

ORACLE DATABASE CLONING MADE SIMPLE, FAST, AND STORAGE-EFFICIENT WITH EMC SNAPSURE AND ORACLE DIRECT NFS CLIENT CLONEDB

A Detailed Review

EMC GLOBAL SOLUTIONS

Abstract

This white paper demonstrates the combination of EMC® snapshots and the new Oracle Database 11g Release 2 feature, Direct NFS Client clonedb. Through the use of these technologies, the EMC Celerra® NS-960 unified storage platform provides great flexibility, high performance, and a highly competitive total cost of ownership (TCO) for replicating an Oracle production database to quickly and easily create test/dev and quality assurance (QA) environments.

June 2011



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.

All trademarks used herein are the property of their respective owners.

Oracle may have substituted components of the original environment in this document with hardware of a similar (or higher) specification to the original equipment used in the EMC Proven Solution. The content contained in this document originated from a validated EMC Proven Solution. The modification introduced by Oracle may have caused changes in performance, functionality, or scalability of the original solution. Please refer to www.emc.com/solutions for further information on validated EMC Proven Solutions.

Part Number H8728

Table of contents

Executive summary	5
Business case.....	5
Solution overview	5
Key results	5
Introduction	6
Purpose	6
Scope	6
Audience	6
Terminology	7
Technology overview	8
Overview.....	8
EMC Celerra unified storage platform	8
EMC Replication Manager	8
EMC Celerra SnapSure	9
How SnapSure works.....	9
Oracle software stack.....	9
Oracle RAC 11g configuration	10
Overview.....	10
Physical environment.....	10
Environment profile.....	11
Server and network architecture.....	11
Storage architecture.....	13
Disk layout.....	14
Testing and validation of the RAC scenario.....	15
Overview.....	15
Configuring EMC Celerra NS-960 storage array	15
Oracle RAC 11g Release 2 installation and database creation	15
Enable dNFS	15
Baseline performance creation.....	16
Creation of Celerra SnapSure checkpoint for the database file system	16
Performance impact.....	17
Using Oracle dNFS clonedb to create the test/dev database with the SnapSure checkpoint as the source.....	20
Oracle Database 11g single-instance configuration.....	24
Overview.....	24

Physical environment.....	24
Environment profile, software, hardware, and so on.....	24
Testing and validation of the single-instance scenario.....	25
Performance workload against dNFS clonedb database	25
Conclusion	26
Summary	26
Findings.....	26
Next steps.....	26
References.....	27
White papers	27
Oracle Metalink Note	27
Oracle online document.....	27

Executive summary

Business case One of the most challenging tasks for Oracle database administrators is to create, deploy, and manage a large number of images of a production environment. Some of those tasks include deploying these images for testing, development, reporting, and so on. Administrators often have to perform these activities frequently and quickly to meet their internal service level agreements (SLAs).

In addition, cost reduction and resource optimization are key requirements for today's companies. Many customers need the capability to quickly and easily provision testing and development (test/dev), quality assurance (QA), or other non-production environments, by leveraging a virtual copy of the production database.

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 test/dev, QA, and other purposes. When combined with powerful tools such as EMC Replication Manager (RM) and Oracle Direct NFS Client (dNFS) clonedb, these platforms can offer a very compelling argument for hosting the non-production database environment with dNFS.

This solution demonstrates:

- How customers can easily create an image of an Oracle production database using network-attached storage (NAS) through the use of EMC snapshot technology.
- How customers can use Oracle's dNFS clonedb feature to clone the production database instantly and provision that clone for test/dev, QA, and other purposes.
- How dNFS clonedb provides a "thinly provisioned" copy of the production database. A "thinly provisioned" copy is a copy that uses minimal space; it consists largely of the changed blocks between the clone and the production database.

This type of "thin provisioning" reduces the amount of storage required for test/dev purposes. In addition, setting up a test/dev database is instantaneous because there is no need to copy the data files from the source database to the target database.

Key results

This white paper demonstrates the following benefits of the solution:

- The low impact of storage-based replication (snapshots), as well as dNFS clonedb, on database transactions per second (TPS) using an industry-standard OLTP workload.
- The speed with which a copy of a production Oracle database can be created using the dNFS clonedb feature for test/dev, QA, or other purposes.
- The low impact of dNFS clonedb in terms of storage space.
- The low impact of dNFS clonedb in terms of performance of the clonedb database itself.

Introduction

Purpose

This white paper covers the following topics:

- Demonstrate rapid provisioning of a virtual copy of an Oracle RAC 11g production environment using EMC Celerra SnapSure™ checkpoint.
- Demonstrate how dNFS clonedb can be used to make multiple “thinly provisioned” writeable copies of the production Oracle RAC 11g environment. These writable copies use both EMC Celerra SnapSure checkpoint and the Oracle RAC 11g production environment with dNFS clonedb.

Scope

The scenario in this white paper emulates the real-world business needs as closely as possible. In this scenario, the Oracle RAC 11g Release 2 production environment is deployed on EMC Celerra NS-960. The database is accessed using Oracle dNFS.

This solution demonstrates the capabilities of EMC hardware and software, including Celerra SnapSure, together with dNFS clonedb, to build a test/dev environment from an Oracle RAC 11g Release 2 production environment. This test/dev environment is mounted to an Oracle Database 11g Release 2 single instance server.

The Celerra SnapSure checkpoint creates a logical point-in-time image of a production file system (PFS). Celerra SnapSure checkpoints can be created on a local or remote Celerra storage array, and then exported and mounted to client computers. These checkpoints can be either read-only or read/write, in a manner similar to normal file systems.

To enhance ease of use, this solution uses RM to manage the creation of Celerra SnapSure checkpoints. With RM, customers can automatically complete all 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

Table 1 defines terms used in this document.

Table 1. Terminology

Term	Definition
Oracle dNFS	Oracle dNFS is the integration of the NFS client functionality into the Oracle RDBMS server. The task of configuring Oracle database for NAS storage becomes much simpler when using dNFS instead of kernel NFS. The major benefits of dNFS include simplicity, ease of administration, load balancing, high availability and cost-effectiveness. The I/O performance is also better because Oracle has optimized the I/O code path by reducing kernel overhead significantly.
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 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).
Oracle dNFS clonedb	A feature introduced in Oracle Database 11g Release 2 that allows non-production databases such as test/dev databases to be created instantly without physically copying the production files. The files in the non-production databases are created based on a “copy-on-first-write” or a “thin provisioned” technology.

Technology overview

Overview

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

- EMC Celerra unified storage platform
- EMC Replication Manager
- EMC Celerra SnapSure
- Oracle software stack, including:
 - Oracle RAC 11gRelease2
 - Oracle Database 11gRelease 2
 - dNFS
 - clonedb

EMC Celerra unified storage platform

The EMC Celerra unified storage platform product family offers a flexible architecture and multi-protocol connectivity. This enables connectivity over IP/Ethernet, iSCSI, and Fibre Channel (FC) SAN environments.

Celerra NS-960 is a unified storage system that brings advanced failover and FAST to multi-protocol 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	NFS and Common Internet File System (CIFS) protocol
iSCSI storage	Celerra Data Mover
FC storage	FC through the back-end EMC CLARiiON storage array

EMC Replication Manager

RM manages EMC point-in-time replication technologies through a centralized management console. RM 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.

Customers can use RM to:

- Create replicas using a variety of technologies to suit customer requirements, either on demand or based on schedules and policies defined by the customer.
- Simplify replica management with application consistency.

RM consists of the following components:

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

EMC Celerra SnapSure

SnapSure creates a logical point-in-time image (checkpoint) of a 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 writable checkpoints per PFS while allowing PFS applications continued access to real-time 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.

Oracle software stack

The Oracle software stack covered by this solution consists of:

- Oracle RAC 11g Release 2 Enterprise Edition (11.2.0.2), which includes Oracle Database 11g Release 2 Grid Infrastructure and Oracle Cluster Ready Services (CRS)
- dNFS
- dNFS clonedb

Oracle RAC 11g configuration

Overview

The Oracle RAC 11g testing scenario includes:

- An Oracle RAC 11g cluster consisting of four nodes running the production database
- A single instance database server running Oracle Database 11g Release 2 for hosting the test/dev database

In both environments, all database objects are stored on an NFS mount. Data files, temp files, control files, online redo log files, and archived logs are accessed using dNFS.

Physical environment

Figure 1 shows the overall physical architecture of the Oracle RAC 11g environment.

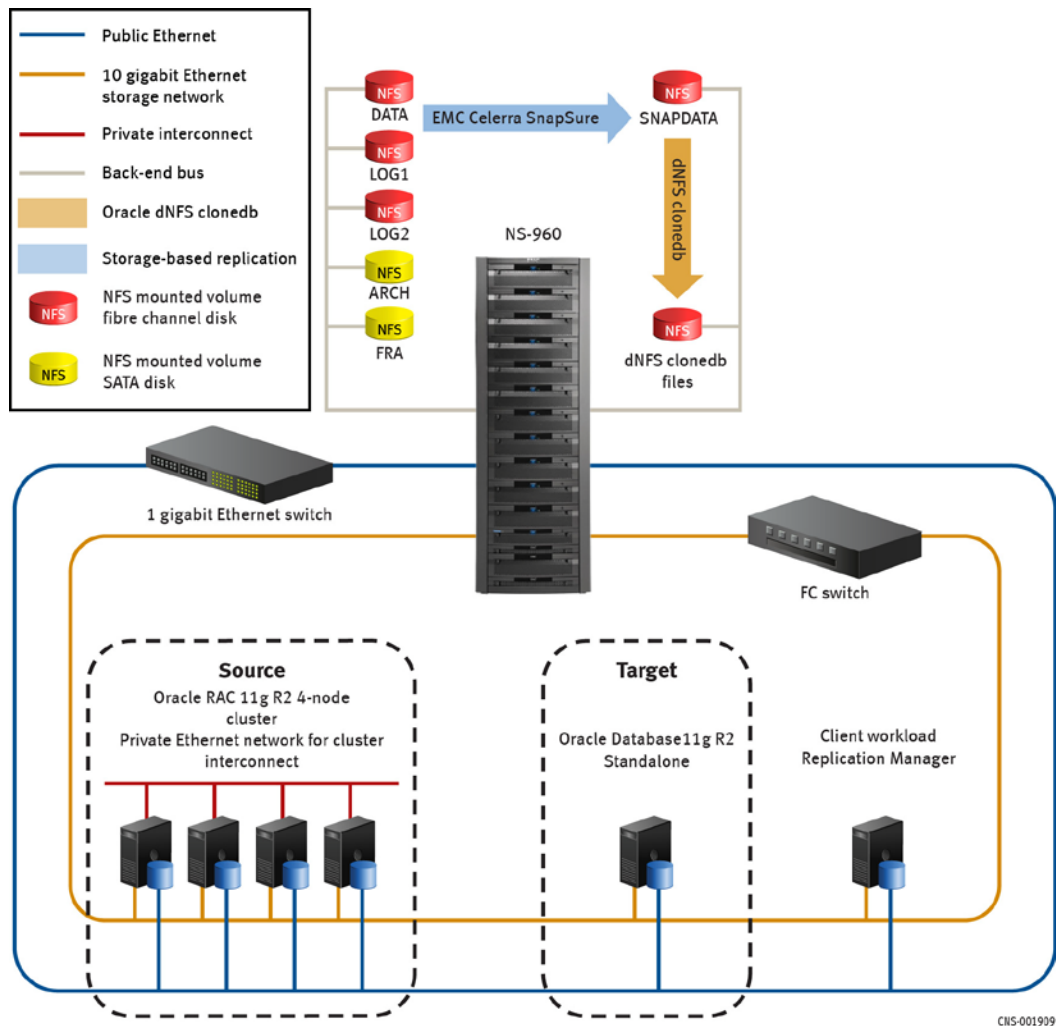


Figure 1. Physical architecture of the Oracle RAC 11g testing scenario

Environment profile

The RAC testing scenario 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 dNFS
User scaling	1,000 – 5,000

Server and network architecture

This solution design consists of two sites:

- The source site consists of a four-node Oracle RAC 11g cluster with shared file systems using Oracle dNFS.
- The target site consists of a single-instance Oracle Database 11g test/dev database, also using Oracle dNFS.

Both the source and target sites are connected to the following networks:

- Client network
- Storage network

dNFS provides file system semantics for both Oracle RAC and Database 11g. The following IP addresses must be considered for both the source and target database servers:

- An IP address and associated host name are registered in DNS for each public network interface.
- For the storage network, EMC recommends private IP addresses.

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

Table 4. Production RAC node components

Hardware component (four computers)	Quantity
2.66 GHz quad-core processor	4
Memory	128 GB
10 GbE Ethernet NIC	2
1 GbE Ethernet NIC	3
Software	Version
Oracle Enterprise Linux	Version 5.5
Oracle RAC 11g Release 2 Enterprise Edition (includes Oracle Database 11g Release 2 and CRS)	Version 11.2.0.2
RM Agent	Version 5.2.3.0

Table 5 lists the hardware and software components of the single-instance test/dev database.

Table 5. Single instance test/dev server components

Hardware component (one computer)	Quantity
2.66 GHz quad-core processor	1
Memory	48 GB
10 GbE Ethernet NIC	2
1 GbE Ethernet NIC	3
Software	Version
Oracle Enterprise Linux	Version 5.5
Oracle Database 11g Enterprise Edition Release 2	Version 11g(11.2.0.2)
RM Agent	Version 5.2.3.0

Table 6 lists the RM Server, RM Console, and client workload components used in this testing scenario.

Table 6. RM Server, RM Console, and client workload components

Software	Version
Microsoft Windows Server 2003 R2 Enterprise Edition	Version 2003R2
RM Server	Version 5.3.0.0
RM Console	Version 5.3.0.0
Quest Benchmark Factory for Databases	Version 5.8.0

In this testing scenario, a Celerra NS-960 array is used for both production and test/dev environments. The storage array consists of the components described in Table 7 and Table 8.

Table 7. Hardware resources

Equipment	Quantity	Configuration
EMC Celerra NS-960 unified storage platform (includes an EMC CLARiiON CX4-960 back-end storage array)	1	<ul style="list-style-type: none"> • 2 storage processors • 4 Data Movers • 1 Control Station • 2 x 10 GbE network connections per Data Mover • 7 FC disk array enclosures (DAEs) • 2 SATA DAEs • 105 x 300 GB 15k FC disks • 30 x 1 TB SATA disks
10 GbE switch	1	24 CEE ports
FC switch	1	16 ports (4 Gb/s)

Table 8. Software resources

Software	Version
EMC FLARE®	04.30.000.5.507
EMC DART	6.0.40.1550692002
EMC Unisphere	1.0.0

Table 9 lists the NFS file systems used in this solution.

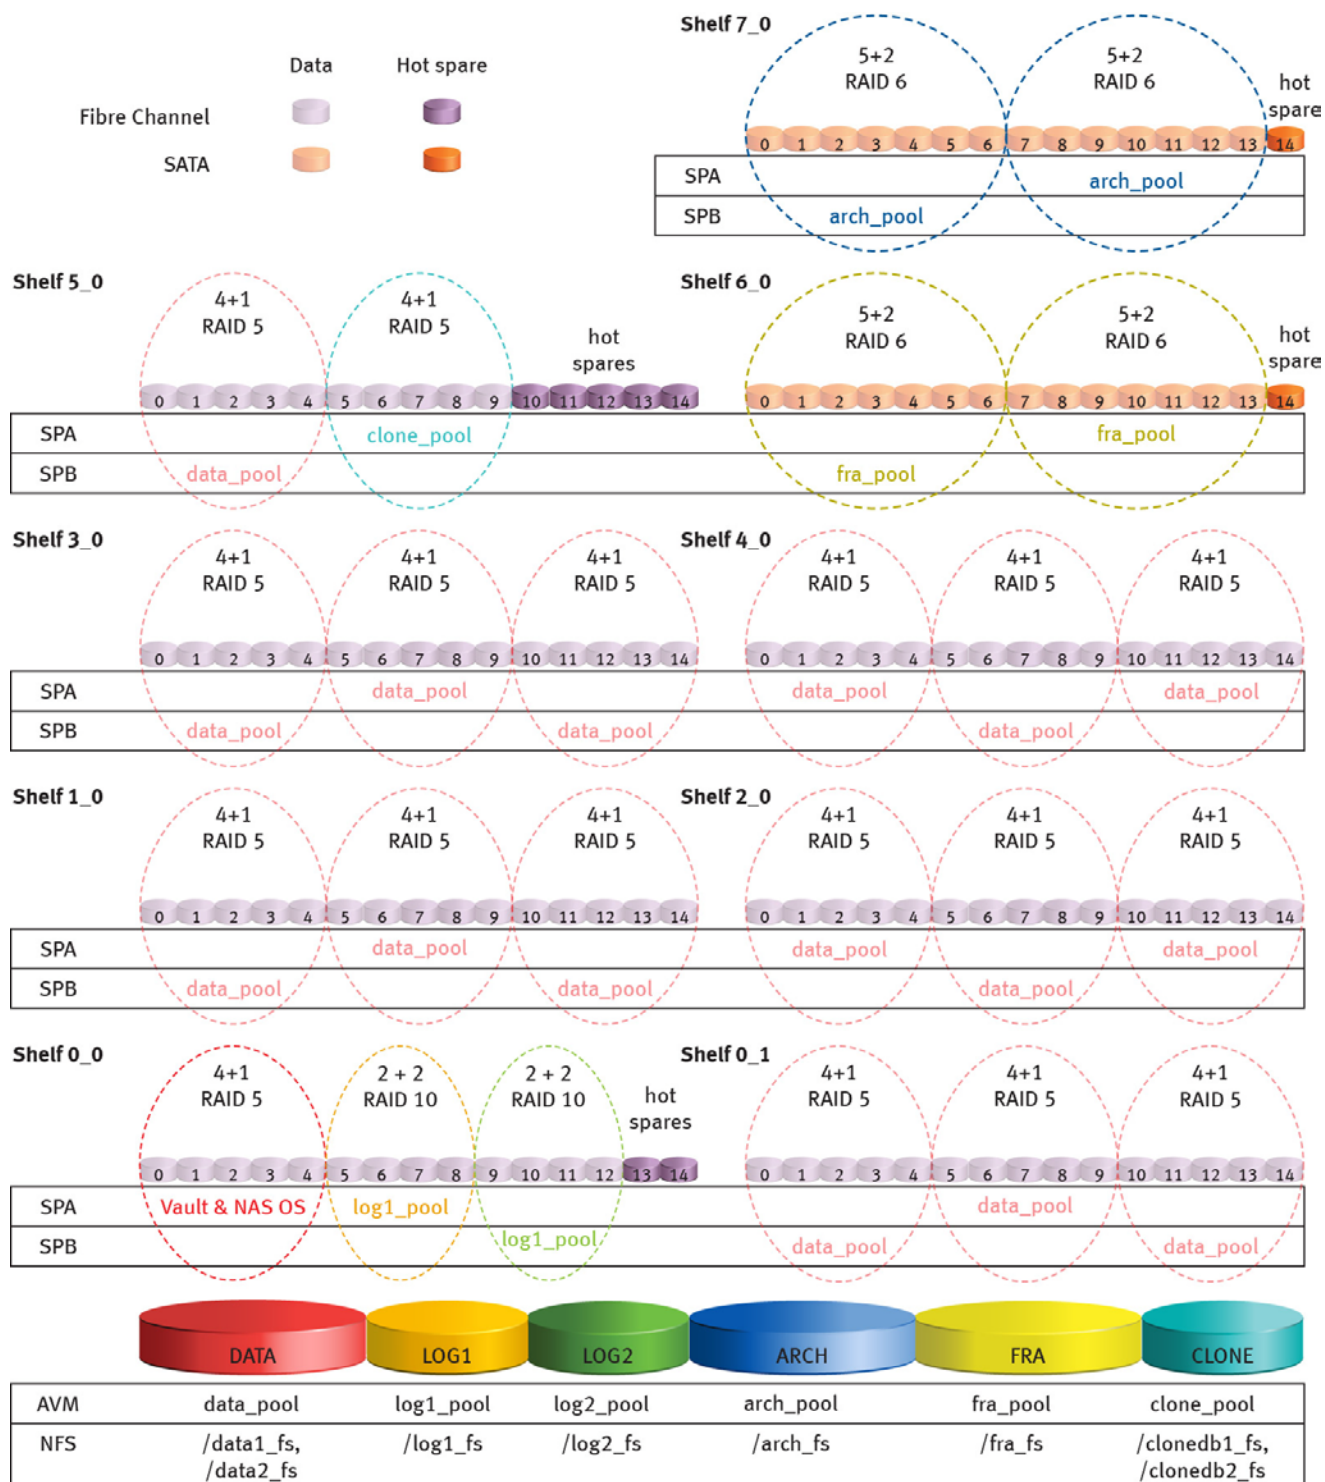
Table 9. Oracle file system allocation on Celerra

Item	File system	Disk type	RAID type
Data and temp files	/data1_fs, /data2_fs	FC	RAID 5
Online redo log, voting disk, OCR, and control files	/log1_fs, /log2_fs	FC	RAID 10
Archived log files	/arch_fs	SATA	RAID 6
Fast recovery area (FRA)	/fra_fs	SATA	RAID 6
dNFS clonedb	/clonedb1_fs, /clonedb2_fs	FC	RAID 5

Production Oracle data files, temp files, online redo log files, control files, archived log files, and fast recovery area are all accessed using dNFS. All target clone database files are also accessed using dNFS. Figure 2 provides the storage configuration overview.

Disk layout

Figure 2 shows the disk layout of the environment.



CNS-001908

Figure 2. Disk layout

Testing and validation of the RAC scenario

Overview

This section describes the high-level project implementation steps for the Oracle RAC 11g scenario:

- Configuration of an EMC Celerra NS-960 storage array
- Installation of Oracle RAC 11g Release 2
- Enabling dNFS
- Creation of a database on shared NFS file systems
- Baseline performance creation
- Creation of a Celerra SnapSure checkpoint for database file systems
- Creation of the dNFS clonedb copy to be used as the test/dev database

Configuring an EMC Celerra NS-960 storage array

Follow these steps when configuring EMC Celerra NFS for RM:

1. Confirm all mount points are exported to the production host.
2. Confirm all mount points are added to the `/etc/fstab` entries.
3. Confirm all file systems are mounted with the `/etc/fstab` entries.

Oracle RAC 11g Release 2 installation and database creation

Follow these steps to install Oracle RAC 11g Release 2:

1. Use Oracle Universal Installer to install Oracle Database 11g Release 2 with CRS.
2. Enable dNFS.
3. Create an Oracle clustered database on shared NFS file systems with Database Configuration Assistant (DBCA).

Notes:

- Refer to the *Storage architecture* section for more details about the NFS file systems used to store the database.
- For the detailed steps, refer to the following Oracle document:
Oracle Real Application Clusters Installation Guide 11g Release 2 (11.2) for Linux and UNIX

Enable dNFS

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

1. Create an `orantfstab` under `$ORACLE_HOME/dbs`.

For example:

```
server: NAS-960
path: <IP address>
export: /data_fs mount: /u02
export: /log1_fs mount: /u03
```

```
export: /log2_fs mount: /u04
export: /arch_fs mount: /u06
```

- 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:

```
cp libodm11.so libodm11.so_stub
ln -s libnfsodm11.so libodm11.so
```

Baseline performance creation

With archived log enabled, a baseline performance test is conducted on the production Oracle database as shown in Figure 3. This test consists of an OLTP workload using an industry-standard database performance benchmark. The following steps are used to create this performance test:

- Initiate the Benchmark Factory console and agents on the client computers.
- Start the Benchmark Factory job with a user load ranging from 1,000 to 5,000 in increments of 100 users.
- Monitor the progress of the test, and capture the results after the test is completed.

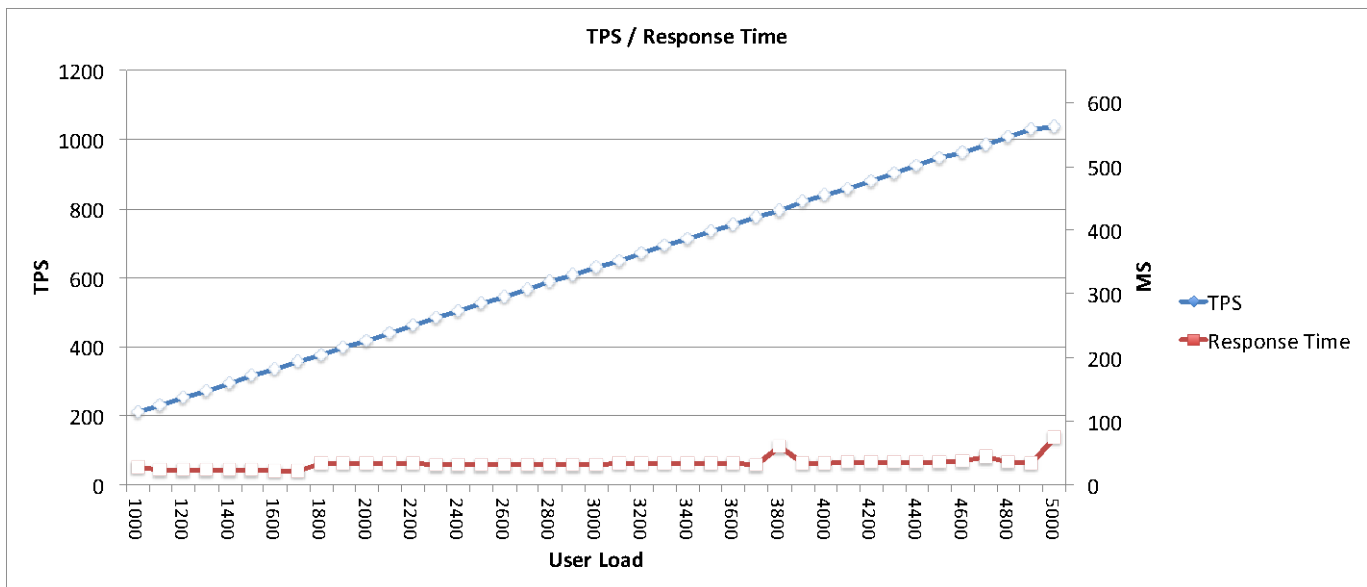


Figure 3. Baseline performance

Creation of Celerra SnapSure checkpoint for the database file system

The performance impact of performing a Celerra SnapSure checkpoint of all the relevant data files using RM is measured in terms of IOPS, response time, and TPS. RM Server and RM Console are installed on a Windows-based computer, and the RM Agent is installed on the source database server. The Celerra SnapSure checkpoint creation and Oracle hot backup operations are entirely managed by RM.

By using RM, the database is placed in hot backup mode before creating the checkpoint, and is taken out of hot backup mode after the checkpoint is completed. As seen in Figure 4, there is a slight performance impact because of this operation. However, there is no impact on TPS, and the impact on response time is minimal.

Follow these steps to create a Celerra SnapSure checkpoint to the database file system:

1. Configure RM.
2. Register the production hosts, mount hosts, and storage in RM.
3. Create the application set in RM for the database to be replicated (this application set only addresses the file systems data1_fs and data2_fs).
4. Create a job with the hot backup mode setting in the RM console to create the SnapSure checkpoint.
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 RM console (which includes the begin-hot-backup/end-hot-backup sequence).
7. Monitor the performance impact on the production database when RM is creating the SnapSure checkpoint.

Performance impact

The response time and TPS are captured when running an OLTP workload on the database. Figure 4 shows the performance of the production RAC database when the SnapSure checkpoint is created.




Figure 4. Performance during SnapSure checkpoint creation

The SnapSure checkpoint is initiated with database hot backup mode enabled at user load 2,000, and a modest increase in the transaction response time is observed during user load 2,100. While the database is in hot backup mode, additional redo is created and an Oracle database-level checkpoint occurs. As a result, the Database Writer (DBWR) process flushes all dirty blocks in the buffer cache to disk. Moreover,

before the SnapSure checkpoint is created, RM archives the current online redo log. As a result of all these operations, the I/O subsystem is heavily loaded at that time.

Figure 5 is captured from Oracle Enterprise Manager Database Control during the process of creating a SnapSure checkpoint.

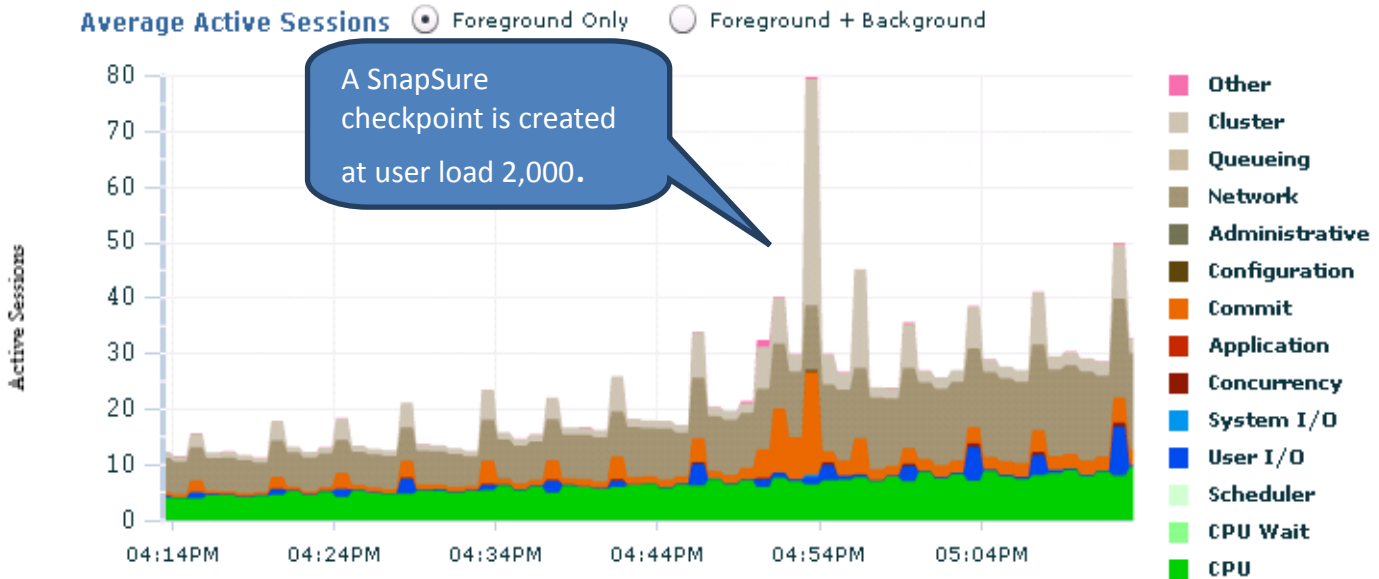
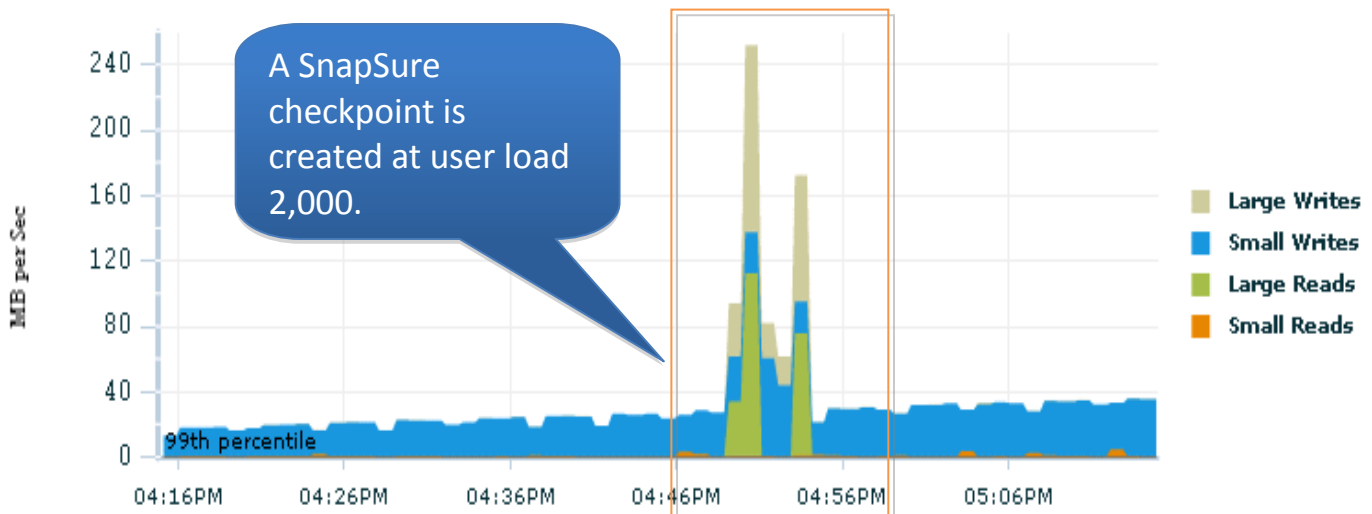


Figure 5. Database Control during SnapSure creation

As seen in Figure 5, while creating the checkpoint with hot backup mode enabled, an increase in the number of average active sessions is observed. This is because Replication Manager archives the current online redo logs before and after creating the checkpoint, and an Oracle database checkpoint is triggered when the database is placed into hot backup mode. This operation forces the Database Writer processes to flush the dirty blocks in the buffer cache to disk, which increases the load on the I/O subsystem.

The two charts shown in Figure 6 confirm that the I/O subsystem is heavily loaded during checkpoint creation:

I/O Requests per Second by I/O Type



I/O Requests per Second by I/O Type




Figure 6. I/O requests per second (from Database Control)

Using Oracle dNFS clonedb to create the test/dev database with the SnapSure checkpoint as the source

The following steps are used to create and configure the test/dev database using dNFS clonedb and the SnapSure checkpoint as the source:

1. Reconfigure dNFS in the single-instance test database by editing the dNFS configuration file `oranfstab` under `$ORACLE_HOME/dbs`.

Note:

For the detailed steps, refer to the *Enable dNFS* section.

2. Mount the SnapSure checkpoint to the target database server.
3. Mount the archive log file system to the target database server.
4. Create a new initialization parameter file by copying the initialization parameter file from a production RAC node onto the target server and then modify that file with the following edits to the parameter settings :

- a. Remove the following RAC related settings:

```
mterac281.instance_number=1
mterac282.instance_number=2
mterac283.instance_number=3
mterac284.instance_number=4
*.remote_listener='fj903-cluster-scan:1521'
*.cluster_database=true
mterac281.sga_max_size=106837311488
mterac282.sga_max_size=104689827840
mterac283.sga_max_size=106837311488
mterac284.sga_max_size=106837311488
mterac281.thread=1
mterac282.thread=2
mterac283.thread=3
mterac284.thread=4
cluster_database_instances
```

- b. Add the following parameter:

```
_no_recovery_through_resetlogs=true
```

- c. Adjust the following parameters to be specific to your environment:

```
*.sga_target
*.pga_aggregate_target
*.shared_pool_size
*.sga_max_size
*.large_pool_size
*.audit_file_dest
```

```
*.control_files
*.db_name
*.diagnostic_dest
*.instance_name
*.log_archive_dest_1
```

5. Start up the instance with the “nomount” option.
6. Re-create the control files.

Notes:

- The name of the database can be changed to another name that is different from the source database name if needed.
 - The names of the file systems that store the dNFS clonedb must not be the same as the file systems that store the SnapSure checkpoint.
7. In the new target database environment, perform dNFS clonedb by running the following command for each copy of the data file in the test/dev database stored in the SnapSure checkpoint file system:

```
dbms_dnfs.clonedb_renamefile(backup_file_name, new_data_filename).
```

For example:

```
exec
dbms_dnfs.clonedb_renamefile('/u04/oradata/mterac28/system.260.725492807.dbf', '/clonedb/uc491tst/system.260.725492807.dbf');
```

This command will create the following data file:

```
/clonedb/uc491tst/system.260.725492807.dbf
```

This file acts as the data file of the test database.

8. Recover the database with the following command in SQLPLUS:

```
SQL > Recover database using backup controlfile until
cancel;
```

When prompted, apply the appropriate archived logs.

9. Open the database with the following command in SQLPLUS:

```
SQL> alter database open resetlogs;
```

10. Re-create the temporary tablespace, with the accompanying temp file.

After you complete these steps, and the test database is up and running, you are able to read and write to the test database. When you start running your workload, 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 when creating a dNFS clonedb clone, an OLTP workload is run on the production database during the creation of the dNFS clone.

As shown in Figure 7, two dNFS clonedb databases are created during the performance run at user load 2,000 and 3,000, respectively. These dNFS clonedb databases are used for the test/dev databases. Each procedure takes 10 seconds.

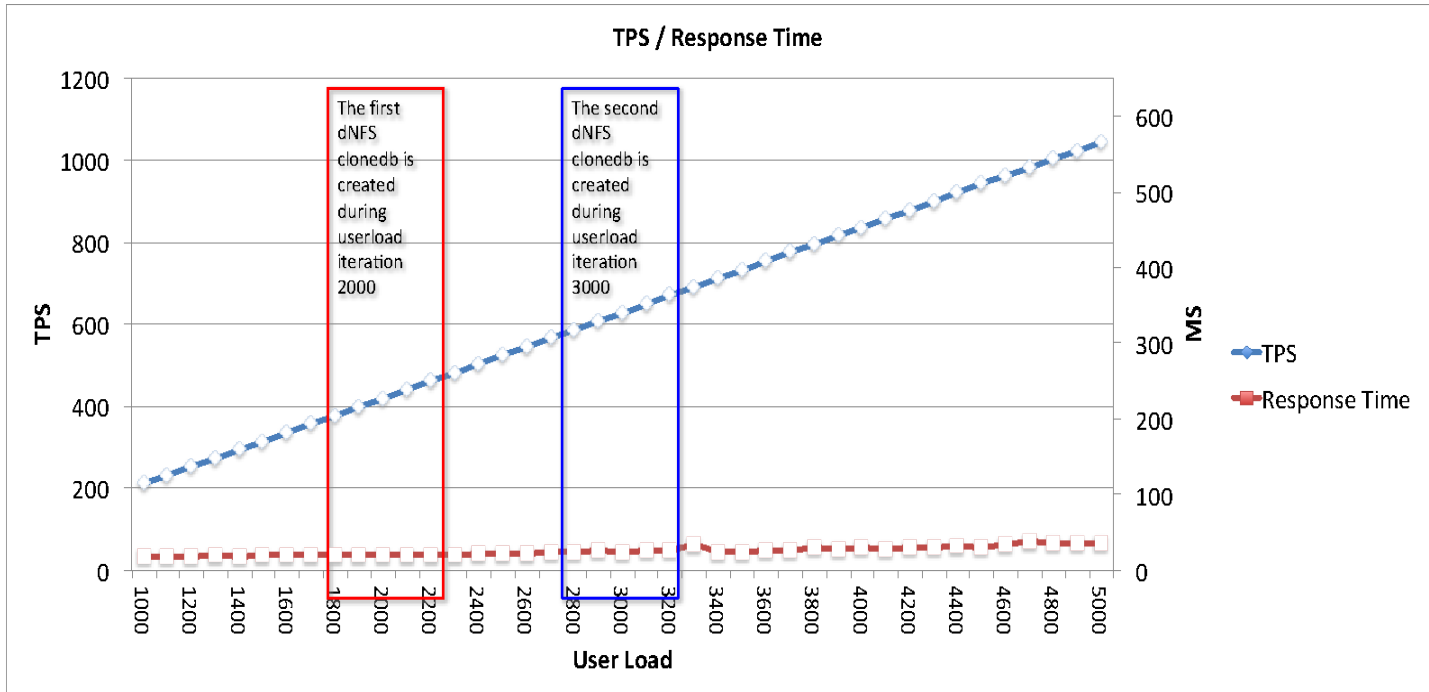


Figure 7. Performance during two dNFS clonedb procedures

Figure 7 looks nearly identical to the baseline performance diagram, which indicates that there is no significant performance impact on the production database when creating dNFS clonedb clones.

Figure 8 compares the disk space used by the production database and the test database on the target database server.

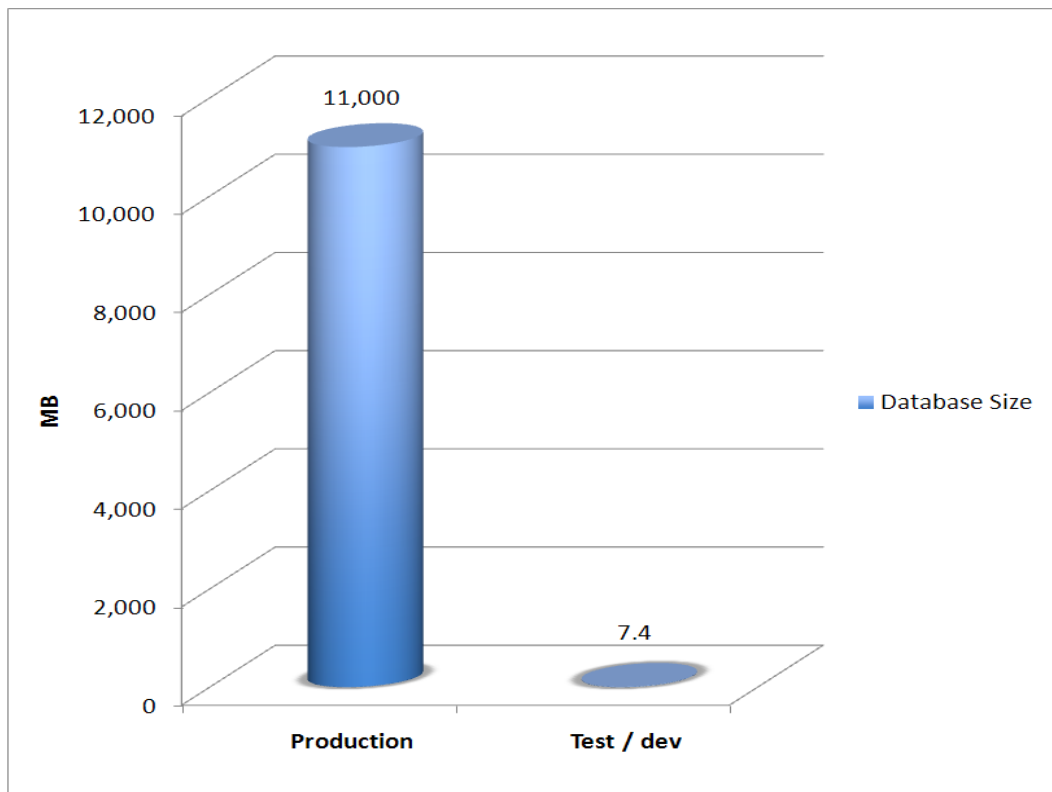


Figure 8. Space consumed by data files on production and test/dev

As seen in Figure 8, after creating a dNFS clonedb clone, the test database uses only 7.4 MB disk space. This is because production data files are not copied into the new database environment; only the blocks that are updated in the test database are written to the data files created by dNFS clonedb (the data files of the test database).

Table 10 lists the time cost summary of this test.

Table 10. 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 dNFS clonedb	10 seconds

Oracle Database 11g single-instance configuration

Overview

An additional scenario was included to determine the performance of the dNFS clonedb database itself. In order to accomplish this, it is necessary to provision a production server with an identical configuration to the test/dev server.

Therefore, in this testing scenario, the production and the test/dev environments consist of an Oracle Database 11g single-instance server. Aside from that, the configurations of the Oracle RAC 11g and Oracle Database 11g single-instance environments are identical.

Physical environment

Figure 9 shows the overall physical architecture of the single-instance environment.

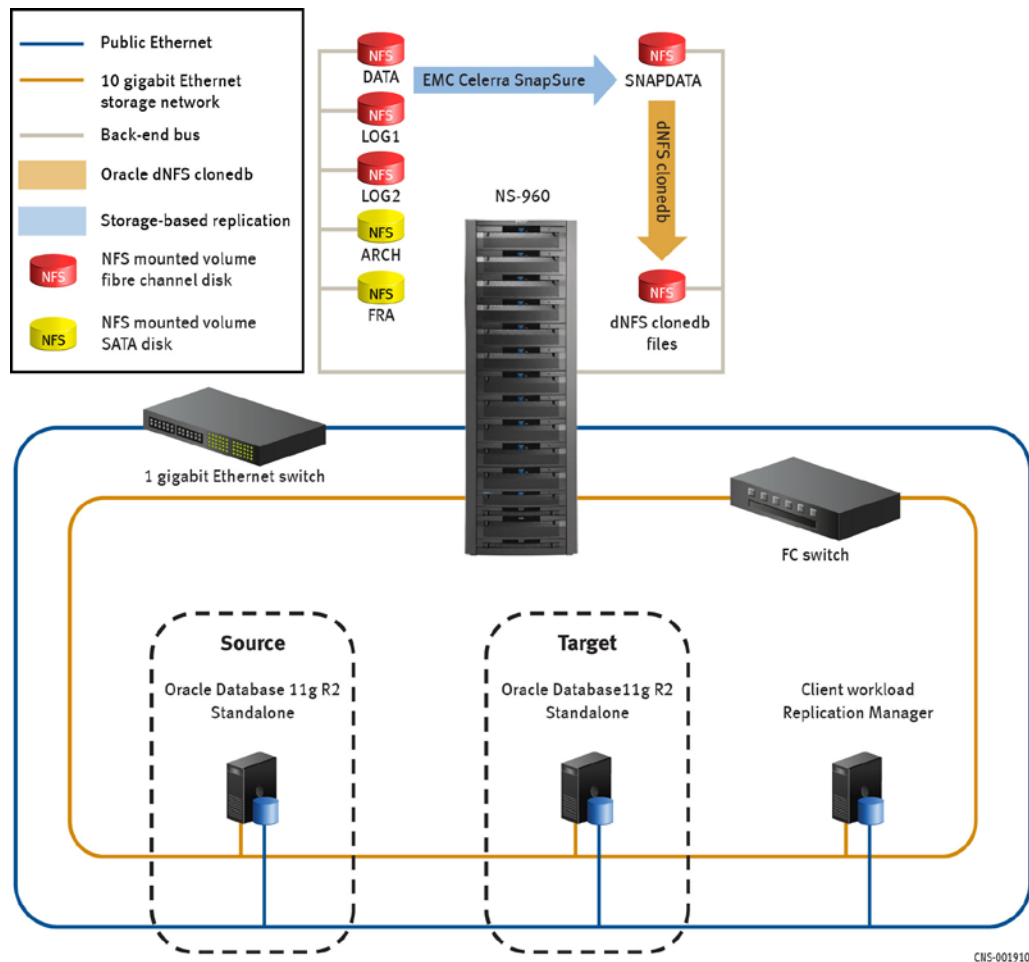


Figure 9. Single-instance physical architecture

Environment profile, software, hardware, and so on

All software versions and settings in the single-instance test environment are identical to the production RAC environment. The only exception is the production database server, which is identical to the test/dev database server configuration shown in the RAC scenario. The storage layout is also identical.

Testing and validation of the single-instance scenario

Performance workload against dNFS clonedb database

The following tests are performed in this scenario:

- A baseline test is run on the production environment, establishing single-instance performance on a normal database.
- A performance test is run against the dNFS clonedb database to compare the clonedb environment to the previous run.

Note:

The workload on the production database is not running during this test. Therefore, the workload is running only against the dNFS clonedb database.

These tests result in the following performance diagram, as shown in Figure 10:

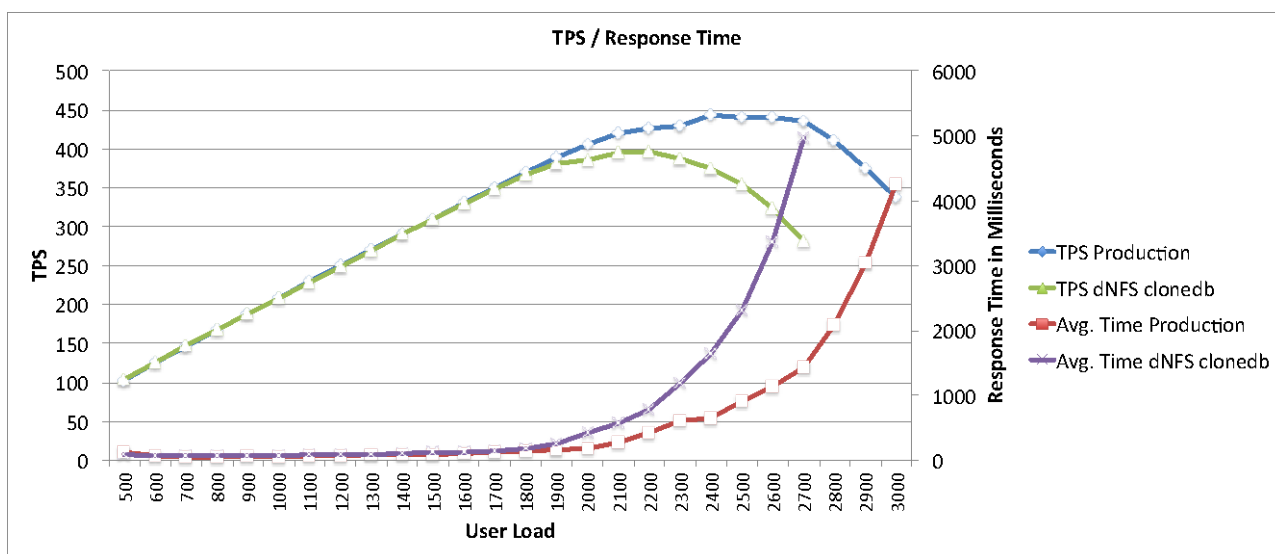


Figure 10. Single-instance performance comparison

Note:

The performance of these runs is in no way comparable to the runs on the RAC testing environment. This is because these runs were against a modest, single-instance database server, as opposed to a larger, four-node RAC cluster. This is required because the clonedb database runs on this equipment, and thus equipment used to run the production environment must be comparable.

The performance at the peak iterations of each database is as follows:

Test/dev performance:	397 TPS
Production performance:	443 TPS
Difference:	10.41%

Conclusion

Summary

This white paper highlights the benefits that EMC Celerra unified storage platforms can offer:

- Great flexibility
- A highly competitive entry price
- Excellent performance
- Multiple ways to repurpose the production environment for test/dev, QA, and other purposes

When Celerra is combined with powerful tools such as RM, they can offer a very compelling argument for hosting non-production databases environment on NFS. Furthermore, this white paper demonstrates that dNFS clonedb, a new feature in Oracle Database 11g Release 2, works well with the EMC Celerra SnapSure checkpoint technology.

Findings

This solution addresses the following key points:

- Centralizing all the non-production environments on the EMC Celerra unified storage platforms offers not only great flexibility but also a highly competitive entry price.
- By using the combination of EMC Celerra SnapSure checkpoint, RM, and Oracle dNFS clonedb, the time taken to clone a production database is greatly reduced. Moreover, the disk space required by the cloned database is minimal because only the data modified after cloning is stored.
- With the EMC SnapSure checkpoint technology and Oracle dNFS clonedb, a production database environment can be cloned in minutes. This cloned environment can be used to seed a test/dev environment, which greatly improves the efficiency of the migration process. Also, the storage overhead of the process is minimal, much less than a full physical copy.
- The performance impact of the Celerra SnapSure and dNFS clonedb creation operations are relatively minor (miniscule in the case of Celerra SnapSure and non-existent in the case of dNFS clonedb). Furthermore, the performance capabilities of the dNFS clonedb database are about 90% of a normal production database.
- Compared with traditional ways of database cloning such as Oracle RMAN and Oracle import/export, this solution minimizes the impact on the performance of production database and dramatically reduces storage requirements.

Next steps

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

References

White papers

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

- *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

For additional information, see the following Oracle Metalink Note:

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

Oracle online document

For additional information, see the following Oracle online document:

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