

VIRTUAL TEST/DEV ENVIRONMENT PROVISIONING FOR ORACLE RAC 11g WITH EMC UNIFIED STORAGE

A Detailed Review

EMC GLOBAL SOLUTIONS

Abstract

This white paper demonstrates the combination of EMC® snapshots, VMware® virtualization, and Oracle Direct NFS (DNFS) clonedb feature technologies. The EMC Celerra® NS-960 unified storage platform provides great flexibility and a highly competitive total cost of ownership (TCO) for replicating an Oracle production database to create Test/Dev and quality assurance (QA) environments quickly and easily.

September 2011

vmware®
PARTNER

TECHNOLOGY
ALLIANCE

Copyright © 2011 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.

VMware, ESX, vMotion, VMware vCenter, and VMware vSphere are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. All other trademarks used herein are the property of their respective owners.

Part Number h7256.1

Table of Contents

Executive summary	5
Business case	5
Solution overview.....	5
Key results	5
Introduction	7
Purpose.....	7
Scope.....	7
Audience.....	7
Terminology	8
Technology overview	9
Introduction	9
EMC Celerra unified storage platform.....	9
EMC Replication Manager.....	10
EMC SnapSure.....	10
VMware vCenter Converter.....	10
VMware vSphere 4.1	11
Performance.....	12
Server consolidation	12
Rapid provisioning	12
Workload management	12
High availability (HA).....	12
Business continuity.....	12
Storage architecture	12
Oracle software stack	13
Configuration	14
Overview	14
Physical environment.....	14
Environment profile	15
Server and network architecture	15
Storage architecture	19
Disk layout	21
Test and validation	22
Introduction	22
Configure the EMC Celerra NS-960 storage array.....	22
Complete VMware vSphere 4.1 tasks.....	22
Install an Oracle 11g R2 RAC database	23
Enable DNFS.....	23

Create a baseline performance	24
Create a Celerra SnapSure checkpoint and perform P2V conversion	24
Use Oracle DNFS clonedb to get the test database	29
Provision multiple virtual machines for Test/Dev from the virtual machine template	31
Test summary.....	32
Conclusion	33
Summary.....	33
Findings	33
References	34
White paper.....	34
Oracle Metalink Note.....	34
Oracle online document	34
VMware online documents	34

Executive summary

Business case

One of the most challenging tasks for Oracle Real Application Cluster (RAC) database administrators is to create, deploy, and manage a large number of production environment images. This includes deploying the environment for testing, development, reporting, and other operations. Administrators often have to perform these activities frequently and quickly to meet internal service-level agreements (SLAs).

On the other hand, cost reduction and resource optimization are key requirements for today's companies. It is increasingly important to build development, quality assurance (QA) testing, or other non-production environments, based on common and reusable components.

Solution overview

EMC® Celerra® unified storage platforms offer great flexibility, a highly competitive entry price, and multiple ways to replicate the production database environment for various purposes, such as testing and development (Test/Dev), QA, or backup. When combined with virtualization technologies and powerful tools such as EMC Replication Manager and Oracle Direct NFS (DNFS) clonedb, these platforms can offer the immediate ability to use the newly provisioned virtual instances, because the replication of a physical production environment into a virtual environment can be achieved within 30 minutes.

This solution demonstrates:

- The ease of creating an image of the Oracle production database using network-attached storage (NAS) through EMC snapshot technology.
- How VMware® virtualization simplifies management and reduces cost:
 - Single-instance architectures save hardware and license costs.
 - VMware tools offer quick ways to repurpose and protect the data.
 - VMware technologies offer performance advantages and additional high availability (HA), migration, load balancing, and other options.
- How to use Oracle's DNFS clonedb feature to clone the test databases instantaneously.

Key results

This white paper demonstrates the following benefits of the solution:

- Low impact of storage-based replication strategies on the normal I/O requests per second (IOPS) and transactions per second (TPS) of the database.
- Quick and reliable to convert local or remote physical machines into virtual machines without any disruption or downtime by using the VMware vSphere™ conversion tool. After the conversion, virtual machines can be rapidly deployed from the virtual machine template.
- Speedy deployment of Oracle DNFS clonedb virtual copies of a production database for Test/Dev, or other purposes.

- The thin provisioning capability offered by the combination of EMC SnapSure™ and the new Oracle DNFS clonedb feature reduces the amount of storage requirements for testing and development purposes. In addition, setting up a test database is virtually instantaneous because there is no need to copy the data files from the source database to the target database.

Introduction

Purpose

This white paper covers the following topics:

- Demonstrate the rapid provisioning of a physical Oracle 11g RAC environment to a virtual Test/Dev single-instance (RAC-collapsed) architecture.
- Demonstrate how protocol migration can be used to host Test/Dev, or other purpose instances in a NAS environment.

Scope

This white paper uses a hypothetical scenario that emulates real-world business needs as closely as possible. In the scenario, the Oracle 11g R2 RAC production environment is deployed on an EMC Celerra NS-960 and four physical servers.

This solution:

- Demonstrates the capabilities of EMC Celerra, EMC SnapSure, VMware ESX[®] 4.1 virtual infrastructure, and Oracle's DNFS clonedb feature, to build an application Test/Dev environment from an Oracle 11g R2 RAC production environment.
- Utilizes EMC Replication Manager to simplify the creation of Celerra SnapSure checkpoints. With Replication Manager, customers can automatically complete all of the steps to create Celerra SnapSure checkpoints for the file systems that contain database data files.

Audience

The intended audience for the white paper is:

- Internal EMC personnel
- EMC partners
- Customers

Terminology

Table 1 defines terms used in this document.

Table 1. Terminology

Term	Definition
Direct NFS Client (DNFS)	Direct NFS Client integrates the NFS client functionality in the Oracle software. It makes the task of configuring the Oracle database for NAS storage much simpler through the Direct NFS client instead of kernel NFS. The major benefits of the Direct NFS client include simplicity, easy administration, load balancing, high availability, and cost-effectiveness. The I/O performance is also better because Oracle can optimize the I/O code path by avoiding the kernel overhead.
DNFS clonedb	A new feature in Oracle 11g R2 that enables non-production databases such as Test/Dev to be created instantaneously without physically copying the production files. The files in the non-production databases are created based on copy-on-first-write technology.
Oracle Grid Infrastructure 11g	Oracle Grid Infrastructure is the software required to run the RAC environment. It provides basic clustering support at the operating system level and enables Oracle software to run in clustering mode. It also provides the cluster-ready services (CRS) layer that enables for cluster coherency, distributed lock management, and transparent node failover.
P2V	A process to convert a local or remote physical machine into a virtual machine on VMware vSphere.
RAC 11g	Real Application Cluster is software that enables customers to use clustered hardware by running multiple Oracle instances that access the same database.
Storage-based replication and recovery	<p>A solution component that provides replication functionality through the storage layer using specialized software and hardware.</p> <p>Storage-based replication provides the following benefits:</p> <ul style="list-style-type: none">• Offload the database server's central processing units (CPUs) from the I/O and processing requirements of the replication operations.• Superior Mean Time to Recovery (MTTR) through the use of logical storage layer replication (commonly referred to as snapshots).• Storage-based replication is faster than other ways of replication (for example, software-based replication)

Technology overview

Introduction

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

- EMC Celerra unified storage platform
- EMC Replication Manager
- EMC SnapSure
- VMware vCenter™ Converter
- VMware vSphere 4.1
- Storage architecture
- Oracle software stack

EMC Celerra unified storage platform

EMC Celerra unified storage platform products offer a flexible architecture and multiprotocol connectivity. This enables connectivity over IP/Ethernet, iSCSI, and Fibre Channel (FC) SAN environments.

EMC Celerra NS-960 is a unified storage system that brings advanced failover and Fully Automated Storage Tiering (FAST) to multiprotocol environments. With Celerra NS-960, customers can:

- Connect to multiple storage networks through NAS, iSCSI, FC SAN, and Celerra Multi-Path File System (MPFS). MPFS improves performance over traditional NAS.
- Move beyond direct-attached storage with an integrated EMC CLARiiON® CX4 that scales up to 960 disks.
- Simplify common administrative tasks with EMC Unisphere™ for greater efficiency with file system deduplication, FAST Cache, Virtual Provisioning™, and Celerra Automated Volume Management.

The key features provided by Celerra are described in Table 2.

Table 2. Celerra key features

Feature	Provided by
NAS	Network File System (NFS) and Common Internet File System (CIFS) protocols
iSCSI storage	Celerra Data Mover
FC storage	FC through the back-end EMC CLARiiON storage array

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. It auto-discovers the replication environment and enables streamlined management by scheduling, recording, and cataloging replica information including auto-expiration.

Specifically for Oracle database administrators, Replication Manager provides the following benefits, when combining with the EMC snapshot technologies (for example, EMC SnapSure):

- Accelerating backups with minimal disruption to production.
- Rapidly recovering Oracle databases for reduced recovery time objectives (RTO).
- Refreshing non-production environments for Test/QA development and reporting.
- Managing an increasingly complex infrastructure.

EMC Replication Manager consists of the following components:

- Replication Manager Server—Stores all information about users, hosts, replicas, and ongoing operations.
- Replication Manager Console—Provides a graphical user interface and command line interface for users to control Replication Manager's functions.
- Replication Manager Agent—Interacts with the application and storage layers to create, mount, restore, or expire replicas of mission-critical data.

EMC SnapSure

EMC SnapSure creates a logical point-in-time image (checkpoint) of a production file system (PFS) that reflects the state of the PFS at the point in time when the checkpoint is created. SnapSure can maintain a maximum of 96 read-only checkpoints and 16 writeable checkpoints per PFS while allowing PFS applications continued access to realtime data.

How SnapSure works

SnapSure uses a copy-on-first-modify principle. A PFS consists of blocks. When a block within the PFS is modified, a copy that contains the block's original contents is saved to a separate volume called the SavVol. Subsequent changes made to the same block in the PFS are not copied into the SavVol. The original block from the PFS in the SavVol and the unchanged PFS blocks remaining in the PFS are read by SnapSure according to a bitmap and blockmap data-tracking structure. These blocks combine to provide a complete point-in-time image that is called a checkpoint.

VMware vCenter Converter

VMware vCenter Converter is a tool used to quickly and reliably convert local and remote physical machines into virtual machines without any disruption or downtime. It can be used as an integrated version in a vSphere Client or as a standalone version.

VMware vCenter Converter can run on a wide variety of hardware and supports most commonly used versions of the Microsoft Windows and Linux operating systems.

VMware vCenter Converter is shown in Figure 1.

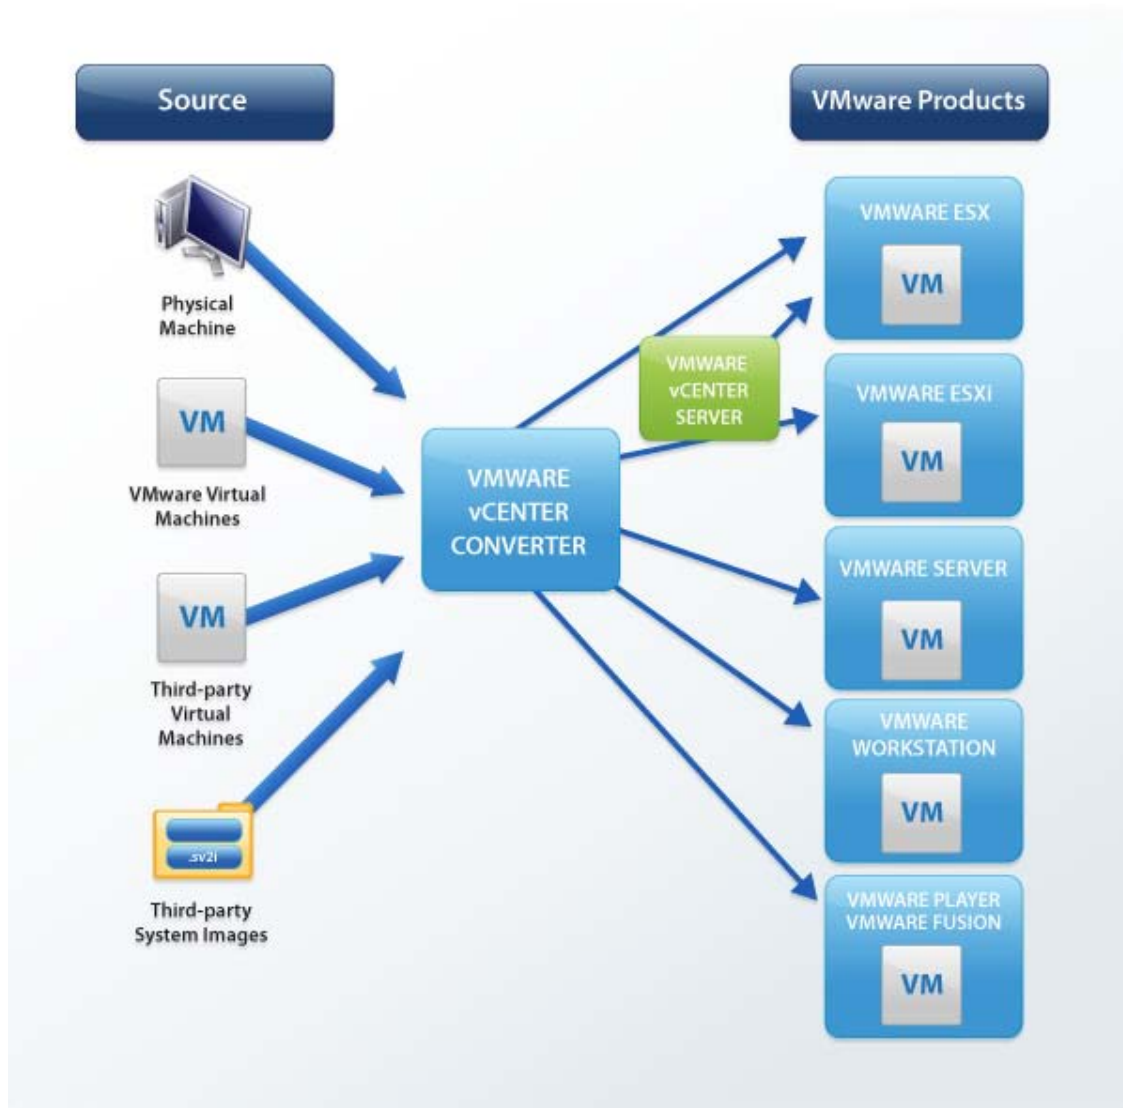


Figure 1. VMware vCenter Converter

VMware vSphere 4.1

VMware vSphere 4.1 is the most reliable virtualization platform that virtualizes servers, storage, and networking in the industry, allowing multiple unmodified operating systems and applications to run independently in virtual machines, while sharing physical resources. vSphere 4.1 enables higher consolidation ratios with unequaled performance by providing groundbreaking new memory-management technology and expanding its resource-pooling capabilities with new granular controls for storage and the network. vSphere 4.1 scales to an unmatched number of virtual machines and virtualized hosts to support the buildout of private and public clouds at even lower operational costs than before.

The benefits of running Oracle databases on VMware are listed below:

Performance

- I/O is not an issue
- Scale up and out
- Newer hardware can increase performance

Server consolidation

- Fully utilize hardware
- Maintain application isolation
- Scale dynamically and right-size infrastructure
- Reduce hardware cost

Rapid provisioning

- Streamline activation, deployment, and validation of servers
- Avoid manual configuration errors

Workload management

- Zero downtime maintenance
- Migrate live databases
- Distributed Resource Scheduler (DRS)

High availability (HA)

- VMware vMotion®
- Without clustering or RAC

Business continuity

- VMware Site Recovery Manager
- Hardware reduction at the failover site
- Comprehensive testing of a Disaster Recovery (DR) solution

Storage architecture

The general recommendations for mechanical disk drives are as follows:

- Drives with higher revolutions per minute (rpm) provide higher overall random-access throughput and shorter response time than drives with lower rpm. For optimum performance, higher-rpm drives are recommended for data files, temp files, and online redo log files. Because of significantly better performance, FC disks are always recommended for storing data files, temp files, and online redo log files.
- SATA II drives have slower response and rotational speed, and moderate performance with random I/O. However, they are less expensive than the FC disks for the same or similar capacity.
- SATA II drives are usually the most cost-effective option and generally provide sufficient performance for storing archived redo logs and the fast recovery area. In the

event of high-performance requirements for backup and recovery, FC disks can also be used for this purpose.

Oracle software stack

The Oracle software stack covered by this solution consists of:

- Oracle Grid Infrastructure 11g R2
- Oracle Database 11g R2 Enterprise Edition (11.2.0.2)
- Oracle DNFS
- Oracle DNFS clonedb

Configuration

Overview

The virtual infrastructure includes two ESX server hosts running VMware ESX Server 4.1. One ESX server is configured as the P2V conversion target, on which the converted virtual machines reside. The other ESX server is configured for the virtual machines. This server hosts the tools used in the processes of test and validation including Replication Manager, VMware vCenter Converter, and VMware vCenter Server. The production Oracle RAC environment includes four physical nodes. All database objects are stored on a Network File System (NFS) mount. Data files, temp files, control files, online redo log files, and archived logs are accessed using the DNFS protocol.

Physical environment

Figure 2 illustrates the overall physical architecture of the environment.

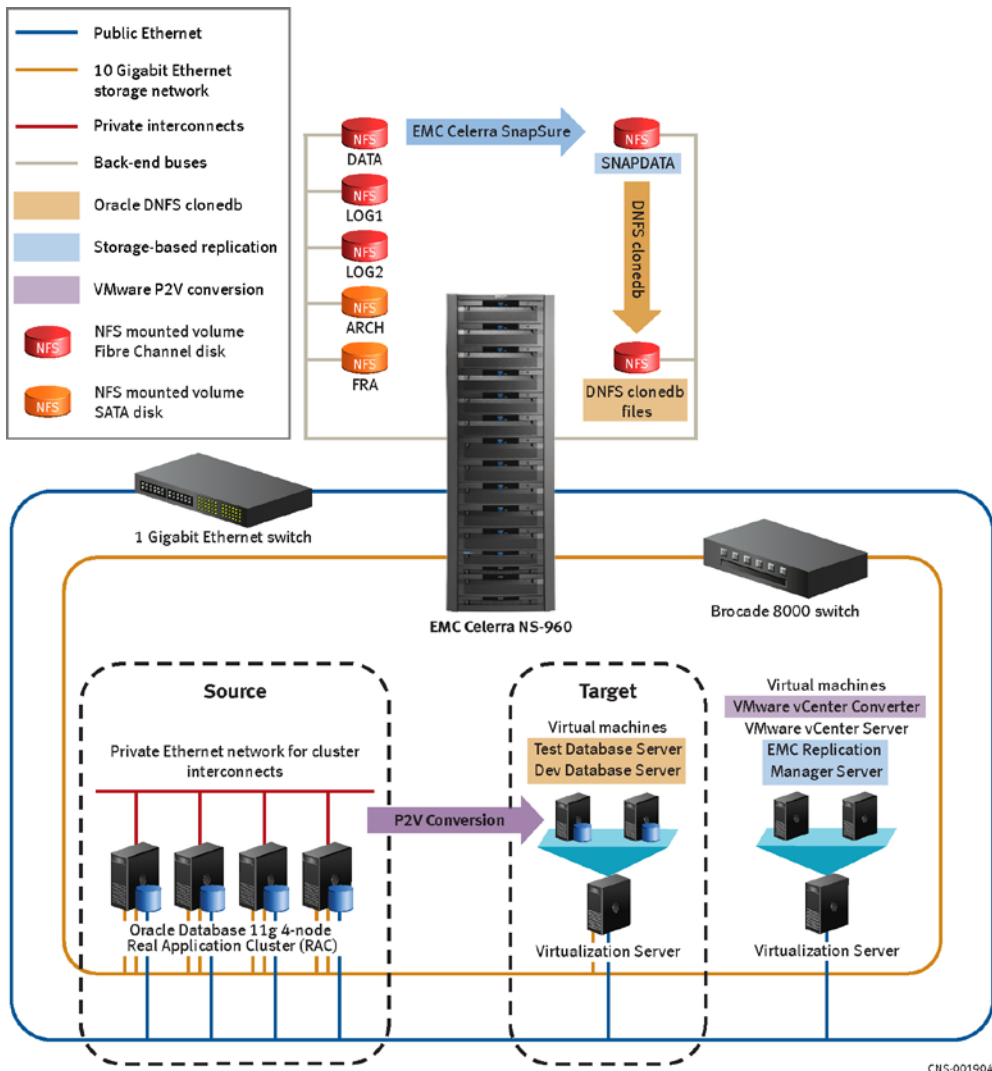


Figure 2. Physical architecture

- Celerra SnapSure: Celerra SnapSure is used to make a storage-based replication that is mounted to the virtual machines. For detailed information, refer to *Create a Celerra SnapSure checkpoint and perform P2V conversion*.
- P2V Conversion: The P2V Conversion is used to replicate the physical machine environment (including the operating system, the server configuration, and other software installed on the server) to the virtual machine. For detailed information, refer to *Create a Celerra SnapSure checkpoint and perform P2V conversion*.
- DNFS clonedb: The DNFS clonedb is used to make a DNFS clone based on the data files presented in the SnapSure checkpoint. For detailed information, refer to *Use Oracle DNFS clonedb to get the test database*.

Environment profile

The white paper is validated with the environment profile listed in Table 3.

Table 3. Environment profile

Profile	Quantity/Type/Size
Database size	1 TB
Database profile	OLTP
Network connectivity	10 Gb Ethernet
Network protocol	NFS over IP using Oracle DNFS client
User scaling	1000 - 5000

Server and network architecture

This solution design consists of two sites:

- The source: consists of a four-physical-node RAC database with a shared file system
- The target: consists of several virtual machines on VMware vSphere 4.1 for Test/Dev, QA, or other purposes.

RAC nodes are configured as follows:

- Each node is a physical machine.
- Each node has at least three network adapters:
 - One for the public network interface
 - One for the private network interface (interconnect)
 - One for the storage network interface
- The private connection between all of the nodes in a RAC cluster should be a dedicated network that is separated from the public network. The interfaces in a private network must support the user datagram protocol (UDP), by using high-speed network adapters and switches that support TCP/IP (Gigabit Ethernet or better is recommended).

- The interface names associated with the network adapters for each network must be the same on all nodes in RAC.
- DNFS provides file system semantics for Oracle Database and RAC 11g.
- The following IP addresses must be considered for every node in RAC:
 - An IP address and associated hostname are registered in the DNS or the /etc/hosts file for each public network interface.
 - One unused virtual IP (VIP) address and an associated VIP name, registered in the DNS or the /etc/hosts file, which are configured for the primary public network interface, are needed for each node. The virtual IP address must be in the same subnet as the associated public interface. After installation, the clients (the application tier nodes) can be configured to use the VIP name or VIP address. If a node fails, its virtual IP address fails over to another node.
 - For the private IP address and optional hostname for each private interface, Oracle recommends that private network IP addresses are used for these interfaces.

Table 4 lists the hardware and software components of the four-node RAC database.

Table 4. RAC database components

Hardware (four machines)	
Component	Quantity
2.66 GHz quad-core processors	4
Memory	128 GB
10 GbE Ethernet NICs	2
1 GbE Ethernet NICs	3
84 GB 15k internal SCSI disk	1
42 GB 15k internal SCSI disk	1
Software	Version
Oracle Enterprise Linux	5.5
Oracle Grid Infrastructure	11.2.0.2
Oracle Database Enterprise Edition	11.2.0.2
Replication Manager Agent	5.2.3.0

Two VMware ESX server nodes are configured. Table 5 shows the hardware and software configuration of the VMware ESX server.

Table 5. VMware ESX server components

Hardware (two machines)	
Component	Quantity and description
2.66 GHz quad-core processor	1
Memory	24 GB
1 GbE Ethernet NICs	2
10 GbE Ethernet NIC	1
84 GB 15k internal SCSI disk	1
42 GB 15k internal SCSI disk	1
Software	Version
VMware ESX Server	4.1

Table 6 shows the hardware and software configuration of the vCenter Server virtual machine.

Table 6. vCenter Server virtual machine components

VMware virtual machine	
Virtual hardware	
Component	Quantity
2.66 GHz single-core processor	1
Virtual memory	4 GB
1 GbE Ethernet virtual NIC	1 (used for public network)
25 GB 15k internal SCSI virtual disk	1 (used for operating system)
Software	Version
Microsoft Windows Server 2008 R2 Enterprise Edition	2008 R2
VMware vCenter Server	4.1.0
vSphere Client	4.1.0

Table 7 shows the hardware and software configuration of the Test/Dev virtual machine.

Table 7. Test/Dev virtual machine components

VMware virtual machine	
Virtual hardware	
Component	Quantity
2.66 GHz single-core processors	2
Virtual memory	4 GB
1 GbE Ethernet virtual NIC	1 (used for public network)
10 GbE Ethernet virtual NIC	1 (used for storage network)
80 GB 15k internal SCSI virtual disk	1 (used for OS)
Software	Version
Oracle Enterprise Linux	5.5
Oracle database software	11.2.0.2

Table 8 shows the hardware and software configuration of the VMware virtual machine.

Table 8. Replication Manager and VMware vCenter Converter virtual machine components

VMware virtual machine	
Virtual hardware	
Component	Quantity
2.66 GHz single-core processor	1
Virtual memory	512 MB
1 GbE Ethernet virtual NICs	2
20 GB 15k internal SCSI virtual disk	1 (used for operating system)
Software	Version
Microsoft Windows Server 2003 R2 Enterprise Edition	2003 R2
EMC Replication Manager Server	5.3.0.0
EMC Replication Manager Console	5.3.0.0
VMware vCenter Converter Standalone	4.1

Storage architecture

In this use case, Celerra NS-960 arrays are used for both the production and Test/Dev environments. Table 9 and Table 10 describe the hardware and software resources of the storage array.

Table 9. Hardware resources

Equipment	Quantity	Configuration
EMC Celerra NS-960 unified storage platform (include an EMC CLARiiON CX4-960 back-end storage array)	1	<ul style="list-style-type: none"> • 2 storage processors • 3 Data Movers • 1 Control Station • 2 x 10 GbE network connections per Data Mover • 7 FC shelves • 2 SATA shelves • 105 x 300 GB 15k FC disks • 30 x 1 TB SATA disks
10 Gigabit Ethernet switch	1	<ul style="list-style-type: none"> • 24 CEE ports

Table 10. Software resources

Software	Version
EMC Celerra Manager Advanced Edition	5.6
EMC Navisphere® Agent	6.29.5.0.37
EMC FLARE®	04.29.000.5.006
EMC DART	5.6.49-3
EMC Navisphere Management Suite	6.29

Table 11 lists the Oracle database file systems allocation in this solution.

Table 11. Oracle file system allocation on Celerra

What	File system	File system type
OCR file	/crs	RAID-protected NFS file system on FC disks
Voting disk file		
Oracle data files	/datafs1, /datafs2	RAID-protected NFS file system on FC disks

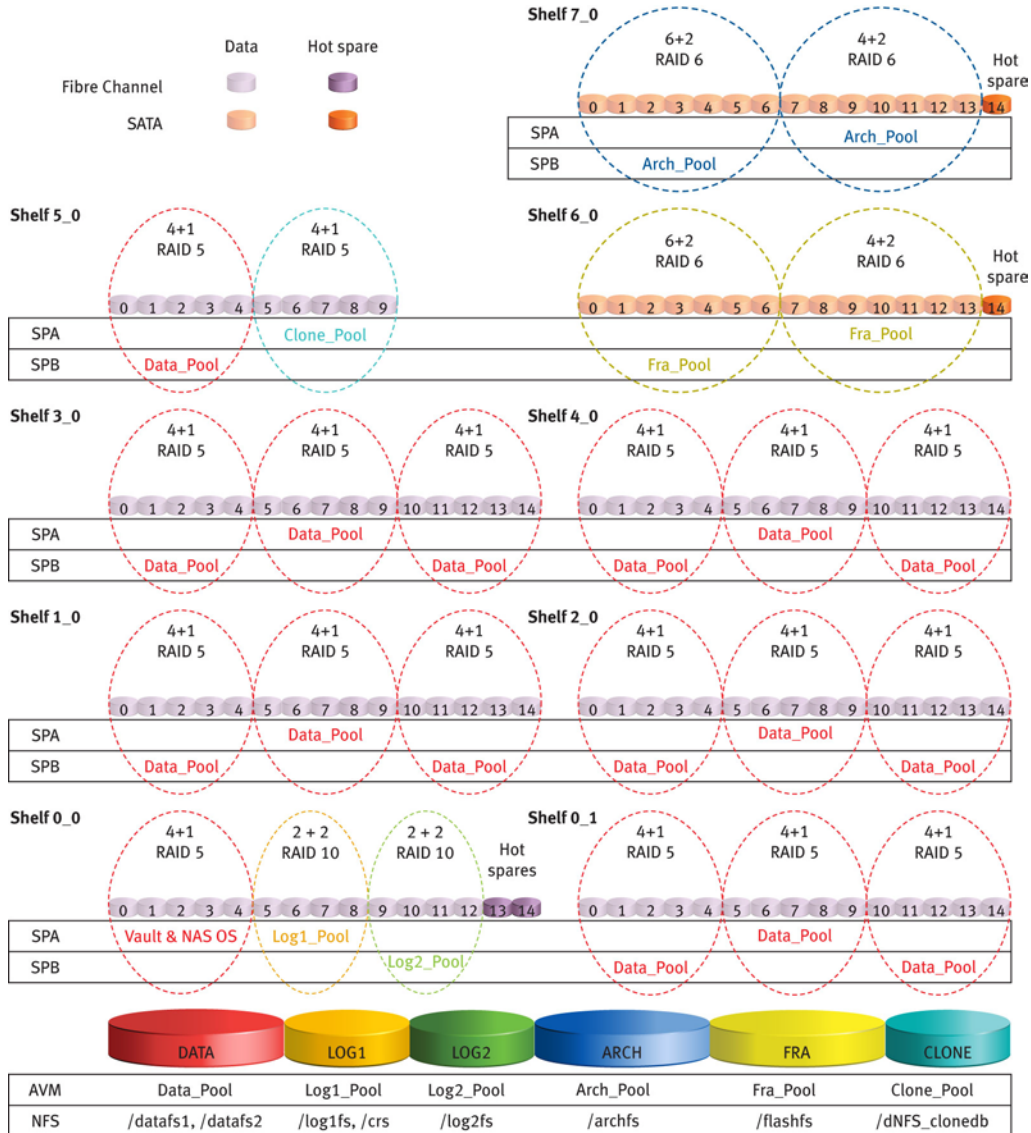
Oracle temp files		
Oracle control files		
Oracle online redo log files	/log1fs, /log2fs	RAID-protected NFS file system on FC disks
Oracle archived log	/archfs	RAID-protected NFS file system on SATA disks

Oracle data files, temp files, online redo log files, control files, and archived log files are accessed using DNFS. The OCR file and the voting disk file are accessed using the Kernel NFS file system.

Oracle data files, temp files, and online redo log files are contained in NFS file systems stored on RAID 5 on FC disks. Archived log files are stored on RAID 6 on SATA disks.

Disk layout

Figure 3 illustrates the disk layout of the environment.



CNS-001897

Figure 3. Disk layout

Test and validation

Introduction

The following list describes the high-level project implementation steps, including Oracle RAC database configuration, migration, and Test/Dev environment provisioning:

- Configure the EMC Celerra NS-960 storage array
- Complete VMware vSphere 4.1 tasks
- Install an Oracle 11g R2 RAC database
- Enable DNFS (Direct NFS)
- Create a baseline performance
- Create a Celerra SnapSure checkpoint and perform P2V conversion
- Use Oracle DNFS clonedb to get the test database
- Provision multiple virtual machines for Test/Dev from the virtual machine template

Configure the EMC Celerra NS-960 storage array

Check the following steps to configure EMC Celerra NFS for EMC Replication Manager.

Step	Action
1	Confirm all mount points are exported to the production host.
2	Confirm all mount points are added to the <code>/etc/fstab</code> entries.
3	Confirm all file systems are mounted with the <code>/etc/fstab</code> entries.

Complete VMware vSphere 4.1 tasks

VMware vSphere 4.1 configuration includes VMware ESX Server 4.1 installation, VMware vCenter 4.1 installation and configuration, and the creation of two virtual machines. One virtual machine is used to host Replication Manager and VMware vCenter Converter; the other one is used to host the VMware vCenter server for centralizing the management of virtual machines.

Note Refer to the relevant VMware vSphere installation documents for detailed information about installing VMware ESX Server and vCenter. Refer to the system administrator documents for how to configure VMware ESX Server and how to create a virtual machine by vSphere Client.

Install an Oracle 11g R2 RAC database

Follow the steps provided to install an Oracle 11g R2 RAC database.

Step	Action
1	Using Oracle Universal Installer to install Oracle RAC
2	Creating Oracle Real Application Clusters Databases with Database Configuration Assistant

Refer to *Storage architecture* for more details about the file systems.

Refer to *Oracle Real Application Clusters Installation Guide 11g Release 2 (11.2) for Linux and UNIX* for the detailed steps.

Enable DNFS

To enable Oracle Database 11g to use DNFS, the NFS file systems must be mounted and available over regular NFS mounts before starting the installation. The mount options used in mounting the file systems are not relevant because DNFS manages the settings after the installation. The following steps describe how to enable DNFS.

Step	Action
1	Create an oranfstab under \$ORACLE_HOME/dbs. Example: <pre>server: NAS-960 path: 192.168.4.190 export: /data_fs mount: /u02 export: /log1_fs mount: /u03 export: /log2_fs mount: /u04 export: /arch_fs mount: /u06</pre>
2	Oracle Database uses an ODM library, libnfsodm11.so, to enable DNFS. The standard ODM library \$ORACLE_HOME/lib/libodm11 should be replaced by the ODM NFS library: libnfsodm11. Change the directory to \$ORACLE_HOME/lib, and run the following commands: <pre>cp libodm11.so libodm11.so_stub ln -s libnfsodm11.so libodm11.so</pre>

Create a baseline performance

With archived log enabled, a baseline performance test is conducted on the Oracle RAC database for further comparison as shown in Figure 4.

Step	Action
1	Initiate the Benchmark Factory console and agents on the client machines.
2	Start the Benchmark Factory job.
3	Monitor the progress of the test, and capture the results after the test is completed.

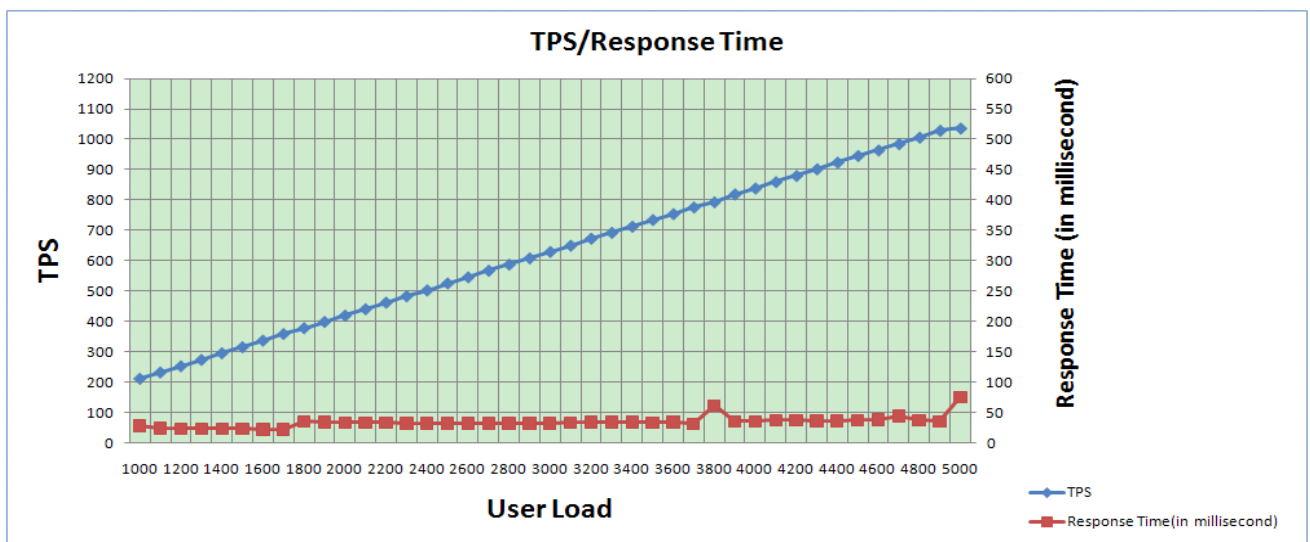


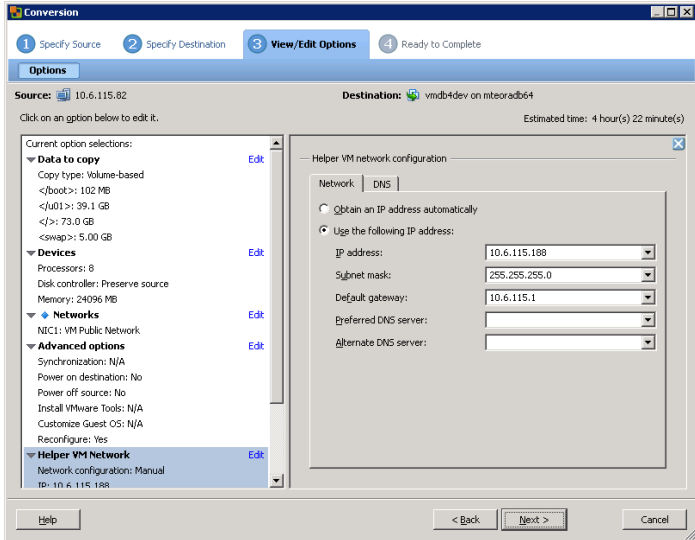
Figure 4. Baseline performance

Create a Celerra SnapSure checkpoint and perform P2V conversion

IOPS/TPS metrics for the database are gathered when creating a Celerra SnapSure checkpoint of all the relevant data files by Replication Manager, as shown in Figure 6.

The performance impacts during the P2V conversion are also captured in Figure 6.

Step	Action
1	Configure Replication Manager.
2	Register the production hosts, mount hosts, and storage in Replication Manager.
3	Create the application set in Replication Manager for the database to be replicated.
4	Create a job with the hot backup mode setting in the Replication Manager console to create the SnapSure checkpoint.

	<p>Note It is not always necessary to use hot backup mode. Refer to the <i>EMC Backup and Recovery for Oracle Database 11g Without Hot Backup Mode using DNFS and Automatic Storage Management on Fibre Channel</i> white paper.</p>
5	Start the Benchmark Factory test with a user load ranging from 1,000 to 5,000.
6	When the user load reaches an iteration of 2,000, create a SnapSure checkpoint for the database by running the job in the Replication Manager console.
7	Monitor the performance impact on the production database when Replication Manager is creating the SnapSure checkpoint.
8	After the checkpoint is created, specify the physical machine as the source in the VMware vCenter Converter console.
9	In the vCenter console, specify the VMware ESX server node that holds the converted virtual machine.
10	<p>Edit VMware Network configuration and assign an available IP address to the virtual machine (temporarily).</p> 
	<p>Figure 5. VMware P2V conversion</p>
11	When the user load reaches an iteration of 3,000, a P2V conversion is performed.
12	Monitor the performance impact of P2V conversion on the production database.
13	After the conversion is completed, start the virtual machine and assign a new IP address and hostname.

Performance impacts

The TPS and response time are captured when running an OLTP workload on the database. Figure 6 shows the performance of the production RAC database when the SnapSure checkpoint is created and the P2V conversion is performed.




Figure 6. Performance during SnapSure checkpoint creation and P2V conversion

The SnapSure checkpoint is initiated with the database hot backup mode enabled at the iteration of user load 2,000. A brief increase in response time and a moderate decrease of TPS are observed during the iteration of user load 2,100. This is because when putting the database into hot backup mode, additional redo data is logged, a checkpoint is initiated, and the Database Writer (DBWR) process is forced to flush buffer cache into the disk. Moreover, before the SnapSure checkpoint is created, the database is forced to archive the current online redo log.

The P2V conversion is initiated at the iteration of user load 3,000 and completed at the iteration of user load 3,600. No performance impact is observed during the P2V conversion, since the converter process only copies the source node's local disks or volumes to the destination virtual machine, it does not affect the Oracle file system located on Celerra.

Figure 7, Figure 8 , Figure 9 , and Figure 10 are captured from the Oracle database control during the processes of creating a SnapSure checkpoint and making the P2V conversion.

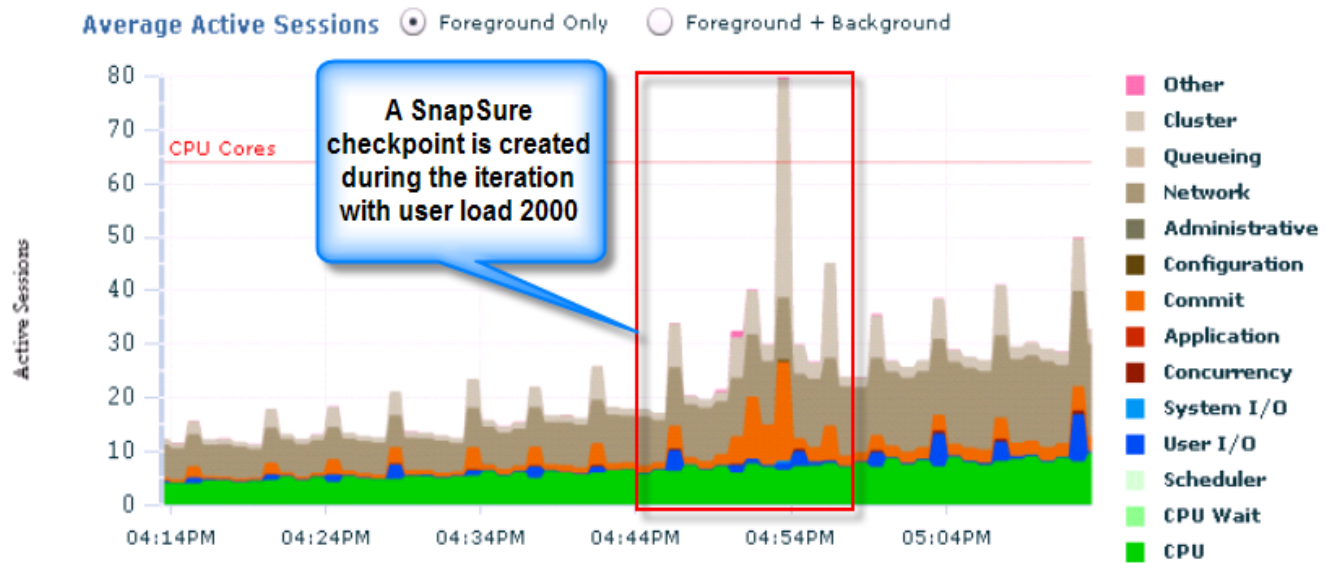


Figure 7. Average Active Sessions when creating a SnapSure checkpoint

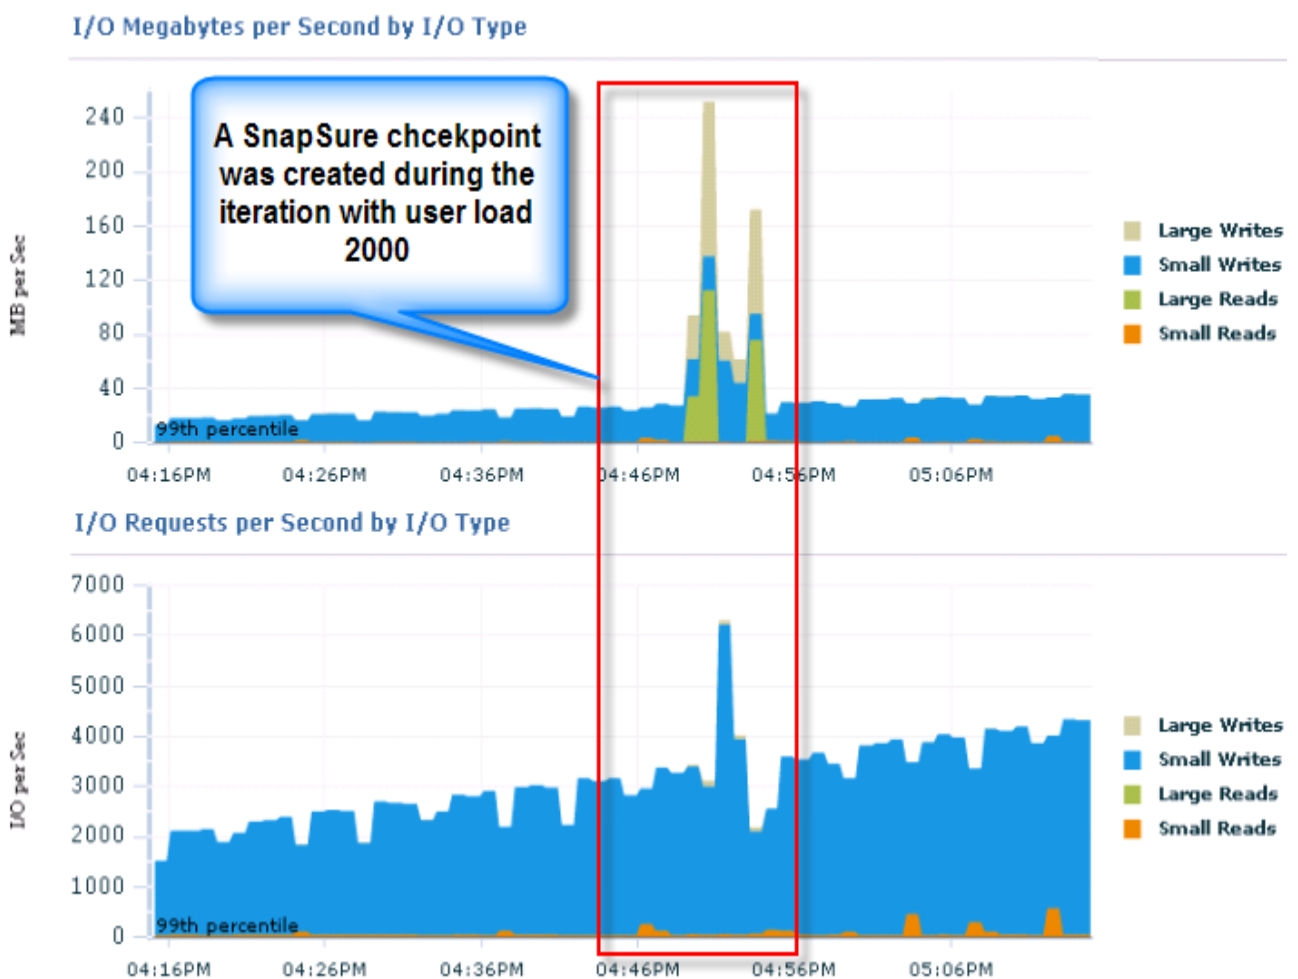


Figure 8. I/O traffic when creating a SnapSure checkpoint

Figure 7 and Figure 8 show the details of how the database is impacted when creating the SnapSure checkpoint.




Figure 9. Average Active Sessions when performing a P2V conversion

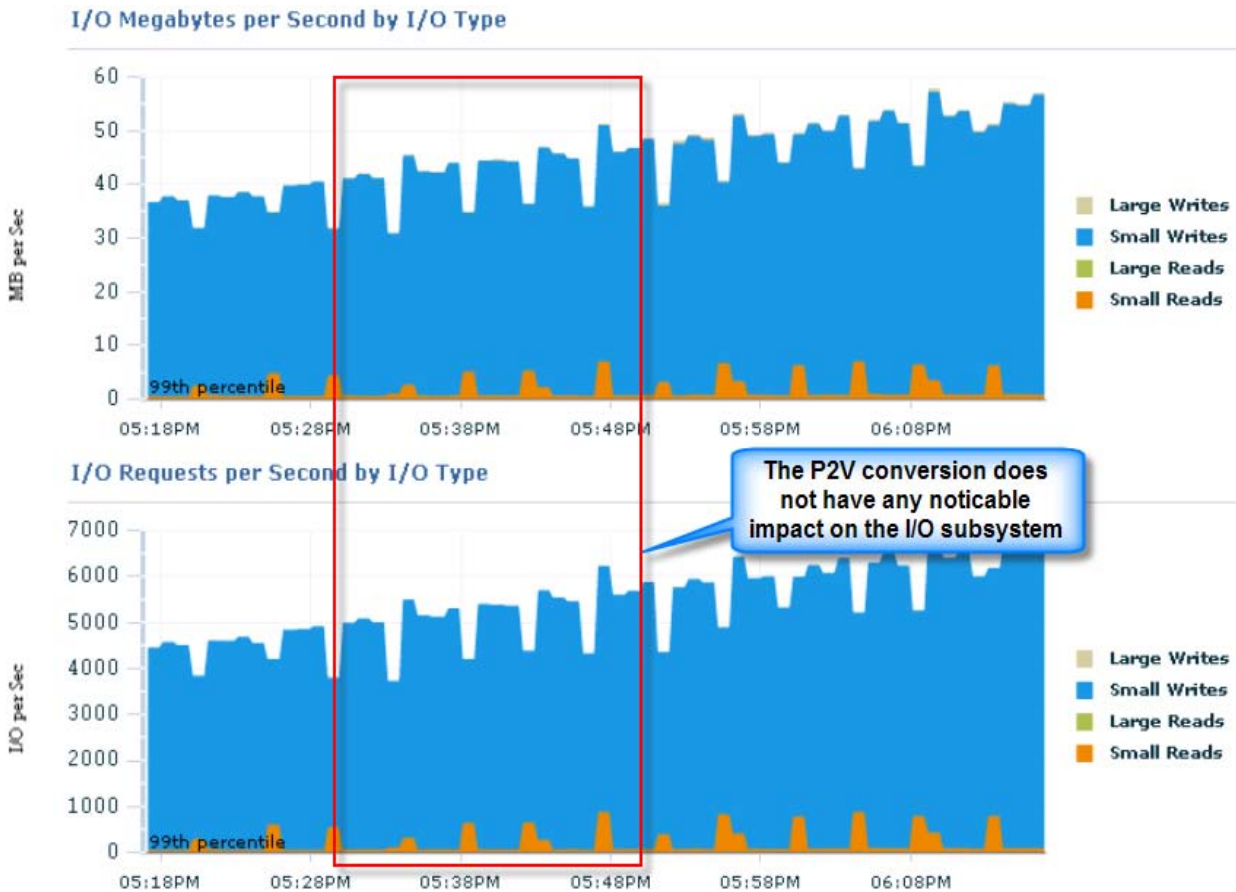


Figure 10. I/O traffic when performing a P2V conversion

Figure 9 shows that there is almost no impact on the database when making the P2V conversion. It is reasonable that the number of average active sessions increase because the load is increased steadily by using the load generate tool during the test.

Figure 10 shows that there is no impact on the I/O subsystem during the P2V conversion process.

Use Oracle DNFS clonedb to get the test database

The following steps describe how to configure and perform the DNFS clonedb after the physical RAC server is converted into a virtual machine.

Step	Action
1	<p>Recompile the ORACLE_HOME with the “rac off” option on the virtual machine created from the P2V conversion.</p> <p>Note The virtual machine used for testing is created by P2V from the production RAC environment. It contains the same CRS_HOME and ORACLE_HOME as the RAC physical machine. The virtual machine is used for a single-instance database, so the oracle binary files must be re-compiled with the “rac off” option.</p>
2	<p>Reconfigure the DNFS in the virtualized single-instance test database by editing the DNFS configuration file oranfstab under \$ORACLE_HOME/dbs.</p> <p>Note For the detailed steps, refer to <i>Enable DNFS</i>.</p>
3	Mount the Celerra SnapSure checkpoint to the target virtual database server.
4	Edit the initialization parameter files with the correct parameter setting.
5	Start up the instance with the “nomount” option.
6	<p>Re-create the control files.</p> <p>This can be accomplished by running the following command in the source database:</p> <pre>alter database backup controlfile to trace;</pre> <p>Then run the create controlfile command (which can be found in the trace file generated by the above command) against the target database.</p> <p>Note The name of the database can be changed to another name that is different from the source database name if needed.</p>
7	<p>In the new target database environment, perform DNFS clonedb by running the following command for each copy of data file in the test database:</p> <pre>dbms_DNFS.clonedb_renamefile(backup_ file_name, new_data_filename).</pre> <p>For example:</p> <pre>exec</pre>

	<pre>dbms_DNFS.clonedb_renamefile('/u04/oradata/mterac28/system.260.725492807.dbf', '/clonedb/uc491tst/system.260.725492807.dbf');</pre>
8	<p>Recover the database with the following command in SQLPLUS:</p> <pre>SQL> Recover database using backup control file until cancel;</pre> <p>When prompted, apply the appropriate archived logs.</p>
9	<p>Open the database:</p> <pre>SQL> alter database open resetlogs;</pre>

The test database is up and running after these steps, so read and write to the test database activities can be operated. When the workload running is started, storage consumption of the test database grows with the speed at which the data is modified.

In order to validate whether there is any performance impact on the production database, a TPC-C workload runs on the production database during the processes of conducting the DNFS clone.

Figure 11 shows the performance when two DNFS clonedb procedures are conducted at the iteration of user load 2,000 and 3,000 respectively, to deploy the database in the virtualized environment for Test/Dev purposes. Each procedure takes 10 seconds.

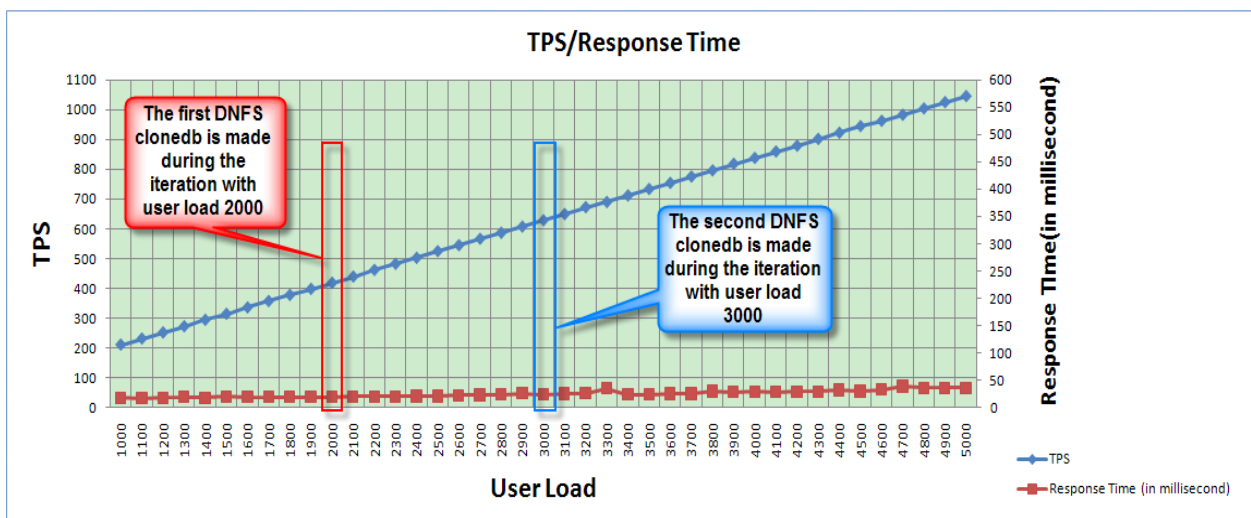


Figure 11. Performance during two DNFS clonedb procedures

Figure 12 looks exactly the same as the baseline performance diagram, which indicates that there is no performance impact on the production database when conducting DNFS clonedb.




Figure 12. Space consumption by data files

Figure 12 compares disk space used by the production database and test database on the virtual machine. After performing DNFS clonedb, the test database uses only 7.4 MB disk space. This is because production data files are not copied into the new test database environment; only the blocks that are updated in the test database are written to the disk.

Provision multiple virtual machines for Test/Dev from the virtual machine template

The following steps indicate how to provision multiple virtual machines for Test/Dev.

Step	Action
1	Create a virtual machine template from the virtualized test database server created by P2V.
2	Create new virtualized database servers from the virtual machine template.
3	<p>Create new writeable databases using the 11gR2 DNFS clonedb feature:</p> <ul style="list-style-type: none"> Mount the Celerra SnapSure checkpoint to the target virtual database server created in step 2. Clone the database on the target virtual database server with the DNFS clonedb feature. <p>For the detailed steps, refer to <i>Use Oracle DNFS clonedb to get the test database</i>.</p>

At this point, multiple virtual machines and the databases are available for use. When it is time to remove the Test/Dev environment, all files in those environments can be deleted without any impact on the production environment.

Note Refer to the *VMware System Administration Guide* for detailed information about how to use a virtual machine template.

Test summary

Table 12 summarizes the test cost time.

Table 12. Time cost summary

Item	Value
User load range	1,000 to 5,000 with intervals of 100
Total time for SnapSure checkpoint	7 minutes
Total time for P2V conversion	17 minutes
Total time for DNFS clonedb	10 seconds
Total time for virtual machine template creation	12 minutes
Total time for deploying virtual machine from template	11 minutes

Conclusion

Summary

This white paper highlights that EMC Celerra unified storage platforms can offer great flexibility, a highly competitive entry price, and multiple ways to replicate the production environment for Test/Dev, QA, and other purposes. Furthermore, this white paper demonstrates that the new feature—DNFS clonedb in Oracle database 11g R2—works well with EMC Celerra SnapSure checkpoint technology.

With the combination of EMC Celerra unified storage platform, Replication Manager, VMware virtualization technologies, and Oracle DNFS clonedb feature, users can replicate their production environments for Test/Dev in less than 30 minutes, offering almost the immediate ability to use the newly provisioned virtual instances.

Findings

This solution addresses the following key points:

- Centralizing all non-production environments on EMC Celerra unified storage platforms offers not only great flexibility but also a highly competitive entry price.
- By using EMC Celerra SnapSure technology with Replication Manager, together with Oracle DNFS cloned and VMware virtualization technology, the time taken to clone a production database is greatly reduced. Moreover, the disk space required by the cloned database is reduced to minimal because only the modified data after cloning is stored.
- With the EMC SnapSure checkpoint technology, Oracle DNFS clonedb feature, VMware vCenter Converter Standalone tool, and virtual machine template, a production database environment can be cloned within one hour, no matter the size of the database. This clone environment can be used to seed a Test/Dev environment within 30 minutes, which greatly improves the efficiency of the migration process.
- Compared with the traditional way of database cloning/migrating (such as Oracle RMAN, Oracle import/export), this solution minimizes the impact on performance of the production database and reduces the storage requirement.

References

White paper

For a white paper that provides similar solutions, see the following documentation:

- *EMC Unified Storage for Oracle E-Business Suite Applications—Virtual Test/Development Provisioning Enabled by EMC Celerra and VMware vSphere on NFS/IP—A Detailed Review*

Oracle Metalink Note

- *Clone your DNFS Production Database for Testing [ID 1210656.1] Doc 1*

Oracle online document

- *Oracle Real Application Clusters Installation Guide 11g Release 2 (11.2) for Linux and UNIX*

VMware online documents

- *vSphere Basic System Administration*
- *ESX and vCenter Server Installation Guide*