



**EMC Unified Storage for
Oracle Database 11g**

Enabled by EMC CLARiiON, EMC Celerra,
VMware vSphere, and
VMware High Availability (HA) Using FCP and NFS

Reference Architecture

EMC Global Solutions



Copyright © 2010 EMC Corporation. All rights reserved.

Published April 2010

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: H6879.2

Table of Contents

Reference architecture overview	4
Solution components.....	7
Technology solution	9
Migration	12
Key elements	17
Physical architecture	20
Validated environment profile	21
Hardware and software resources	22
Conclusion	24

Reference architecture overview

Document purpose

This document provides an overview of the architecture of a virtualized solution that uses the EMC[®] Celerra[®] unified storage platform consisting of:

- Celerra Network Server
- CLARiiON[®] CX4 SAN array (CX4)
- Oracle Database 11g on Linux over the Fibre Channel Protocol (FCP) and Network File System (NFS)

Information in this document can be used:

- As the basis for a solution build, white paper, best practices document, or training.
 - By other EMC organizations (for example, the technical services or sales organization) as the basis for producing documentation for a technical services or sales kit.
-

Solution purpose

The purposes of this solution are to:

- Demonstrate the functional, performance, resiliency, and scalability capabilities of an Oracle software stack that:
 - Is booted on virtualized hardware and incorporates a four-node VMware High Availability (HA) cluster
 - Uses FCP and Oracle ASM to access high-demand, low-latency storage elements
 - Uses NFS to access all other storage elements
 - Compares the performance between Virtual Machine File System (VMFS) and Raw Device Mapping (RDM).
- Integrate VMware High Availability (HA) and Distributed Resource Scheduler (DRS) clustering with Oracle Database 11g and 10g.
- Enable the physical backup and recovery of an Oracle 11g production database while offloading all performance impacts of the backup operation from the production server, using:
 - EMC Replication Manager
 - EMC SnapView[™]

This demonstrates significant performance and manageability benefits in comparison to normal Oracle Recovery Manager (RMAN) backup and recovery.

- Provide disaster recovery capability using:
 - EMC RecoverPoint with CLARiiON splitters
 - VMware vCenter Site Recovery Manager

This demonstrates significant performance and manageability benefits in comparison to normal Oracle Data Guard disaster recovery.

- Provide the capability to clone a running production database with minimal performance impact and no downtime using EMC SnapView clones.
-

The business challenge

All enterprises face the same challenges when it comes to managing database environments. These challenges include:

- Rising costs
- Control over resource utilization and scaling
- Lack of sufficient IT resources to deploy, manage, and maintain complex environments at the departmental level
- The need to reduce power, cooling, and space requirements

All enterprises are constrained by limited budgets and many cannot afford a custom, one-off solution for every business function. This makes the process of creating a database solution for these enterprises even more challenging.

Traditional database deployments have relied on utilizing only an FCP storage infrastructure in a physically booted environment.

Through significant technology innovations and close collaboration with both Oracle and VMware, EMC is able to offer new advanced solutions, incorporating flexible, tiered storage functionality and multiple protocols, as well as optimized use of virtualization.

Customers' choice to leverage different storage protocols, functionality, and virtualization should be based on their individual application requirements for utilization, performance, flexibility, and cost.

EMC designed, tested, and documented this solution to specifically address the following challenges:

- Rising costs
Focus on deploying Oracle over blended FCP and NFS protocol environments to reduce cost and improve efficiency of networked storage
- Resource utilization
Align the technology strengths of networked storage with virtualization software to increase database server and storage utilization, with higher leverage at scale
- Resource efficiency
Simplify database resource and infrastructure administration utilizing a combination of FCP and NFS plus virtualization. Empower DBAs with enhanced solutions for backup, data movement, and data protection.

This solution documents implementation procedures and best practices for deploying a single platform solution to quickly and effectively to achieve these benefits.

The technology solution

This solution demonstrates how organizations can:

- Deploy a solution using a combination of the NFS and FCP protocols on the Celerra.
 - FCP is used for high-I/O and low-latency database objects (notably the datafiles, tempfiles, online redo logfiles, and controlfiles).
 - NFS is used for all other database objects (consisting basically of the flashback recovery area, disk-based backups, and archive logs).
 - Manageability advantages are obtained by using a combination of FCP and NFS. Specifically, archived logs and backups can be accessed through a normal file system interface rather than ASM. This simplifies administration of the database servers.
 - Avoid investing in additional FC infrastructure by implementing a blended solution that uses both FCP and NFS to access storage elements.
 - Get the most out of hardware by using VMware DRS to automatically reallocate virtual machines, according to need, across servers.
 - Maximize the use of the database-server CPU, memory, and I/O channels by offloading performance impacts from the production server during:
 - Backup and recovery by using EMC Replication Manager
 - Disaster recovery operations by using EMC RecoverPoint combined with VMware vCenter Site Recovery Manager
 - Implement a disaster recovery solution with EMC RecoverPoint, combined with VMware vCenter Site Recovery Manager that reduces costs and complexity by using IP as the WAN network protocol.
 - Reduce the complexity of backup operations and eliminate the need to implement scripted solutions by using EMC Replication Manager.
 - Use SnapView to free up the database server's CPU, memory, and I/O channels from the effects of operations relating to backup, restore, and recovery. SnapView clones also help in creating test/development systems without any impact on the production environment.
 - Migrate a running production Oracle database across storage protocols in a virtualized environment.
-

Solution components

Introduction This section describes the components of the solution, and explains some of the key terminology and concepts that are used in this document.

Key terms defined The following table describes several terms used in this paper:

Term	Description
Solution	A solution is a complete, validated, and tested stack of hardware and software upon which a customer would choose to run their entire business or business function.
Solution attribute	A solution attribute addresses the entire solution stack, but does so in a way relating to a discrete area of testing. For example, performance testing is a solution attribute.
Solution component	A solution component addresses a subset of the solution stack that consists of a discrete set of hardware or software, and focuses on a single IT function. For example, backup and recovery, and disaster recovery are solution components. A solution component can be either “basic” or “advanced.”
Basic solution component	A basic solution component uses only the features and functionality provided by the Oracle stack. Examples are RMAN for backup and recovery, and Data Guard for disaster recovery.
Advanced solution component	An advanced solution component uses the features and functionality of EMC hardware or software. Examples are EMC SnapView for backup and recovery, and EMC MirrorView™ for disaster recovery.

Solution attributes The following table describes the solution attributes that are included in this solution:

Attribute	Description
Scale-out OLTP	This is the baseline performance test presently used for virtualized solutions. It consists of an industry-standard TPC-C workload running on a cluster of VMware ESX servers. Multiple database images are used for the workload, emulating a software-as-a-service (SaaS) customer. This is a scale-out, because all users in the database are not able to access all rows in the combined databases but can only see the data in the database to which it is connected. Thus, the database is federated.
Resiliency	The purpose of resiliency testing is to validate the fault-tolerance and high-availability features of the hardware and software stack. Faults are inserted into the configuration at

	<p>various layers in the solutions stack. Some of the layers where fault tolerance is tested include:</p> <ul style="list-style-type: none"> • VMware HA cluster interconnect port • Storage processor
--	--

Solution components

The following table describes the solution components that are included in this solution:

Component	Description
Basic Backup	This is backup and recovery using Oracle RMAN, the built-in backup and recovery tool provided by Oracle.
Advanced Backup	<p>This is backup and recovery using EMC value-added software or hardware.</p> <p>In this solution the following are used to provide Advanced Backup functionality:</p> <ul style="list-style-type: none"> • EMC Replication Manager • EMC SnapView
Basic Protect	This is disaster recovery using Oracle Data Guard, Oracle's built-in remote replication tool.
Advanced Protect	<p>This is disaster recovery using EMC value-added software and hardware.</p> <p>In this solution the following is used to provide Advanced Protect functionality:</p> <ul style="list-style-type: none"> • EMC RecoverPoint with CLARiiON splitters • VMware vCenter Site Recovery Manager
Test/Dev	<p>A running production OLTP database is cloned with minimal, if any, performance impact on the production server, as well as no downtime. The resulting dataset is provisioned on another server for use for testing and development.</p> <p>EMC Replication Manager is used to automate the test/dev process.</p>
Migration	The ability to migrate a running production Oracle database from one storage protocol to another (FCP/ASM to DNFS or KNFS, or vice versa) is a frequent customer request. This can include full production database cut-over, or the creation of a clone on another storage protocol for backup, test/dev, or other purposes.

Technology solution

Overview

Two sites connected by a WAN are used in the solution, one site is used for production and the other site is used as a disaster recovery target. A Celerra unified storage platform is present at each site. Oracle Database 11g for x86-64 is run on Red Hat Enterprise Linux or on Oracle Enterprise Linux.

The Oracle software stack is booted on virtualized hardware and incorporates a four-node VMware HA cluster. For the purpose of high availability, VMware VMotion is used to migrate a virtual machine running the production database from one ESX server to another without any downtime. Each ESX server is connected to both the production storage and to the cluster-interconnect TCP/IP networks. VMware DRS is used to manage the allocation of resources automatically to different virtual machines based on the need.

FCP storage networks consisting of dedicated, redundant FCP switches are present at both sites.

The solution also includes virtualized servers for use as Test/Dev, Advanced Backup, Basic Protect, and Advanced Protect targets. Virtualization of the test/dev and disaster recovery (DR) target servers is supported using a VMware ESX server.

Blended FCP/NFS

This is a blended FCP/NFS solution. Depending on the nature of the database object, either FCP or NFS is used to access it. FCP is used to access high-demand, low-latency objects. NFS is used to access low-demand, high-latency objects.

The *Storage layout* section describes in detail how each Oracle file type and database object is stored and accessed for this solution.

Production site

The following components are present at the production site and are connected to the production FCP storage network and to the WAN:

- Virtualized Oracle Database 11g scale-out servers running on virtual machines hosted on a four-node VMware HA cluster. The virtualized servers are connected to the production FCP storage network using virtualized connections.
- A Celerra unified storage platform is connected to the four-node VMware HA cluster through the production storage network. EMC SnapView is used to provide an advanced backup solution, and SnapView clones are used to create a test/dev target database.
- A RecoverPoint cluster connected to the FCP storage network and to the WAN.
- A production side vCenter server is present to manage the process of disaster recovery using VMware vCenter Site Recovery Manager.

The virtualized Oracle Database 11g servers are also connected to the client and VMware HA cluster interconnect networks.

Disaster recovery target site

The disaster recovery target site consists of:

- A virtualized single-instance Oracle Database 11g server that is used as:
 - The disaster recovery target for Basic Protect
 - The target for Test/Dev

The virtualized single-instance server is connected to the client, WAN, and target storage networks through virtualized connections on the virtualization server.

- A Celerra unified storage platform connected to the VMware ESX server through the disaster recovery storage network.
 - The Oracle Database 11g single-instance target. The server accesses this network through a virtualized switch on the ESX server. The CX4 contains the RecoverPoint splitter driver.
 - A RecoverPoint cluster connected to the FCP storage network, and to the WAN.
 - A Disaster Recovery (DR) target vCenter is also present to manage the process of failover, using VMware vCenter Site Recovery Manager.
-

Connected to both sites

A VMware ESX server is connected to both the production and target FCP storage networks.

The virtualized single-instance Oracle 11g target server accesses both the production and target FCP storage networks and is connected to the client WAN through virtualized connections on the virtualization server.

The LUNs on the Celerra unified storage platform are formatted using the VMFS file system; VMDK files are provisioned on the target virtual machines within this file system.

Storage layout

The following table describes how each Oracle file type and database object is stored and accessed for this solution:

What	Object type	Protocol	Stored on	File-system type
Oracle datafiles	High demand Low latency	FCP	FC disk (LUNs)	RAID-protected ASM diskgroup
Oracle tempfiles				
Oracle online redo logfiles				
Oracle controlfiles				
VMDK pointer files				
Archived logfiles	Low demand High latency	NFS	SATA II	RAID-protected NFS
Flashback recovery area				

Disk-based backups				
--------------------	--	--	--	--

- High-performance database objects are accessed over an FCP network using redundant network switches.
- FCP storage networks consist of dedicated, redundant FCP switches. Four separate sets of LUNs (for data and redo logs) are created on the back-end storage and each corresponding set is presented to the virtual machine using RDM and VMFS. The FCP LUNs were presented to the virtual machine using RDM on one of the configurations and using VMFS on the other. The performance on these configurations was validated and a comparative analysis performed.
- Advanced Backup and Test/Dev solutions were validated by integrating Replication Manager on the configuration where the LUNs were presented using RDM.
- The disaster recovery solution by integrating RecoverPoint was performed on the configuration, where the LUNs were presented using VMFS. VMware vCenter Site Recovery Manager was validated as a disaster recovery solution using RecoverPoint with CLARiiON splitters.

Oracle ASM is used as the file system and volume manager for the FCP mounted database objects. ASMLib is used to virtualize the LUNs on the database server.

- Oracle datafiles, tempfiles, and online redo logfiles are stored on separate LUNs that are mounted on the database server using ASM over FCP.
- Three ASM diskgroups are used - one diskgroup for datafiles and tempfiles, and two diskgroups for online redo logfiles.
- The online redo logfiles are mirrored across the two ASM diskgroups using Oracle software multiplexing. The controlfiles are mirrored across the online redo log ASM diskgroups.
- Each ASM diskgroup and its underlying LUNs are designed to satisfy the I/O demands of individual database objects. For example, RAID 5 is used for the datafiles but RAID 10 is used for the logfiles. All of these diskgroups are stored on FC disks.

IP network architecture

The design implements the following physical IP connections:

- TCP/IP and NFS provide network connectivity and file system semantics for NFS file systems on Oracle Database 11g.
- Client virtual machines run on the VMware ESX server. They are connected to a client network.
- Client, VMware HA cluster interconnect, and redundant TCP/IP storage networks consist of dedicated network switches and virtual local area networks (VLANs).
- The VMware HA cluster interconnect and storage networks consist of trunked IP connections to balance and distribute network I/O. Jumbo frames are enabled on these networks.

Migration

Introduction

Customers often request the ability to migrate a virtualized Oracle Database across storage protocols. In response to this, the Oracle Consulting (CSV) group has validated that customers who have an Oracle Database operating in a virtualized environment can migrate data from:

- An FCP/ASM to an NFS-mounted file system
- An NFS-mounted file system to an FCP/ASM

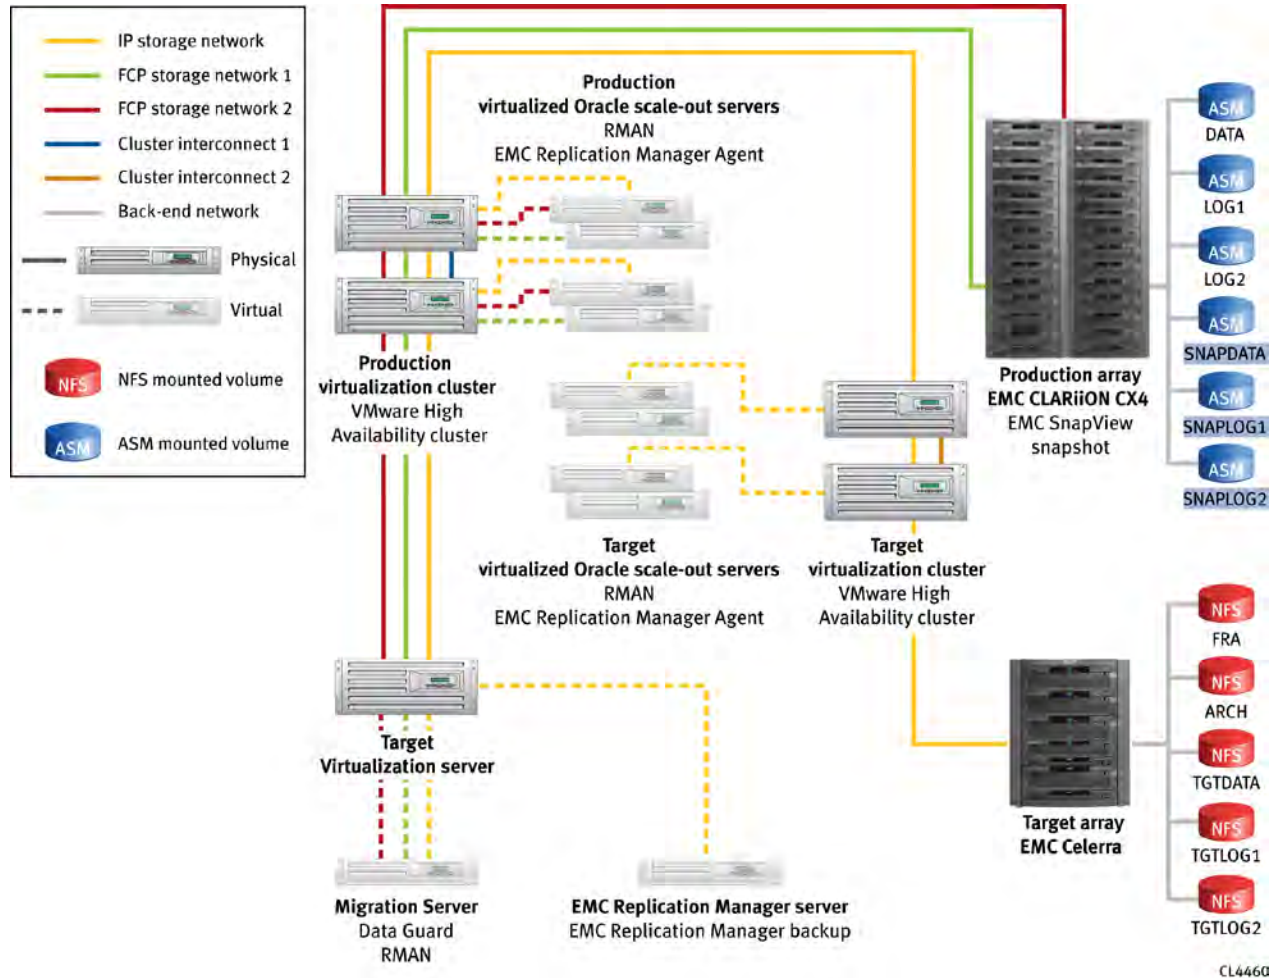
Customers might be interested in cross-protocol migration because testing and development activities can be supported on less expensive, simpler NFS storage, while production can be carried out on higher performing ASM/FCP storage.

Customers may also want to use the storage replication technologies available on a different protocol from the production database. For example, customers may wish to use RecoverPoint to support remote replication but have production on NFS.

The migrations were performed with minimal performance impact and no downtime on the virtualized production database.

SAN to NAS migration diagram

The following diagram is a high-level view of the migration component:



Migrating an online Oracle Database from SAN to NAS

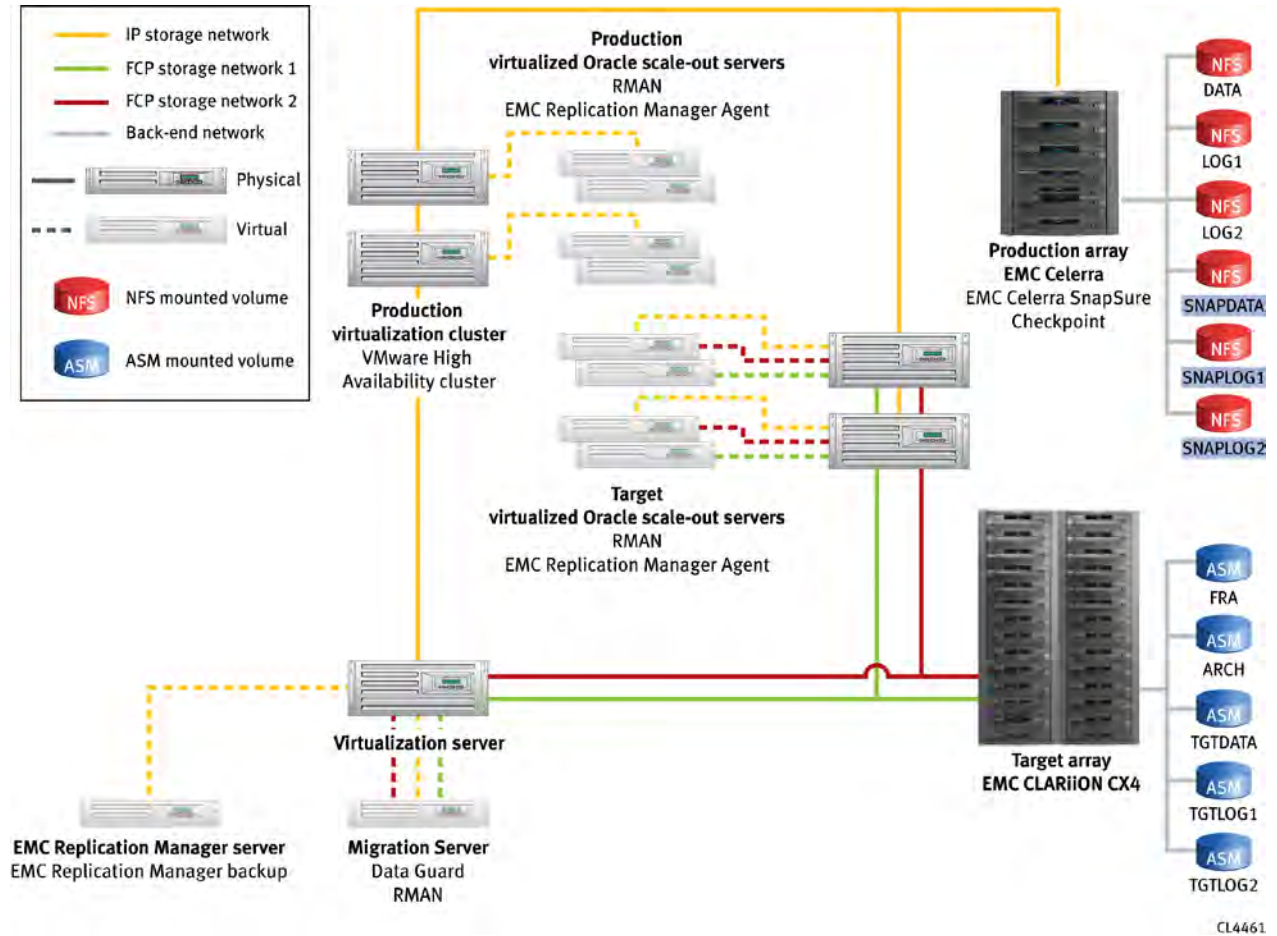
These steps were followed to perform the SAN to NAS migration operation:

Step	Action
1	Using EMC Replication Manager, a consistent backup of the running virtualized production database is performed on the CLARiiON using a SnapView snapshot.
2	This backup is mounted (but not opened) on the migration server, in this case a VMware virtual machine (a physically booted server would also work). The NFS target array is also mounted on the migration server.
3	Using Oracle Recovery Manager (RMAN), a backup of this database is taken onto the target location. This backup is performed as a database image, so that the datafiles are written directly to the target NFS mount.
4	The migration server is then switched to the new database, which has been copied by RMAN to the NFS mount.
5	The target database is set in Data Guard continuous recovery mode, and Data Guard log ship/log apply is used to catch the virtualized target database up to the virtualized production version.
6	Once the virtualized target database is copied to production, Data Guard failover can be used to retarget to the virtualized target database. If an appropriate networking configuration is performed, clients will not see any downtime when this operation occurs.

The result is that the virtualized production FCP-mounted database can be migrated to NFS with minimal performance impact and no downtime.

NAS to SAN migration diagram

The following diagram is a high-level view of the NAS to SAN migration component:



**Migrating an
online Oracle
Database from
NAS to SAN**

These steps were followed to perform the NAS to SAN migration operation:

Step	Action
1	Using EMC Replication Manager, a consistent backup of the running virtualized production database is performed on the EMC Celerra [®] , using a SnapSure checkpoint snapshot.
2	This backup is mounted (but not opened) on the migration server, in this case a VMware virtual machine (a physically booted server would also work). The FCP/ASM target LUNs are also mounted on the migration server.
3	Using Oracle Recovery Manager (RMAN), a backup of this database is taken onto the target location. This backup is performed as a database image, so that the datafiles are written directly to the target FCP/ASM LUNs.
4	The migration server is then switched to the new database, which has been copied by RMAN to the FCP/ASM LUNs.
5	The virtualized target database is set in Data Guard continuous recovery mode, and Data Guard log ship/log apply is used to catch the virtualized target database up to the production version.
6	Once the virtualized target database is caught up to production, Data Guard failover can be used to retarget to the virtualized target database. If appropriate networking configuration is performed, clients will see no downtime when this operation occurs.

The result is that the virtualized production NFS-mounted database can be migrated to FCP/ASM with minimal performance impact and no downtime.

Key elements

Introduction This section provides an overview of the technologies that are used in this solution.

EMC Celerra unified storage platform Celerra includes a world-class network-attached storage (NAS) array combined with the functionality and performance of the leading midrange storage area network (SAN) array. Celerra provides both NAS and SAN functionality and performance without compromise.

The key features provided by Celerra are described in the following table:

Feature	Provided by
SAN storage	Fibre Channel Protocol (FCP) through the back-end EMC CLARiiON CX4 storage array
NAS	Network File System (NFS) and Common Internet File System (CIFS) protocols
iSCSI storage	Celerra's Data Movers

EMC SnapView EMC SnapView lets you create local point-in-time snapshots and complete data clones for testing, backup, and recovery operations. With SnapView, you can create multiple copies of production data on your CLARiiON quickly and easily.

SnapView enables you to:

- Carry out nondisruptive backup operations with space-saving SnapView snapshots
 - Restore data quickly and easily with SnapView clones
 - Accelerate application development by quickly providing the most current data for testing
-

EMC Replication Manager EMC Replication Manager manages EMC point-in-time replication technologies through a centralized management console. Replication Manager coordinates the entire data replication process – from discovery and configuration to the management of multiple application consistent disk-based replicas. Replication Manager can auto-discover your replication environment and enable streamlined management by scheduling, recording, and cataloging replica information including auto-expiration.

Replication Manager integrates with the Oracle Database server and provides an easy interface to create and manage Oracle replicas.

EMC RecoverPoint

EMC RecoverPoint provides integrated continuous remote replication (CRR) and continuous data protection (CDP) for on-demand protection and recovery at any point in time. RecoverPoint's advanced capabilities include policy-based management, application integration, and WAN acceleration.

RecoverPoint can be implemented to provide a single, unified solution to protect and replicate data across heterogeneous storage. This enables the customer to simplify management and reduce costs, recover data at a local or remote site at any point in time, and ensure continuous replication to a remote site without impacting performance.

EMC RecoverPoint CLARiiON splitter

The CLARiiON (CX) splitter is a driver that can be directly installed on a CLARiiON. When an application writes to storage, the CX splitter splits the data and sends one copy of the data over FCP to the source-site RecoverPoint appliance (RPA), and a second copy to the storage array. These requests are acknowledged by the source RPAs and are then replicated to the remote RPAs over the IP WAN network.

The main advantages of CLARiiON array-based splitters are that they:

- Reduce the cost associated with the additional host-based splitter agents or specialized fabric
 - Provide concurrent local and remote (CLR) protection. CLR replication can be performed on the same LUNs.
-

Oracle software stack

The Oracle software stack covered by this solution consists of:

- Oracle Database 11g
 - Oracle Enterprise Linux
 - Automatic Storage Management (ASM)
-

VMware ESX

VMware ESX is the flagship virtualization server product from VMware and is the market leader in server virtualization and cloud computing. VMware vSphere provides a high-performance, robust, fault-tolerant and high-availability virtualization solution. This also helps in leveraging virtualization features to enable cloud computing, which will help in better manageability of IT infrastructure and lowest total cost per application workload.

VMware Distributed Resource Scheduler

VMware Distributed Resource Scheduler (DRS) is a cluster feature of VMware vCenter Server; it provides dynamic load balancing and resource sharing for multiple virtual machines across ESX servers. DRS uses VMotion as the underlying transport feature to move virtual machines from one ESX server to another.

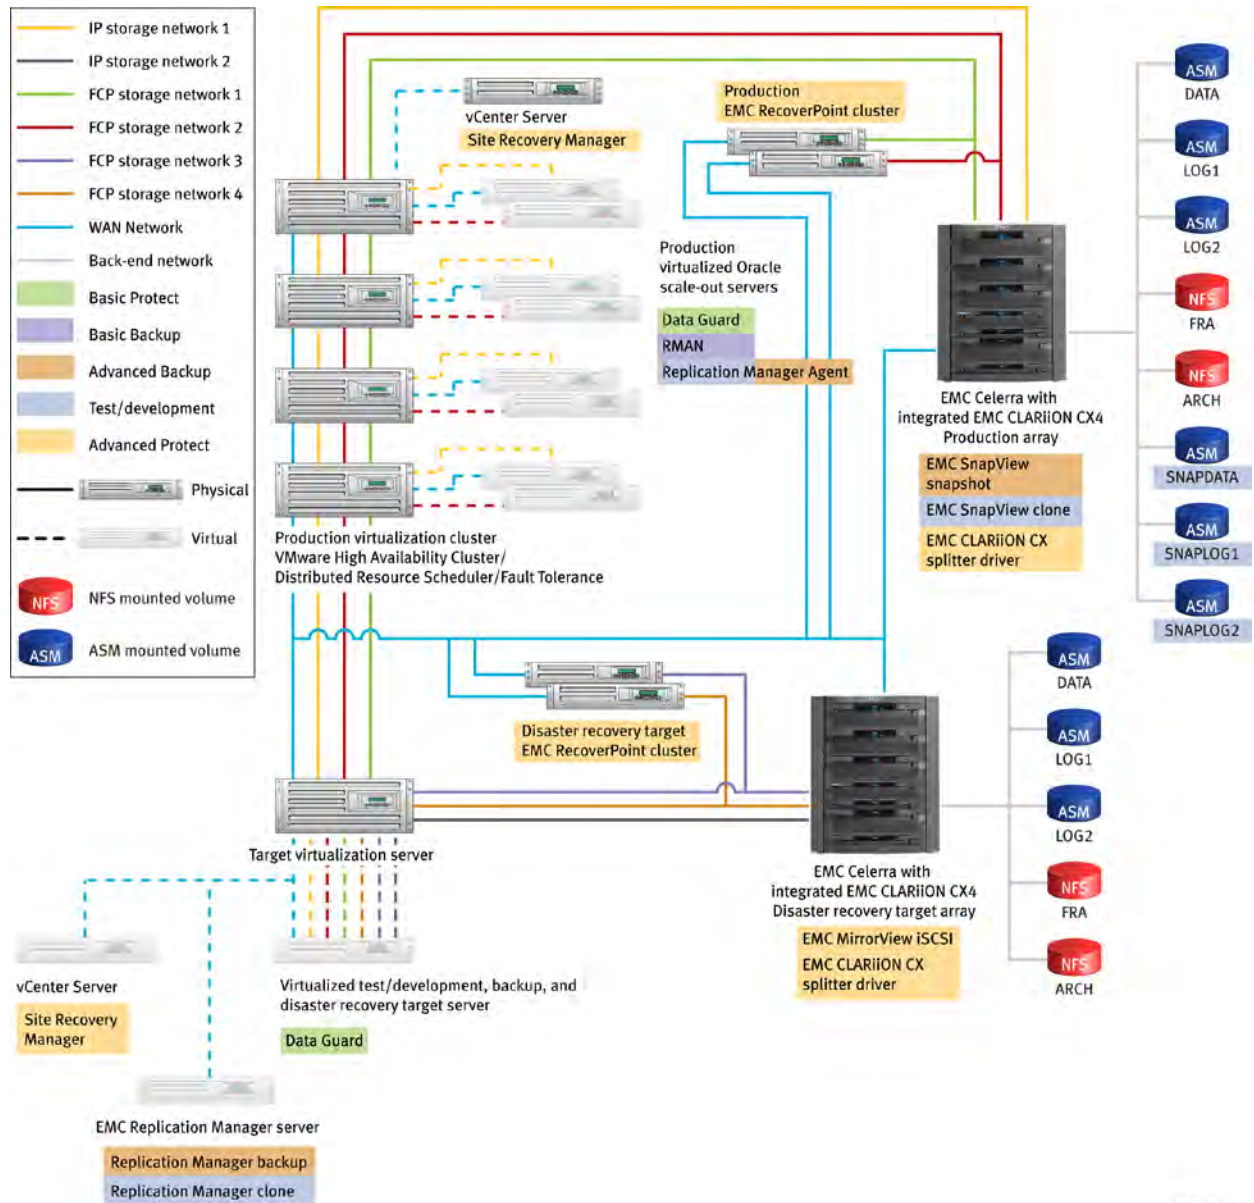
VMware Site Recovery Manager

VMware vCenter Site Recovery Manager is a complete, full stack, remote replication solution for VMware virtual machines. It uses VMware vCenter servers on the production and target sides of the replication to coordinate the process of replicating the operating system (OS) image, as well as the database or other application data. The process of replicating the full database server OS, plus the Oracle Database, using VMware vCenter Site Recovery Manager, was validated and tested. VMware vCenter Site Recovery Manager uses software plug-ins called Storage Replication Adapters (SRAs) in order to facilitate the replication. These allow VMware vCenter Site Recovery Manager to interface with storage layer replication technology in order to perform the replication. The RPA used in this testing was RecoverPoint, using the CLARiiON splitter.

Physical architecture

Reference architecture diagram

The following diagram depicts the overall physical architecture of the solution:



CNS-001575

Validated environment profile

Profile characteristics

For information on the validated environment profile and performance results, refer to the Proven Solution Guide for this solution. This information can be accessed on EMC Powerlink[®], EMC.com, and the EMC|KB.WIKI.

Hardware and software resources

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

Equipment	Quantity	Configuration
EMC Celerra unified storage platforms (included an EMC CLARiiON CX4 back-end storage array)	2	<ul style="list-style-type: none"> • 2 Data Movers • 4 GbE network connections per Data Mover • 2 or 3 FC shelves • 1 SATA shelf • 30 or 45 73 GB FC disks (depending on configuration) • 15 1 TB SATA disks • 1 Control Station • 2 storage processors • DART version 5.6.44-4
Gigabit Ethernet switches	5	24 ports per switch
FCP switches	2	<ul style="list-style-type: none"> • 16 ports • 4 Gb throughput
VMware ESX HA cluster servers	4	<ul style="list-style-type: none"> • 2 2.66 GHz Intel Pentium 4 quad-core processors • 24 GB of RAM • 2 146 GB 15k internal SCSI disks • 2 onboard GbE Ethernet NICs • 2 additional Intel PRO/1000 PT quad-port GbE Ethernet NICs • 2 SANblade QLE2462-E-SP 4 Gb/s dual-port FC HBAs (4 ports in total)
Virtualization server (VMware ESX server)	1	<ul style="list-style-type: none"> • 4 2.86 GHz AMD Opteron quad-core processors • 32 GB of RAM • 2 146 GB 15k internal SCSI disks • 2 onboard GbE Ethernet NICs • 3 additional Intel PRO/1000 PT quad-port GbE Ethernet NICs • 2 SANblade QLE2462-E-SP 4 Gb/s dual-port FC HBAs (4 ports in total)

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

Software	Version
Oracle Enterprise Linux	4.7
VMware ESX Server/VMware vSphere	4.0
VMware vCenter	4.0
VMware Site Recovery Manager	4.0.0
Microsoft Windows Server 2003 Standard Edition	2003
Oracle Database Standard Edition	11g (11g Ver. 11.1.0.7) or 10g
Quest Benchmark Factory for Databases	5.8.1
EMC Celerra Manager Advanced Edition	5.6
EMC Navisphere [®] Agent	6.26.0.2.24
EMC PowerPath [®]	5.0.0 (build 157)
EMC FLARE [®]	04.28.000.5.504
EMC DART	5.6.44-4
EMC Navisphere Management	6.28
EMC MirrorView	6.7
EMC RecoverPoint	3.0 SP1
EMC RecoverPoint SRM RPA	1.0 SP2
EMC CLARiiON splitter driver	04.28.006.6.003
EMC Replication Manager	5.2.0.0

Conclusion

Summary This section provides a summary of the solution and of the business challenges that it addresses.

Reduced total cost of ownership In any reasonable configuration, the database server's CPU is the most precious component of the entire architecture. Therefore, the over-arching principle of EMC's Oracle Database 11g solutions is to free up the database server's CPU (as well as memory and I/O channels) from utility operations such as backup and recovery, disaster recovery staging, test/dev, and cloning. The highest and best use of the database server's CPUs is to parse and execute the SQL statements that are required by the application user.

CPU usage

This solution reduces the load on the database server CPU by using:

- EMC Replication Manager and EMC SnapView to carry out a physical backup of an Oracle 11g production database, while offloading all performance impacts of the backup operation from the production server.
 - EMC RecoverPoint for disaster recovery operations.
-

Blended FCP/NFS architecture The use of a blended FCP/NFS architecture provides very significant ease-of-use advantages. Provided that the customer has access to a Celerra unified storage platform, the amount of software that the customer is required to install and configure on the database server is significantly reduced.

The NFS protocol provides access to storage elements such as archive logs and disk-based backups using normal file-system semantics. This is a significantly simpler form of management than that provided through FCP and ASM.

The use of FCP allows for high-performance, low-latency I/O for databases that require high-performance storage.

Virtualization The use of virtualization for the production database server provides manageability and ease-of-use advantages. In a scale-out context, virtualization can provide superior performance and scalability compared to physically booted configurations – even when using hardware identical to that used in the physically booted configuration.

Utility servers, such as a test/dev target and basic protect target, are more easily and conveniently managed as virtual machines than as physically booted Oracle Database servers. The advantages of consolidation, flexible migration and so on, which are the mainstays of virtualization, apply to these servers as well.

Reduced costs

One of the main challenges faced by the customer is to reduce cost by utilizing infrastructure effectively. Virtualization enables reduction in the number of servers

and related IT hardware in the data center. The other main feature is the ability to move a running virtual machine production database from one physical sever to another physical server without any downtime.

Distributed resource scheduling

VMware DRS provides the ability to distribute the workload across multiple ESX servers by migrating virtual machines according to resource consumption and demand. The VMware DRS helps in dynamically managing the resources pooled together from multiple ESX servers. VMware helps in aggregating and centrally managing the resources of multiple hosts in the form of resource pools. Different sets of rules and priority can be defined for individual Oracle virtual machines based on the criticality. When an Oracle Database server virtual machine operating within a resource pool experiences increased utilization of resources, DRS will try to allocate the resources from the centralized pool based on the rules and priorities.

High availability

VMotion

The use of VMware VMotion to migrate an Oracle Database instance from one piece of hardware to another was achieved in our tests with virtually no performance impact and no downtime on the running Oracle Database instance. This provides a very high level of manageability and downtime reduction for tasks such as software and hardware upgrades.

VMware HA cluster

The use of a VMware HA cluster provides further high-availability advantages to the virtualized solution. EMC has validated the use of a VMware HA cluster with Oracle Database 11g in a single-instance scale-out environment. In addition, a VMware HA cluster can be used to automate failover in the event of a hardware or software failure.

**Business
continuity**

Advanced backup and recovery

EMC Replication Manager is a comprehensive graphical application that provides Oracle storage replication using EMC storage technology. This eliminates the requirement for the customer to write scripts or to manually perform replication tasks. These tasks can now be fully automated and managed by Replication Manager.

Advanced Protect using EMC RecoverPoint and CX splitters with VMware Site Recovery Manager

The advantages of Advanced Protect using RecoverPoint through the CX splitter driver with VMware vCenter Site Recovery Manager are as follows:

- **Cost savings:** RecoverPoint provides data compression of between two and 15 times before shipping data over the wire (typical compression is approximately five times).
- **Superior data integrity:** RecoverPoint logs every change, and maintains those changes as long as the journal volume is capable of tracking them (this depends on change rate and the size of the journal).
- **No disruption to replication of production data:** RecoverPoint allows the customer to rapidly provision a copy of the database to test disaster recovery or other tasks, without any disruption to the ongoing replication of the production data.
- **Simultaneous local and remote replication:** RecoverPoint is unique in that it allows for simultaneous replication targets from the same source.
- **Minimal performance impact:** The performance testing carried out by EMC demonstrates that RecoverPoint has minimal, if any, impact on performance.
- **Full stack replication:** vCenter Site Recovery Manager is unique among remote replication technologies in that it provides replication of the entire stack, rather than simply the database layer. This provides the customer with a greater assurance that the entire database environment can be brought up successfully on the target side in the event of a disaster.

Test/Dev

The ability to deploy a writeable copy of the production database is required by many customers. The process of provisioning this copy must create minimal, if any, performance impact on the production database server. Also, absolutely no downtime can be tolerated. The Test/Dev solution component documented here provides this deployment ability.

Robust performance and scaling

The resiliency testing carried out by EMC ensures that the database configuration is reliable. High availability is used at every major layer of the solution, including the database server, NAS file server, and back-end CX4. By testing the fault tolerance of all of these layers, the ability of the application to withstand hardware failures with no downtime is assured.

The performance testing carried out by EMC utilizes an industry-standard OLTP benchmark but does so without exotic tunings that are not compliant with best practices. In addition, real-world configurations that the customer is likely to deploy are used. This enables the customer to be reasonably assured that the configuration that they choose to run their application will do so predictably and reliably.

Next steps

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

To learn more about this and other solutions contact an EMC representative or visit <http://www.emc.com/solutions/application-environment/oracle/solutions-for-oracle-database.htm>.
