dramatically cut costs. Tremendous opportunity exists to reduce operational and capital expenses that result from information growth and unused resources.

Organizations want to accelerate virtualization to include mission-critical applications, but not at the expense of service level agreements (SLAs), performance, or availability. In addition to keeping data center operations running continuously, IT must reduce complexity by deploying a solution that simplifies management.

There are three main drivers for virtualization based on cost savings:

- **Server consolidation** – reduce the number of servers. For the most part, servers are under-utilized. Using virtualization can consolidate servers into a pool of computing resources - virtual servers - which can then be optimally sized for various tasks. A ratio of four or five to one (virtual servers to physical servers) can typically be achieved. This saves on the number of servers, which means less air conditioning, power, physical space, and fewer IT administrators, leading to overall savings in infrastructure and operations costs.
- **Ease of application deployment** - virtualization promises real cost savings and agility to customers. Through virtualization, a company can rapidly deploy applications where the underlying technology components can expand and

contract with the flow of the business life cycle. Traditionally, once an application was deployed it was bound to a particular infrastructure, until the infrastructure was upgraded. The result was low efficiency, poor utilization, and inflexibility. Virtualization allows applications to be dynamically deployed onto the most suitable infrastructure at run time. Therefore, it decouples the application from the hardware infrastructure.

- **Reducing data center downtime** - Data center downtime is generally categorized as either “planned” or “unplanned.” The ratio (and therefore cost) of unplanned downtime to planned downtime is usually estimated as four to one.

Among the various reasons for planned downtime, hardware maintenance is the most common. Regardless of the cause of downtime, when mission-critical applications are not accessible the business experiences significant negative impact, including loss of productivity, revenues, and customer satisfaction. Any effective solution that can eliminate or significantly reduce planned downtime offers real operational benefits for system administrators. Virtualization technologies can provide the ability to avoid data center downtime or outages associated with physical infrastructure.

IT departments and infrastructure providers are under increasing pressure to provide computing infrastructures at the lowest possible cost (operating and capital). In order to do this, the concepts of resource pooling, virtualization, dynamic provisioning, utility, and commodity computing must be leveraged to meet these needs.

This solution demonstrates how vSphere, in conjunction with the Symmetrix V-Max platform, using consolidation and virtualization, can reduce operating and capital expenditure, as follows:

- **Operating Expense (OPEX)** - accelerated methods of server and storage provisioning with VMware and Symmetrix V-Max technologies
 - Administration overhead reduction
 - Less resources (DBAs / SAs) to support the OS and database
 - Improved database availability via VMware HA
 - Better storage efficiency with migration of storage nondisruptively using Symmetrix V-Max VLUN technology
 - Increased server efficiency using load balancing across the cluster
 - Data center overheads – cooling/electrical, etc.
 - Physical servers removed from data center
 - **Capital Expense (CAPEX)** – reduced infrastructure expenses
 - Virtual server density. Many virtual machines to one physical server
 - Movement of storage between tiers, based on cost / performance using Symmetrix V-Max VLUN technology
 - Reduction in the rate of physical database purchase
-

The technology solution

This solution describes the virtualized deployment of Oracle 11g databases on a VMware vSphere environment consisting of two ESX 4 servers. The ESX servers were configured in a VMware HA cluster managed by VMware vCenter Server residing on the EMC Symmetrix V-Max storage system. VMware vCenter Server was installed on a virtual machine (VM) and can be configured to use an embedded SQL Server database for small environments, with SQL Server or an Oracle database as the primary repository.

The following table describes the key components and their configuration details used within this environment.

Quantity	Component	Configuration	Software
2	VMware vSphere™ 4	2 x Xeon X5570 processor, 2.93 GHz, 8 MB 144 GB RAM with two 146 GB SAS internal drives	ESX 4.0
1	Symmetrix V-Max	64 GB mirrored cache 80 x 300 GB 15k 20 x 1 TB SATA	Enginuity™ 5874
2	Oracle 11g OLTP database system	300 GB OLTP	Oracle 11.1.0.7
1	Oracle 11g OLTP database system	150 GB OLTP	Oracle 11.1.0.7
1	EMC ControlCenter® Server	VMware virtual machine	EMC ControlCenter 6.1
1	EMC Symmetrix® Management Console	VMware virtual machine	EMC Symmetrix Management Console 7.0.1.3
1	VMware vCenter Server	VMware virtual machine	Windows 2003 Server
1	DHCP / DNS	VMware virtual machine	Windows 2003 Server

Key components

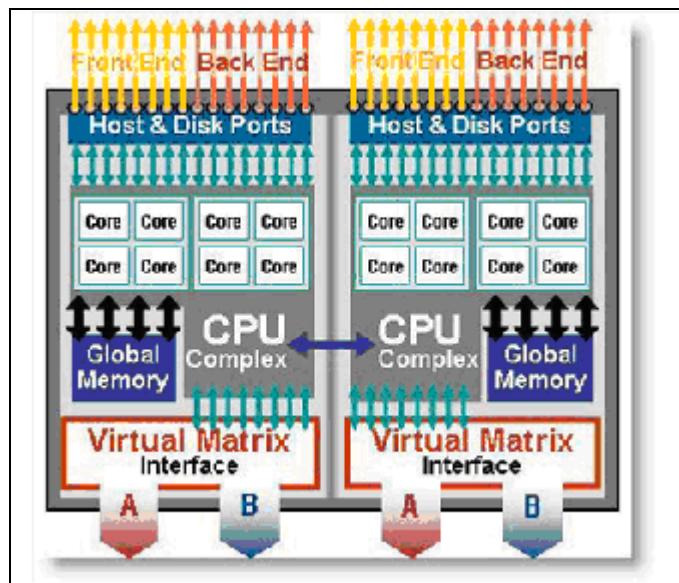
Introduction

This section briefly describes the key solutions components. For details on all of the components that make up the reference architecture, refer to the hardware and software sections.

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, or engine, configuration to an eight-engine configuration and with a maximum of 10 storage bays. Online system upgrades are achieved by adding HA engines. Each HA engine contains two integrated director boards with multi-core CPU processing power, cache memory, front-end ports, and back-end ports.

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.



The image above shows the components that make up a single V-Max Engine or HA node. Current configurations support up to eight V-Max Engines to be configured in one array.

Auto-provisioning Groups

Solutions Enabler version 7.0 introduces Auto-provisioning Groups, an easier, faster way to provision storage in the Symmetrix V-Max system.


Auto-provisioning provides the storage administrator with a method to logically group resources (host initiators (HBAs), front-end storage array ports, and logical devices (LUNs)) and to use these grouped resources to greatly simplify storage management operations. In addition to simplifying storage management operations, auto-provisioning greatly reduces the amount of time required to perform provisioning operations.

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.

Virtual LUN

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. A virtual LUN migration takes place within the storage array and therefore does not use host or SAN resources to migrate data from one tier of storage to another. All data migrations are completely transparent and non-disruptive to the host, application, and database.

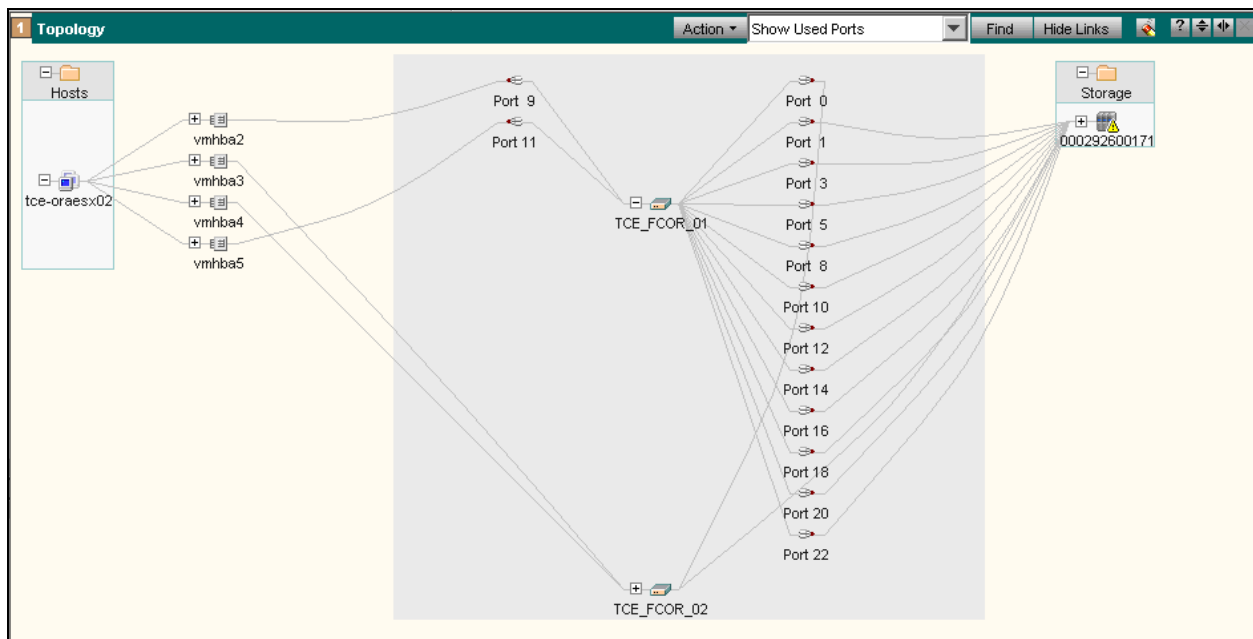
The following image illustrates the ability of Symmetrix V-Max Virtual LUN technology to move data to where it is needed, without users being impacted via nondisruptive migration.



EMC ControlCenter

EMC ControlCenter (ECC) is designed to be the centralized control for an entire distributed storage environment. It is a powerful, flexible, unified framework and suite of tools that provides end-to-end management of storage networks, storage devices, and other storage resources.

The following image from EMC ControlCenter shows the SAN mapping from a VMware ESX server through a Fibre Channel switch to the Symmetrix V-Max storage array.

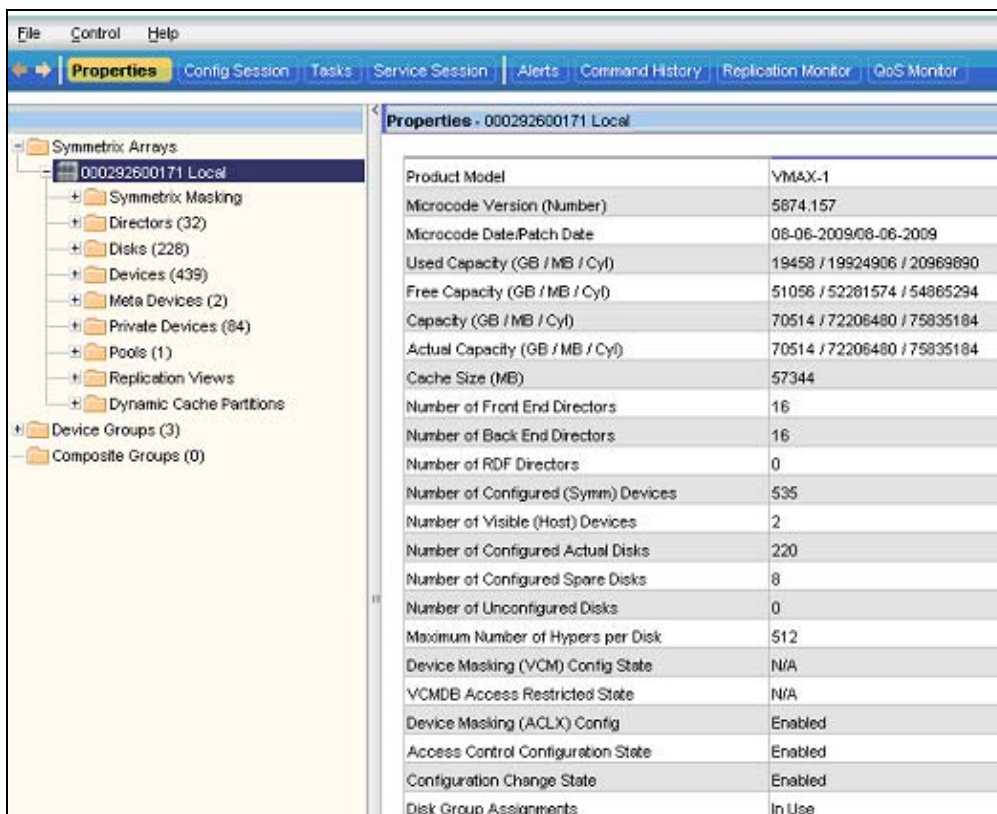


EMC Symmetrix Management Console

Symmetrix Management Console (SMC) is an intuitive, browser-based user interface that configures and manages Symmetrix system arrays. SMC presents the functionality of the Solutions Enabler SYMCLI (command line interface) in a browser interface.

SMC is preinstalled on the Symmetrix V-Max system service processor, and can also run on a data center host.

The following image shows the Symmetrix Management Console.



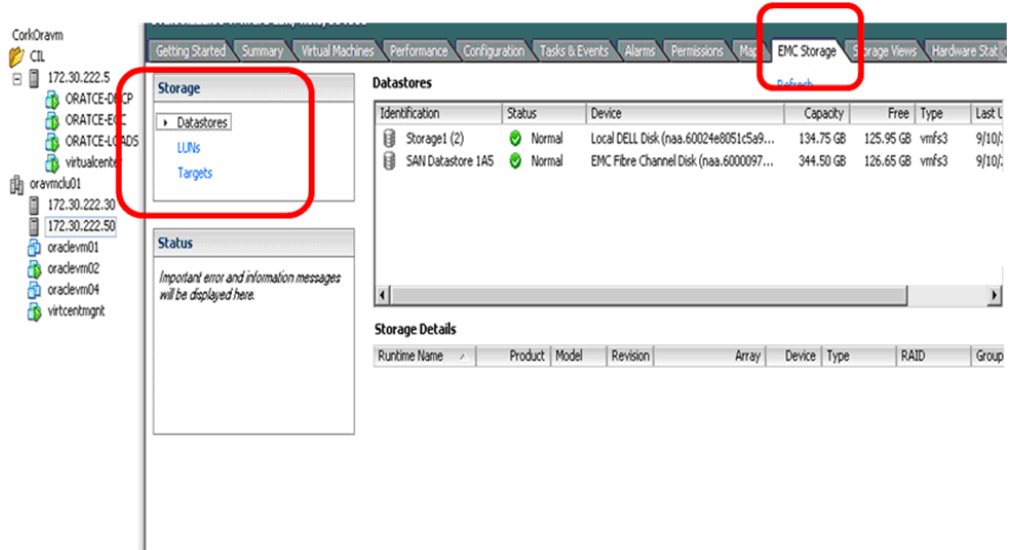
EMC Storage Viewer for vSphere Client

The EMC Storage Viewer for vCenter provides the vCenter administrator with the tools to both enable understanding of the underlying storage infrastructure as well as a common language with which to speak to the storage administrators. The goal is to facilitate communications and to enable faster and more precise coordination of efforts should issues arise in the environment.

The Storage Viewer for vSphere Client brings the storage details up to the virtual infrastructure administrator, merging all the functionality of several different applications that perform storage mapping-related functions into a single application that integrates seamlessly with the vSphere Client.

Using a plug-in to the vSphere Client interface, the EMC Storage Viewer for vSphere Client enables users to resolve the underlying storage of Virtual Machine File System (VMFS) datastores and virtual disks, as well as raw device mappings (RDM). In addition, the user is presented with lists of storage arrays and devices that are accessible to the ESX servers in the virtual infrastructure.

The following image illustrates the integration of the EMC Storage Viewer from a VMware standpoint. Following installation, a new tab, EMC Storage, is seen on the vSphere Client. This feature provides an easy-to-use method that enables vCenter administrators to map VMware datastores to the logical volumes presented by the EMC Symmetrix V-Max storage array.



PowerPath/VE

EMC PowerPath[®]/VE is server-resident software that enhances performance and application availability. PowerPath/VE works with the storage system to intelligently manage I/O paths, and supports multiple paths to a logical device. In this solution, PowerPath provides:

- Automatic failover in the event of a hardware failure. PowerPath/VE automatically detects path failure and redirects I/O to another path.
- Dynamic multipath load balancing. PowerPath/VE intelligently 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.

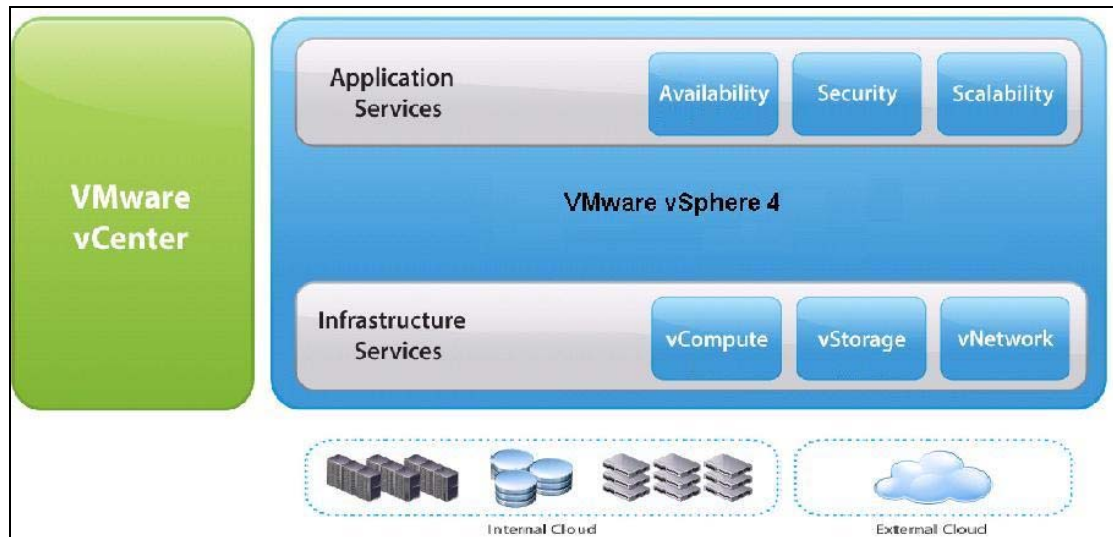
PowerPath/VE enables customers to standardize on a single multipathing solution across their entire environment. For customers standardizing on VMware, this means access to flexible and automatic I/O load balancing to manage the complexity of virtual machines and I/O-intensive applications in hyper-consolidated environments.

VMware vSphere

vSphere 4

VMware vSphere™ leverages the power of virtualization to transform data centers into simplified cloud computing infrastructures. VMware vSphere virtualizes and aggregates the underlying physical hardware resources across multiple systems and provides pools of virtual resources to the data center.

As a cloud operating system, VMware vSphere manages large collections of infrastructure (such as CPUs, storage, and networking) as a seamless and dynamic operating environment, and also manages the complexity of a data center.

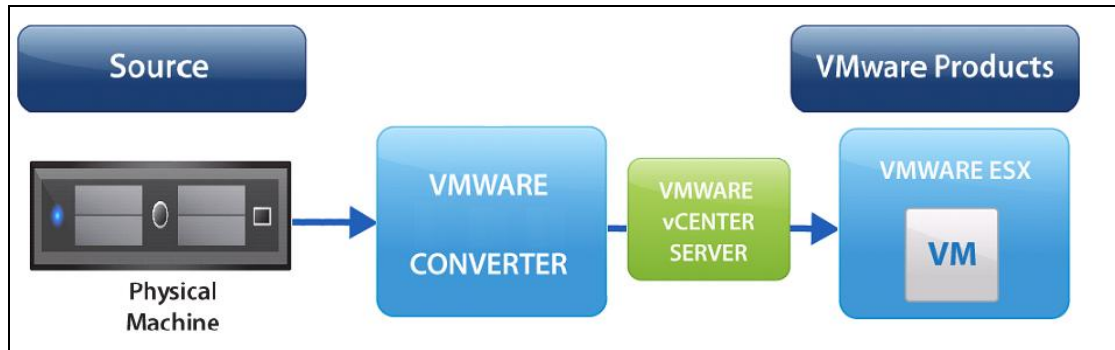


VMware vCenter Server

VMware vCenter Server provides a single point of control of the data center. It provides essential data center services such as access control, performance monitoring, and configuration. It unifies the resources from the individual computing servers to be shared among virtual machines in the entire data center. It does this by managing the assignment of virtual machines to the computing servers and the assignment of resources to the virtual machines within a given computing server based on the policies that the system administrator sets.

VMware Converter

VMware vCenter Converter is an add-on plug-in for vCenter Server 4.0, and provides an integrated solution for migrating both physical and virtual machines to VMware vSphere. vCenter Converter can be used to create, edit, and monitor physical-to-virtual and virtual-to-virtual migration tasks through the vCenter Server user interface. In addition, you can use vCenter Converter to directly import supported third-party images to virtual machines that vCenter Server manages.



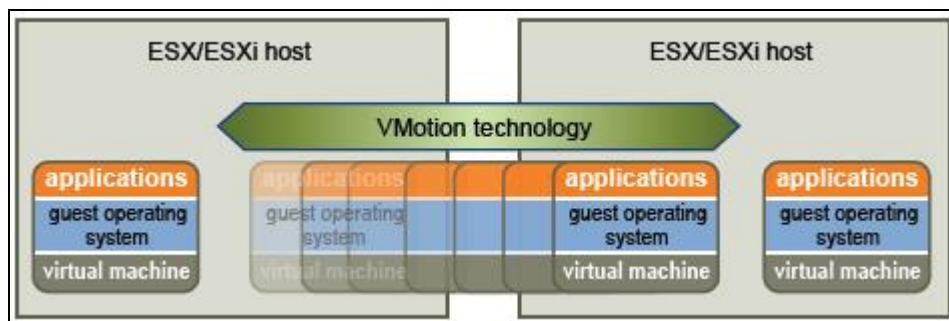
Managing VMware ESX Server disk access

VMware ESX Server provides two choices for managing disk access in a virtual machine, VMFS and RDM:

- **Virtual Machine File Storage (VMFS).** VMFS is a distributed file system that allows concurrent access by multiple hosts to files on a shared VMFS volume. VMFS offers high I/O capabilities for virtual machines. It is optimized for storing and accessing large files such as virtual disks.
- **Raw device mapping (RDM).** RDM is a mapping file in a VMFS volume that acts as a proxy for a raw physical device. The RDM file contains metadata used to manage and redirect disk accesses to the physical device. This technique provides the advantage of direct access to the physical device in addition to some of the advantages of a virtual disk on VMFS storage.

VMware VMotion

VMotion enables the migration of running virtual machines from one physical server to another without service interruption, tailored to the resource requirements of the applications. With VMotion, resources can be dynamically reallocated to virtual machines across physical servers.



VMware High Availability

VMware HA enables quick restart of virtual machines on a different physical server within a cluster automatically if a host fails. All applications within the virtual machines have the benefit of high availability, through application clustering. HA monitors all physical hosts in a cluster and detects host failures. An agent placed on each physical host maintains a heartbeat with the other hosts in the resource pool. Loss of a heartbeat initiates the process of restarting all affected virtual machines on other hosts. HA ensures that sufficient resources are available in the cluster at all times to restart virtual machines on different physical hosts in the event of host failure.

Oracle Database 11g**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 Enterprise Edition comes with a wide range of options to help grow your business and meet users' performance, security, and availability service level expectations.

Oracle ASM

Oracle Automatic Storage Management (ASM) is an integrated database filesystem and disk manager. ASM filesystem and volume management capabilities are built into the Oracle database kernel. And it can reduce 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 database files across an ASM diskgroup after disks have been added or removed.

ASM CMD

The **asmcmd** utility can be used by Oracle database administrators to query and manage their ASM systems. ASM-related information can be retrieved easily for diagnosing and debugging purposes.

Below is an example of the **asmcmd** command **lsdgm**, which can be used to display the ASM diskgroup managed by the ASM including the name, type, size, and so on.

```

ASMCMD> lsdgm
State Type Rebal Sector Block AV Total_MB Free_MB Req_mir_free_MB Usable_file_MB Offline_disks Name
MOUNTED EXTERN N 512 4096 1048576 589712 228105 0 228105 0 DATA/
MOUNTED EXTERN N 512 4096 1048576 98302 98250 0 98250 0 FRA/
MOUNTED EXTERN N 512 4096 1048576 98302 94579 0 94579 0 REDO/
MOUNTED EXTERN N 512 4096 1048576 98302 65480 0 65480 0 TEMP/

```

Swingbench

Swingbench is a publicly available load generator (and benchmark tool) designed to stress test Oracle databases. Swingbench consists of a load generator, a coordinator, and a cluster overview. The software enables a load to be generated and the transactions/response times to be charted.

Swingbench is provided with four benchmarks:

- OrderEntry – TPC-C-like workload.
- Calling Circle – Telco-based self-service workload.
- Stress Test – Performs simple insert/update/delete/select operations.
- DSS – A DSS workload, based on the Oracle Sales History schema.

The Swingbench workload used in case testing was OrderEntry. The Order Entry (PL/SQL) workload models the classic order entry stress test. It has a profile similar to the TPC-C benchmark. It models an online order entry system, with users being required to log on before purchasing goods.

Storage design and layout

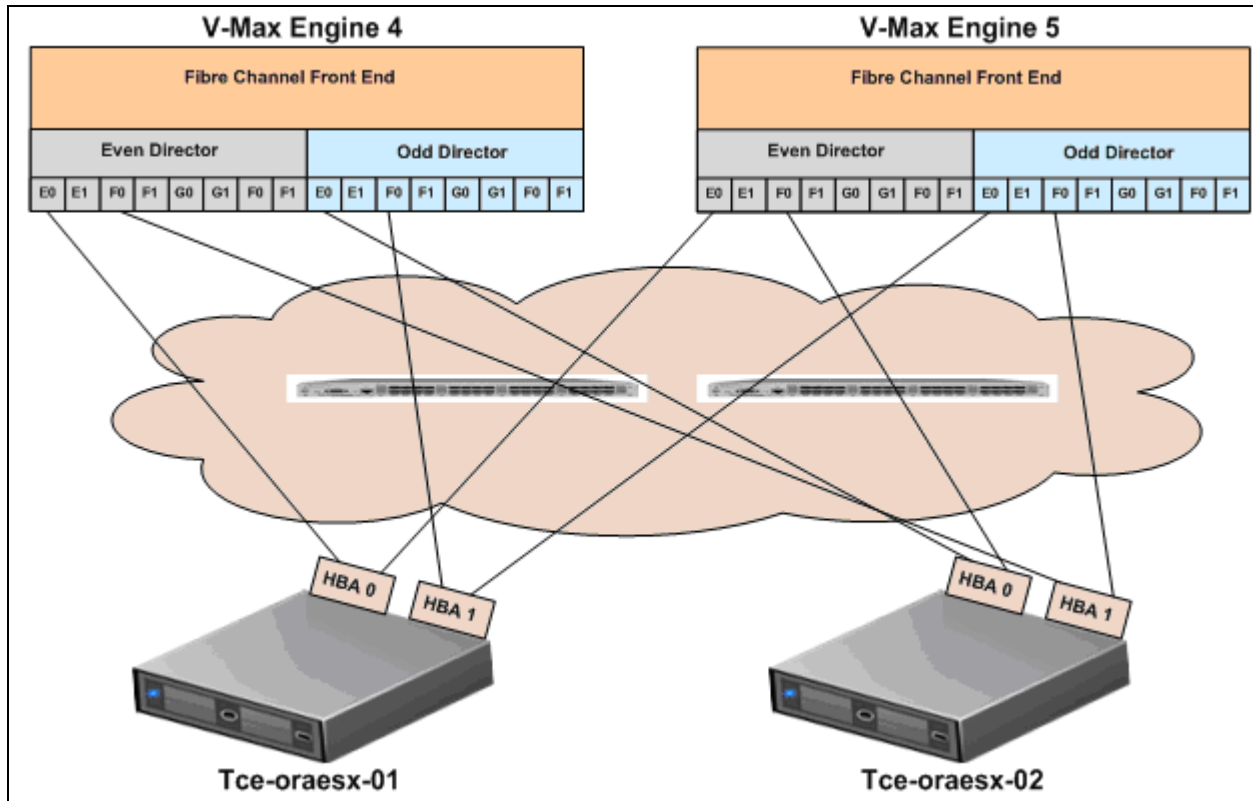
Three Oracle databases were used in this solution, all residing on Linux virtual machines. The databases sizes used were:

- 2 x 300 GB
- 1 x 150 GB

When building this solution the primary goal was to concentrate on functionality and the integration of Oracle 11g on VMware running on the Symmetrix V-Max. As RDM and VMFS were being tested, the same sized devices were created for both databases. All devices (LUNs) created on the Symmetrix V-Max were configured with RAID 1 protection, which offers the best possible performance.

Storage connectivity

The following image illustrates the physical topology connectivity showing protection (no single point of failure) and load balancing for performance.

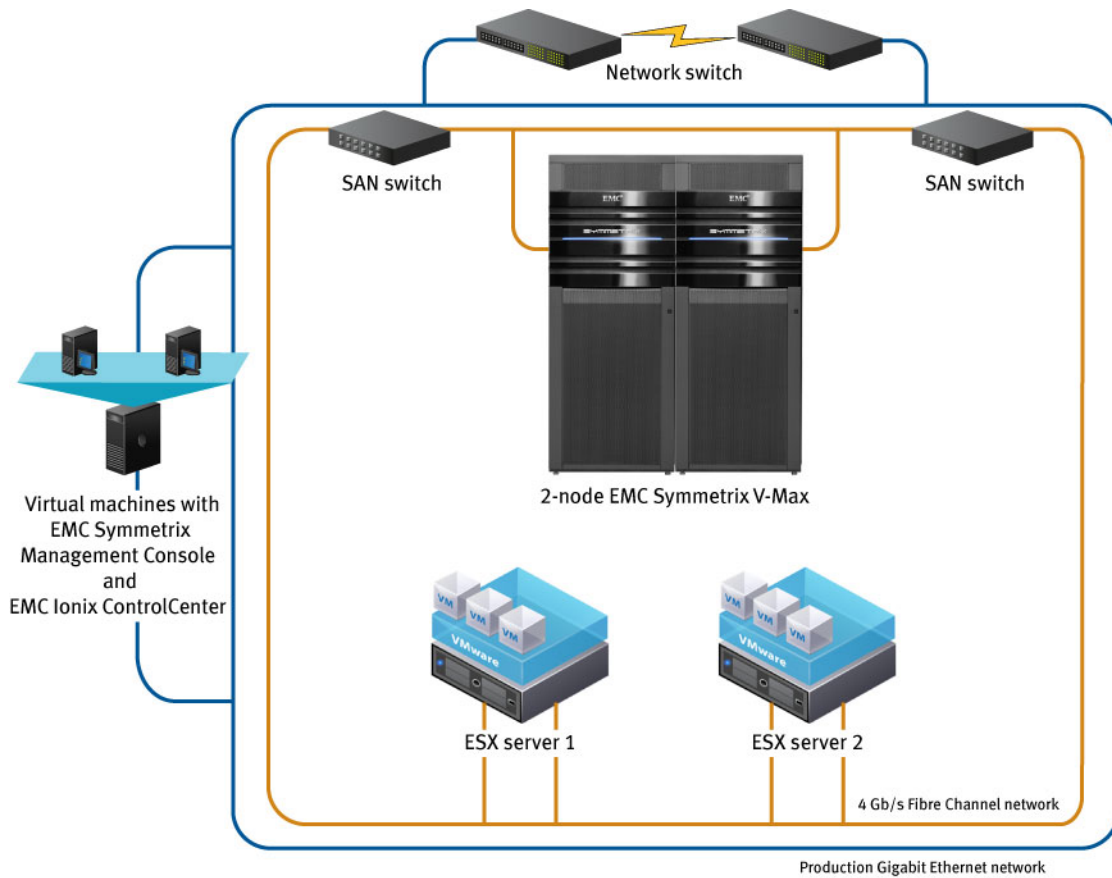


Each VMware vSphere ESX host had two dual-port HBAs installed. Each HBA was connected to two switches to ensure the highest levels of redundancy and performance. As the V-Max used for the solution was two-node (two engines), the connections were put in place to both load balance and protect, in case of component failure. Each ESX host is connected to four different storage processors, across both even and odd directors on the two V-Max Engines.

Physical architecture

Architecture diagram

The following image depicts the overall physical architecture of the solution.



SYM-002227

Validated environment profile

Profile characteristics

The solution was validated with the following environment profile.

Profile characteristic	Value
Database characteristic	OLTP
Benchmark profile	Swingbench OrderEntry - TPC-C-like benchmark
Target Oracle transaction response time	< 10 ms
Size of databases	2 x 300 GB and 1 x 150 GB
Number of databases	3
Array drives: size and speed	300 GB 15k rpm / 1 TB SATA 7.2k rpm

Hardware and software resources

Hardware The hardware used to validate the solution is listed below.

Equipment	Quantity	Configuration
Storage array	1	Symmetrix V-Max: <ul style="list-style-type: none"> • 64 GB mirrored cache • 80 x 300 GB 15k • 20 x 1 TB SATA
SAN	2	4 Gb capable FC switch
ESX servers	2	2 x Xeon X5570 processor, 2.93 GHz, 8 MB 144 GB RAM with two 146 GB SAS internal drives
Virtual machines:		
EMC ControlCenter	1	Win2k3
DHCP / DNS	1	Win2k3
Virtual Center	1	Win2k3
Load Simulator	1	Win2k3
Oracle Database 11g servers	3	Linux Red Hat 5.1

Software The software used to validate the solution is listed below.

Software	Version
Red Hat (OS for database)	5.1
VMware vSphere	4.0
VMware vCenter Converter	4.0
Oracle Database/Clusterware/ASM	11g Release 1 (11.1.0.7.0)
Swingbench load tool	2.3
EMC PowerPath	5.4
EMC Storage Viewer	2.0
EMC Solutions Enabler	7.0
Enginuity	5874
EMC ControlCenter	6.1 UB5
EMC Symmetrix Management Console	7.0

Testing and validation

Introduction

For this solution a number of Oracle databases were run on VMware vSphere virtual machines, connected to an EMC Symmetrix V-Max. At a high level the tests performed were:

VMware Converter P2V (physical to virtual conversion): VMware Converter was used to demonstrate the process of converting physical servers to virtual machines.

VMware RDM vs VMFS comparison: Oracle database loads were run, separately, against both RDM and VMFS storage configurations for Oracle 11g databases on VMware. The benefits of each configuration are detailed on the next page.

VMware VMotion: VMware VMotion was used to migrate a virtual machine from one ESX to another while under Oracle database load. The steps to run the migration were documented and the impact of running database transactions during the migration was identified.

VMware High Availability (HA): VMware HA was used to show automatic failover from one ESX server to another in the event of a physical server failure. For this test the primary ESX server was shut down, forcing a failover to occur. The recovery steps taken by HA for the entire process (including Oracle 11g database startup and availability) were documented.

Symmetrix V-Max Virtual LUN: EMC's Virtual LUN technology was used to seamlessly migrate or move the underlying storage to either different protection types or storage tiers. During the migration the Oracle 11g databases were constantly available and continued to operate under load.

Results and findings

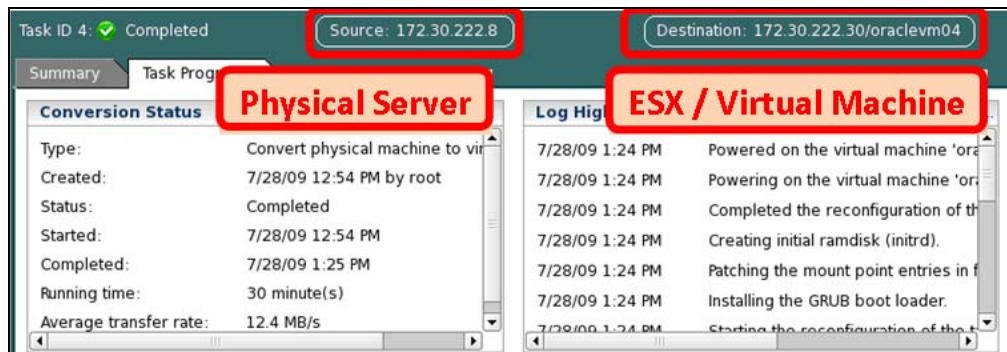
VMware Converter P2V testing

This solution validated the process of converting the physical database server to a virtual machine. The VMware Converter Conversion wizard was used to set up a migration task and Converter Standalone performed all the cloning tasks. The process was accelerated using EMC Symmetrix Management Console (SMC) to allocate the SAN storage to the new ESX destination server.

The high level steps were as follows:

- The Converter Standalone agent connects to the physical server to retrieve source information. The Converter Standalone agent creates an empty helper virtual machine based on the conversion task settings.
- The helper virtual machine is powered on, booting from the Linux image, and connects to the physical server. It begins retrieving the selected data from the source. Once the conversion task is set up, local file systems are selected to be copied to the destination virtual machine.
- The Converter Standalone agent shuts down the helper virtual machine. The conversion process is now complete.

- The database devices are unmasked from the source physical server, using EMC SMC, and allocated to the new ESX destination server.



RDM vs. VMFS comparison testing

Oracle database loads were run separately against both RDM and VMFS storage configurations for Oracle 11g databases on VMware. For these tests there were two VMware virtual machines, one utilizing an RDM storage configuration, and the second virtual machine used a VMFS storage configuration. The Oracle 11g databases for both tests were 300 GB in size.

RDM datastore configuration

	LUN Number	Size (GB)	Datafile Type
1	LUN 1	50	Contains database DATA device 1
2	LUN 2	50	Contains database DATA device 2
3	LUN 3	50	Contains database DATA device 3
4	LUN 4	50	Contains database DATA device 4
5	LUN 5	50	Contains database DATA device 5
6	LUN 6	50	Contains database DATA device 6
7	LUN 7	50	Contains database DATA device 7
8	LUN 8	50	Contains database DATA device 8
9	LUN 9	50	Contains database DATA device 9
10	LUN 10	50	Contains database DATA device 10
11	LUN 11	50	Contains database DATA device 11
12	LUN 12	50	Contains database DATA device 12
13	LUN 13	50	Contains database FRA device 1
14	LUN 14	50	Contains database FRA device 2
15	LUN 15	50	Contains database REDO device 1
16	LUN 16	50	Contains database REDO device 2
17	LUN 17	50	Contains database TEMP device 1
18	LUN 18	50	Contains database TEMP device 2

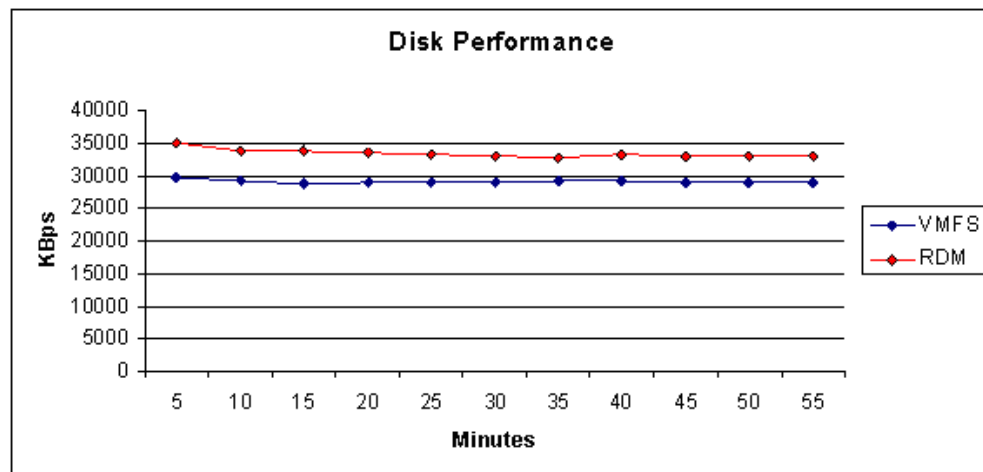
VMFS datastore configuration

	Datastore Type	Description
1	DATA datastore	Contains all database DATA devices
2	FRA datastore	Contains all database FRA devices
3	REDO datastore	Contains all database REDO devices
4	TEMP datastore	Contains all database TEMP/UNDO devices

The Symmetrix V-Max devices were presented to the virtual machine in a logical way that separated the database components based on their performance requirements and profile. The database files that are most active, and demand the best response times, are the redo logs and the datafiles. The database operations that are most response time-sensitive are writes to the redo logs and reads from datafiles and indexes. Writes to the datafiles are asynchronous and are, therefore, not as sensitive.

A database load was run over a one-hour period against both the databases residing on RDM configured devices and VMFS configured devices. There is a slight increase of 6 percent in TPM when running the database on RDM devices compared to VMFS devices. It was also observed in this configuration that on average there was a 13 percent benefit in Kb/s when using RDM compared to VMFS.

In summary, defining the best storage configuration option with VMware depends on the workload. Generally, RDMs will give better performance for small block size transactions and as can be seen from this use case, with a random read and write database profile, RDM yields slightly better I/O performance over VMFS.



VMware VMotion testing

A database load was run against a virtual machine with a 300 GB Oracle 11g database residing on VMware VMFS volumes. While under load the virtual machine was moved, online, using VMotion to a second ESX server. The objective was to observe the effect of carrying out an online migration of a virtual machine between ESX servers.

When activated, VMware VMotion performs the following actions:

- VMware's clustered Virtual Machine File System (VMFS) ensures the state of a virtual machine is encapsulated by a set of files stored on shared SAN storage. VMFS allows multiple installations of ESX Server to access the same virtual machine files concurrently.
- Using VMware VMotion the active memory and precise execution state of the virtual machine are rapidly transferred over a high-speed network. This allows the virtual machine to instantaneously switch from running on the source ESX server to the destination ESX server.

The networks being used by the virtual machine are also virtualized by the underlying ESX server, ensuring that even after the migration, the virtual machine network identity and network connections are preserved.


Note

VMotion also manages the virtual MAC address as part of the migration process.

VMotion migration process - Swingbench data

The following Oracle Swingbench transactions per minute (TPM) chart was captured while running database load during a VMotion migration process. In summary:

- There was a dip in transactions per minute during the migration process.
- The chart also shows the start and end point of the VMotion migration process
- The virtual machine was migrated to the second ESX server and the entire migration completed in a matter of minutes
- When the migration completed database user transactions returned to normal



VMware HA testing

VMware HA continuously monitors all servers in a SAN cluster and detects physical server and network failures. HA leverages shared storage and the VMware vStorage VMFS to enable the other servers in the cluster to safely access the virtual machine for failover. In this solution, failover was tested on a two-node VMware HA cluster with the forced shutdown of a single ESX server. The entire failover operation was automated through VMware HA. Customized Oracle database startup scripts enabled automatic restart of both Oracle ASM and the Oracle 11g database on the virtual machine at boot-up time.

These failover steps were logged in VMware vCenter:

1. The primary ESX server shuts down.
2. The vCenter Server sends an alert that there was a possible host failure of the primary server.
3. All virtual machines are relocated to the secondary ESX server and restarted automatically.

Note

Once the virtual machines are powered up the Oracle database may be configured to start automatically and made available for user connections.

The following image shows the failover as logged in VMware vCenter.



Symmetrix V-Max Virtual LUN

During testing both the SMC (Symmetrix Management Console) and Solutions Enabler command line options were available to carry out the Virtual LUN migration.

To show the functionality of the Virtual LUN it is important to understand the disk placement of the Oracle databases on the Symmetrix V-Max. Using the EMC Storage Viewer for vSphere client it is possible to identify exactly which Symmetrix LUN is being used by which virtual machine and application.

Using the SMC LUN Migration wizard, LUNs, on which an Oracle 11g database resided, were moved in a manner that was nondisruptive to the host or database.

For this test, the pre-configured option for the migration was used. The aim of the virtual LUN migration was to move the underlying storage from RAID 1 300 GB FC disks to RAID 6 1 TB SATA disks.

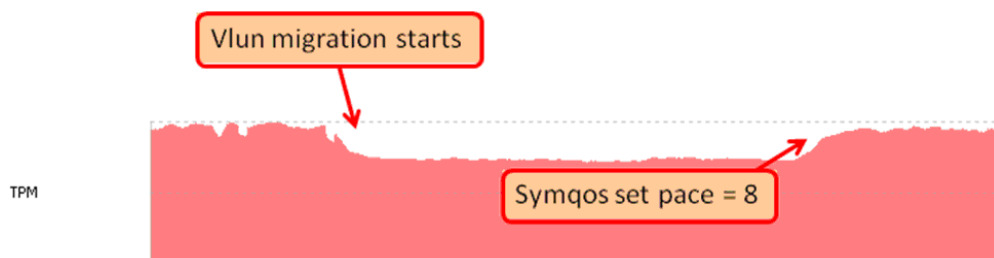
The source devices for the migration were LUNs A9 – BA. These were already placed in a device group, **oraVM01**, using Solutions Enabler. To limit the impact on transactions and response times, Symmetrix priority controls were used. **Symqos** is the command line syntax used to speed up or slow down the transfer rate of the

migration. By default the setting is 1, which is the fastest; settings up to 16 (slowest) can be set. In this testing, load was running against the Oracle 11g database, the Virtual LUN migration was then initiated. At this stage the transactions decreased and response times increased.

The **symqos** command sets the Symmetrix priority control, and in this instance was used to set the pace of the migration to 8:

symqos -g oraVM01 set MIRR pace 8

Once the priority controls were set, the transactions and the response times returned to the normal rates.



Conclusion

Summary

Companies today are under pressure to lower both the operational and capital costs related to their Oracle databases and critical applications systems. Virtualization offers real savings to IT departments in improving flexibility and lowering cost. Two of the most important technology areas that can help reduce the total cost of an Oracle 11g environment are storage and server virtualization. When combined, these technologies make Oracle 11g deployments more cost effective.

This proven solution reference architecture details a virtualized Oracle infrastructure design leveraging an EMC Symmetrix V-Max array and VMware vSphere 4 virtualized servers. This guide also includes various test results, configuration practices, and recommended specific Oracle storage design layouts that meet both capacity and consolidation requirements. Described in this document are many of the technologies that enable the benefits outlined below.

Summary of benefits

This solution provides the following benefits:

Reduction of cost – server consolidation: The virtualized configuration using VMware vSphere 4 allowed multiple Oracle database environments to be consolidated into a two-node VMware cluster configuration. Multiple database environments were run across a highly available infrastructure. The optimization of physical server hardware leads directly to decreased hardware power, cooling, and licensing costs.

Ease of application deployment: In order for customers to be able to fully realize the benefits of virtualization they first need to be moved from physical infrastructure to virtual infrastructure. Through the use of the VMware vCenter 4.0.1 Converter and EMC SAN management tools, conversion from a physical environment to a virtual database infrastructure can be done quickly. This decouples the application/database from the physical infrastructure.

VMware High Availability - reducing data center downtime with VMware VMotion and VMware HA: Infrastructure downtime is inevitable in today's complex IT environments because it is impossible to plan for every scenario. However, planned downtime can be significantly reduced with effective planning and the use of technology such as VMware VMotion. This study clearly demonstrates that VMotion provides a powerful way to eliminate or greatly reduce planned downtime at a lower cost than traditional approaches.

Unplanned downtime can be avoided by successfully implementing VMware HA across the virtualized cluster. In the event of a failure, all virtual machines can be automatically relocated to another secondary ESX server. This failover can be fully automated down to the restart of the Oracle 11g database on the virtual machine.

Storage efficiencies - increased flexibility: Multiple Oracle database deployments can be easily configured. The design outlined the RAID types and volume management options such as RDM and VMFS, as well as showing how to specify LUN sizes to suit the needs of the environment profile. Also, VLUN technology was utilized to move between certain RAID types nondisruptively.

This solution also includes several other components that lower IT costs even further, for example, EMC Symmetrix V-Max and VMware vSphere 4.

Included in the tested configuration is:

- EMC Symmetrix V-Max — This efficient primary storage array comfortably handles the performance and storage capacity requirement of multiple Oracle databases.
- Oracle Database 11g — Oracle Database 11g Enterprise Edition delivers industry-leading performance, scalability, security, and reliability on a choice of clustered or single-servers running Linux.
- Red Hat Linux — Red Hat Enterprise Linux Enterprise Linux 5 provides integrated server and storage technologies. In a single, easily managed product.
- VMware vSphere — The industry's first cloud operating system leverages the power of virtualization to transform data centers into dramatically simplified cloud computing infrastructures and enables IT organizations to deliver the next generation of flexible and reliable IT services, using internal and external resources, securely, and with low risk.

Next steps

EMC can help to accelerate assessment, design, implementation, and management while lowering the implementation risks and costs of a virtual infrastructure solution for an Oracle Database 11g environment.

To learn more about this and other solutions contact an EMC representative or visit www.EMC.com/solutions/oracle.
