

# BACKUP AND RECOVERY FOR ORACLE DATABASE 11g WITH EMC DEDUPLICATION

## A Detailed Review

### EMC GLOBAL SOLUTIONS

#### Abstract

This white paper provides guidelines for the use of EMC® Data Domain® deduplication for Oracle Recovery Manager backup over a Network File System mount. The use of Data Domain deduplication is compared to a normal (non-deduplicated) NFS file system in terms of backup time, file system space, restore time, recovery time, and impact of the operation on the performance of the production database during the backup operation.

February 2011

Copyright © 2011 EMC Corporation. All Rights Reserved.

EMC believes the information in this publication is accurate 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](http://EMC.com).

VMware 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 h8511

## Table of Contents

<b>Executive summary</b> .....	<b>4</b>
Overview .....	4
Business case .....	4
Key results.....	4
<b>Introduction</b> .....	<b>6</b>
Purpose.....	6
Scope.....	6
Audience.....	6
<b>Configuration</b> .....	<b>7</b>
Environment profile .....	7
Physical environment .....	8
Disk layout .....	9
Hardware resources.....	10
Software resources.....	11
<b>Test and validation</b> .....	<b>12</b>
Backup test scenarios .....	12
Backup test results.....	12
Result analysis of backup time .....	13
Restore and recovery test scenarios .....	21
Restore and recovery test results .....	22
<b>Conclusion</b> .....	<b>23</b>
Summary.....	23
Integration of EMC technologies and Oracle RMAN .....	23
Next steps .....	23
<b>References</b> .....	<b>24</b>

## Executive summary

### Overview

Oracle Recovery Manager (RMAN) is a very common tool for backing up Oracle databases. Furthermore, the use of disk-to-disk (D2D) backup, as an alternative to traditional tape backup, is also becoming commonplace. In addition, using Network File System (NFS) as the backup storage protocol provides a simple and manageable alternative to Virtual Tape Library (VTL) or other block-based approaches. The use of NFS is very simple; it requires no complicated API or additional software to interface between the NFS mount and RMAN. An NFS mount is simply a mount, nothing more or less. Therefore, RMAN can be used as the backup software directly, and the RMAN backup can be saved directly onto the NFS mount. This is the approach taken in this white paper.

EMC® Data Domain®, as a D2D backup target, provides compelling advantages for Oracle database administrators (DBAs). Data Domain enables array-based deduplication, which eliminates the redundant data inherently stored by multiple full backups of an Oracle database. Moreover, because full backups are inherently deduplicated, there is no need to implement a complex incremental or differential backup strategy. The space savings, as a result of such approaches, are already achieved transparently by Data Domain deduplication. This benefit has already been amply documented in the following EMC publications:

- [\*EMC Backup and Recovery for Oracle 11g OLTP - Enabled by EMC CLARiiON, EMC Data Domain, EMC NetWorker, and Oracle Recovery Manager using Fibre Channel\*](#)
- [\*EMC Backup and Recovery for Oracle 11g OLTP — Enabled by EMC CLARiiON, EMC Data Domain, EMC NetWorker, and Oracle Recovery Manager using NFS\*](#)

### Business case

In addition to space utilization, businesses using D2D RMAN full backups with a conventional NFS file system experience slow database server performance during backup, restore, and recovery operations. The performance impact on the database server during the RMAN backup operation frequently overshadows any other performance issues in the system.

For this reason, the approach taken in this white paper is different from the documents referenced above. Instead of testing and documenting file system space savings provided by deduplication (as those documents have already thoroughly done), the testing documented in this white paper shows that Data Domain deduplication provides performance advantages for D2D full backups when using Oracle 11g RMAN.

### Key results

The benefits shown by the testing performed are as follows:

- The overall performance impact of the backup operation on the database server is significantly reduced with Data Domain compared to a conventional NFS file system.
- RMAN backups complete more quickly with Data Domain. On average, Data Domain RMAN backups complete 33 percent faster than the conventional NFS file system backups.

- RMAN restore time, as a critical metric for all Oracle DBAs, is significantly improved with Data Domain.

All these results identify the benefits to the customer in terms of cost: By improving RMAN backup performance, the impact on the database server is reduced, and thus the database server is able to process more transactions. This significantly reduces cost per transaction. Furthermore, by reducing the restore and recovery time, the mean time to recovery (MTTR) is improved, which leads to better uptime and availability. Because database downtime is a serious issue in all customer environments, this is a critical improvement.

# Introduction

## Purpose

This white paper documents backup, restore, and recovery performance testing using Data Domain deduplication as opposed to a conventional NFS file system. In particular, this white paper focuses on the performance impact of RMAN backup operations on the running database server and on the time required to restore and recover the database from an RMAN backup.

## Scope

This white paper contains the results of performance tests when using the Data Domain deduplication array to back up Oracle database data with RMAN to an NFS mount. The scope of this white paper is bound by the following parameters:

- To establish a baseline, multiple full backups were performed while taking an RMAN full database backup to a normal (non-deduplicated) NFS mount point. Metrics were recorded for this operation including backup data rate, backup time, file system space occupied by the backup, and the performance impact of the backup operation on the production database. While the backup operation was being performed, an industry-standard OLTP benchmark workload was running against the production database server, to create changes to the database and to provide stress on the database server itself.
- The production database was then destroyed and a full restore was performed using the last full backup. After this operation, a complete database recovery was performed. The restore time, data change rate, recovery time, and last transaction recovered were all recorded.
- These tests were then repeated against a Data Domain deduplication array mounted on the database server using NFS. All relevant metrics were collected on these tests as well. This white paper contains a comparison of these two tests, and shows the performance advantages of the Data Domain deduplication array as compared to a normal (non-deduplicated) NFS mount.

## Audience

This white paper is intended for Oracle DBAs who are involved in planning, architecting, or administering an environment with EMC VNX™ and Data Domain storage platforms, and also for those who are planning to implement backup and replication solutions.

# Configuration

## Environment profile

The white paper was validated with the environment profile shown in Table 1.

**Table 1. Environment profile**

Profile characteristic	Value
Benchmark profile	Quest Benchmark Factory – TPC-C-like benchmark
I/O response time	< 10 ms
Read/write ratio	70/30
Database scale	1,000 warehouses (TPC-C standard)
Number of databases	1
Celerra® NS-480 array drives	30 x 300 GB FC (primary database storage) 15 x 2 TB SATA, Fast Recovery Area (FRA)
Data Domain DD660 array drives	12 x 1 TB SATA (FRA)

## Physical environment

Figure 1 shows the overall physical architecture of the environment.

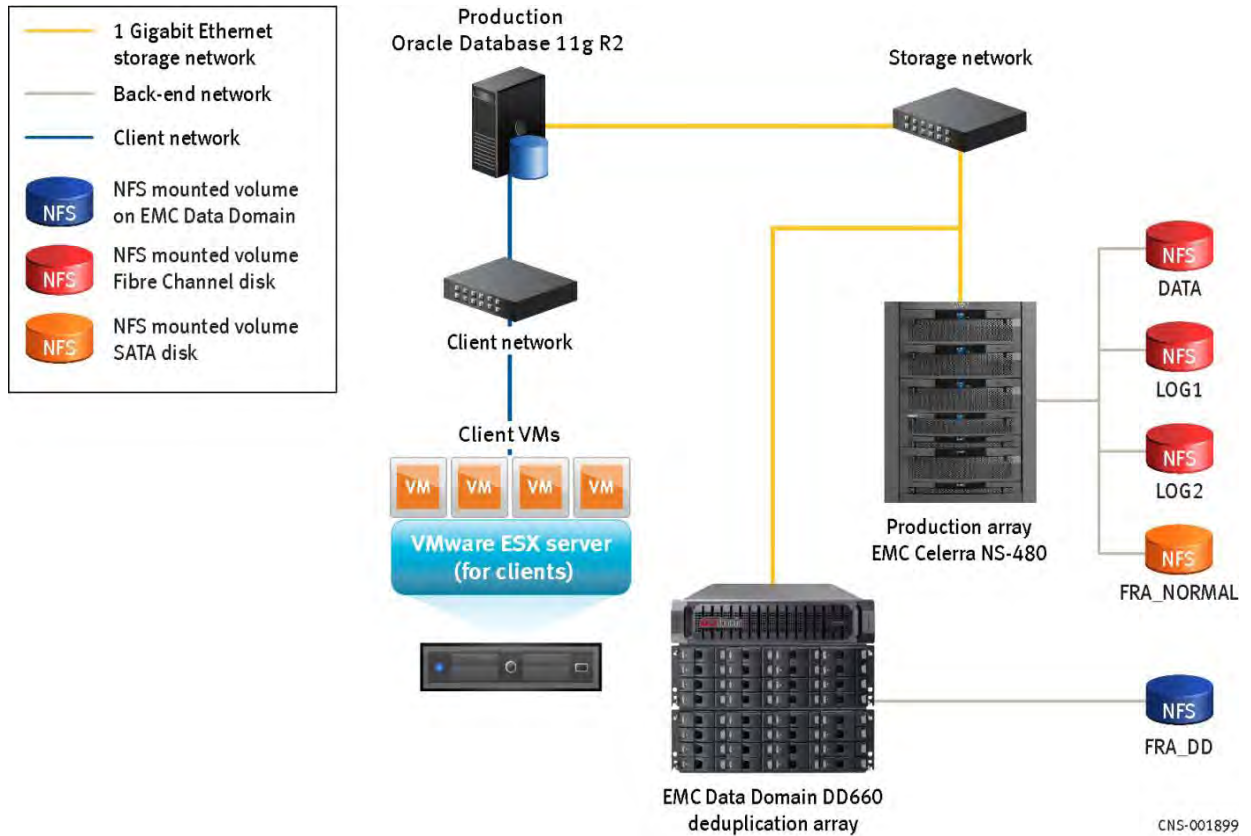
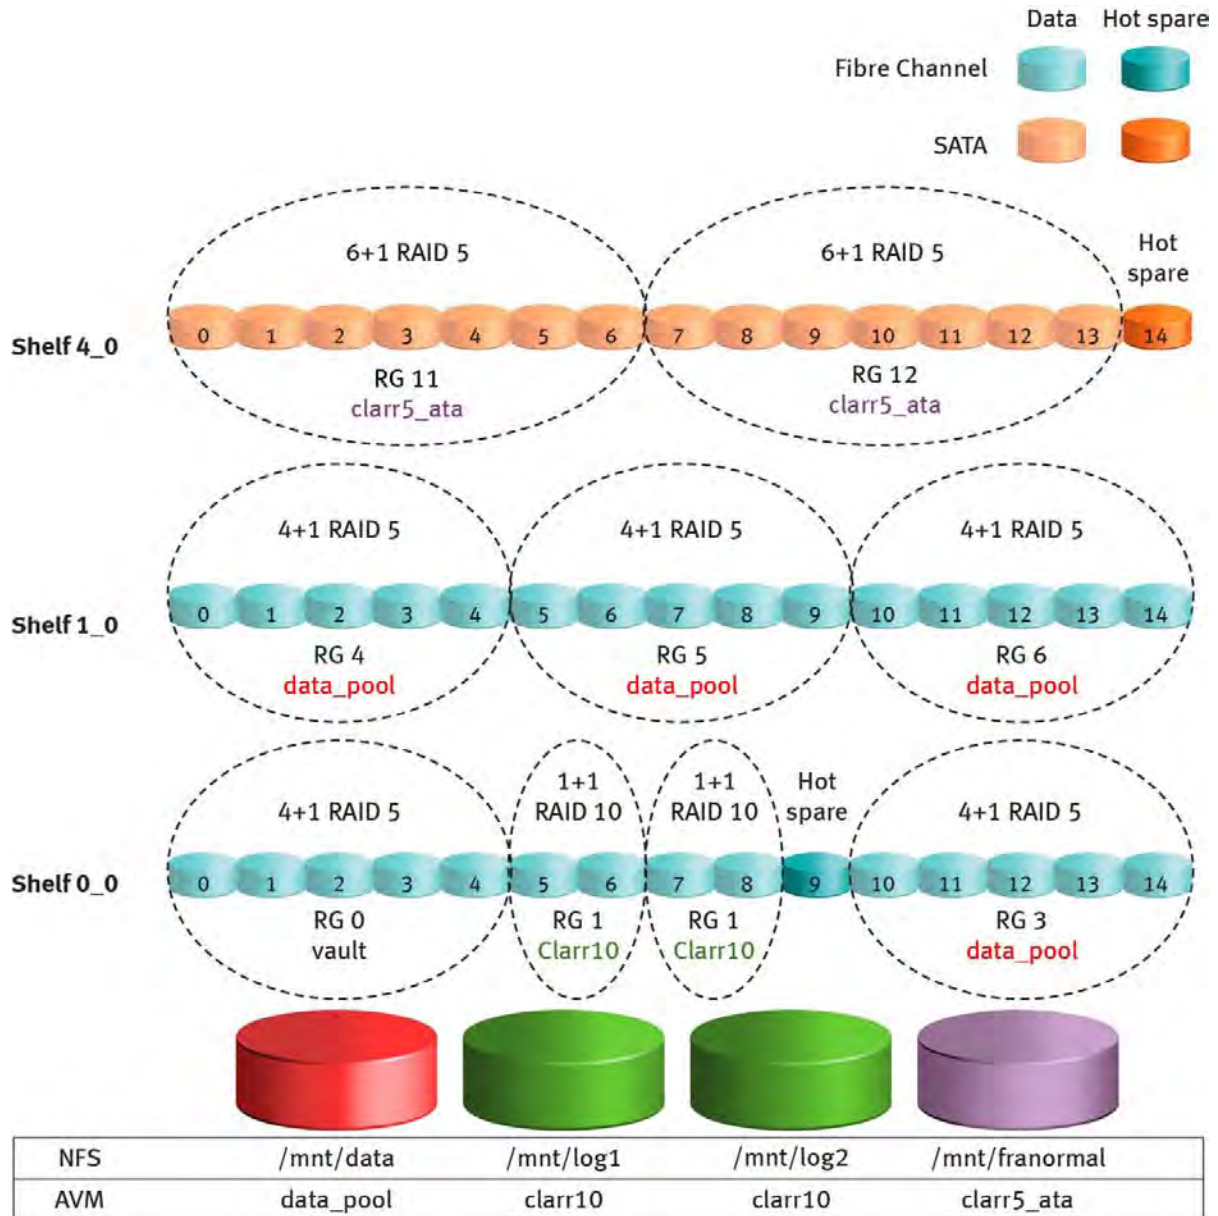


Figure 1. Overall physical architecture of the environment

## Disk layout

Figure 2 shows the disk layout on the Celerra platform.



CNS-001903

Figure 2. Celerra NS-480 disk layout

On the Data Domain DD660, the disk layout is very simple. The Data Domain DD660 contains twelve 1 TB SATA disks. During the testing, the Data Domain DD660 was enabled to configure and manage these disks automatically. When the file system exported by the Data Domain DD660 was mounted onto the host, the full capacity available was seen instantly. Therefore, the Data Domain DD660 was also tested in a default, out-of-the-box configuration.

## Hardware resources

Table 2 lists the hardware resources used to validate this solution.

**Table 2. Hardware resources**

Equipment	Quantity	Configuration
Data Domain DD660	1	<ul style="list-style-type: none"> <li>• 3 x IP connections (1 GbE)</li> <li>• 1 x SATA shelf (12 SATA 1 TB 7200 rpm disks)</li> </ul>
EMC Celerra NS-480	1	<ul style="list-style-type: none"> <li>• 2 x Data Movers</li> <li>• 1 x Control Station</li> <li>• 2 x storage processors</li> <li>• 3 x IP connections (1 GbE) per Data Mover</li> <li>• 2 x FC shelves (30 FC 300 GB 15k rpm disks)</li> <li>• 1 x SATA shelf (15 SATA 2 TB 7200 rpm disks)</li> </ul>
1 GbE switches	2	<ul style="list-style-type: none"> <li>• 16 ports per switch</li> <li>• 1 GbE throughput</li> </ul>
Database server	1	<ul style="list-style-type: none"> <li>• 2 x 2.66 GHz quad-core processors</li> <li>• 24 GB of RAM</li> <li>• 146 GB 15k internal SCSI disks</li> <li>• 2 x onboard GbE Ethernet NICs</li> <li>• 2 x additional quad-port GbE Ethernet NICs</li> <li>• 2 x 4 Gb/s dual-port FC HBAs (4 ports in total)</li> </ul>
Virtualization server	1	<ul style="list-style-type: none"> <li>• 4 x 2.86 GHz quad-core processors</li> <li>• 32 GB of RAM</li> <li>• 2 x 146 GB 15k internal SCSI disks</li> <li>• 2 x onboard GbE Ethernet NICs</li> <li>• 3 x additional quad-port GbE Ethernet NICs</li> <li>• 2 x 4 Gb/s dual-port FC HBAs (4 ports in total)</li> </ul>

## Software resources

Table 3 lists the software resources used to validate this solution.

**Table 3. Software resources**

Software	Version
Oracle Enterprise Linux	5.5
VMware vSphere™	4.1
Oracle Database 11g Enterprise Edition	11.2.0.1
Quest Benchmark Factory for Databases	5.8.1
EMC Celerra Manager Advanced Edition	5.6
EMC DART	5.6.47-11
EMC FLARE®	04.29.000.5.003
EMC Data Domain operating system	4.6.3.7-123873

## Test and validation

### Backup test scenarios

The testing performed in preparation for this white paper consisted of seven performance test scenarios. A TPC-C-like benchmark was run using Quest Benchmark factory during each scenario. The scenarios were constructed as follows:

- Scenario 1: The baseline. In this scenario, no other operations were performed. A TPC-C-like workload was run to establish the baseline performance.
- Scenarios 2 - 4: RMAN full database backup tests on the normal non-deduplicated NFS file system. RMAN disk-to-disk full database backup was completed to the fast recovery area (FRA) while the TPC-C workload was running. For confirmation of repeatability, all three of these tests were identical.
- Scenarios 5 - 7: RMAN full database backup tests on the Data Domain DD660. This test was identical to that used in scenarios 2-4 for the normal non-deduplicated NFS file system, except that the FRA was located on the Data Domain DD660 instead. For confirmation of repeatability, all three of these tests were identical.

The only difference between scenarios 2 - 4 and 5 - 7 was that the FRA was configured on the normal non-deduplicated NFS file system for the first set of scenarios, and then was configured on the Data Domain DD660 for the next set of scenarios. During each of these scenarios, the following script was executed at the beginning of the benchmark run:

```
sleep 1200
time rman <<eof
connect target /
backup database;
backup current controlfile;
exit
eof
```

### Backup test results

The Data Domain deduplication array provides compelling advantages over a normal, non-deduplicated NFS mount point for RMAN backup of an Oracle database in terms of backup, restore, and recovery performance.

## Result analysis of backup time

Figure 3 shows the result of the backup time testing. It demonstrates that the overall backup time is better with a Data Domain deduplicated NFS file system as the backup target.

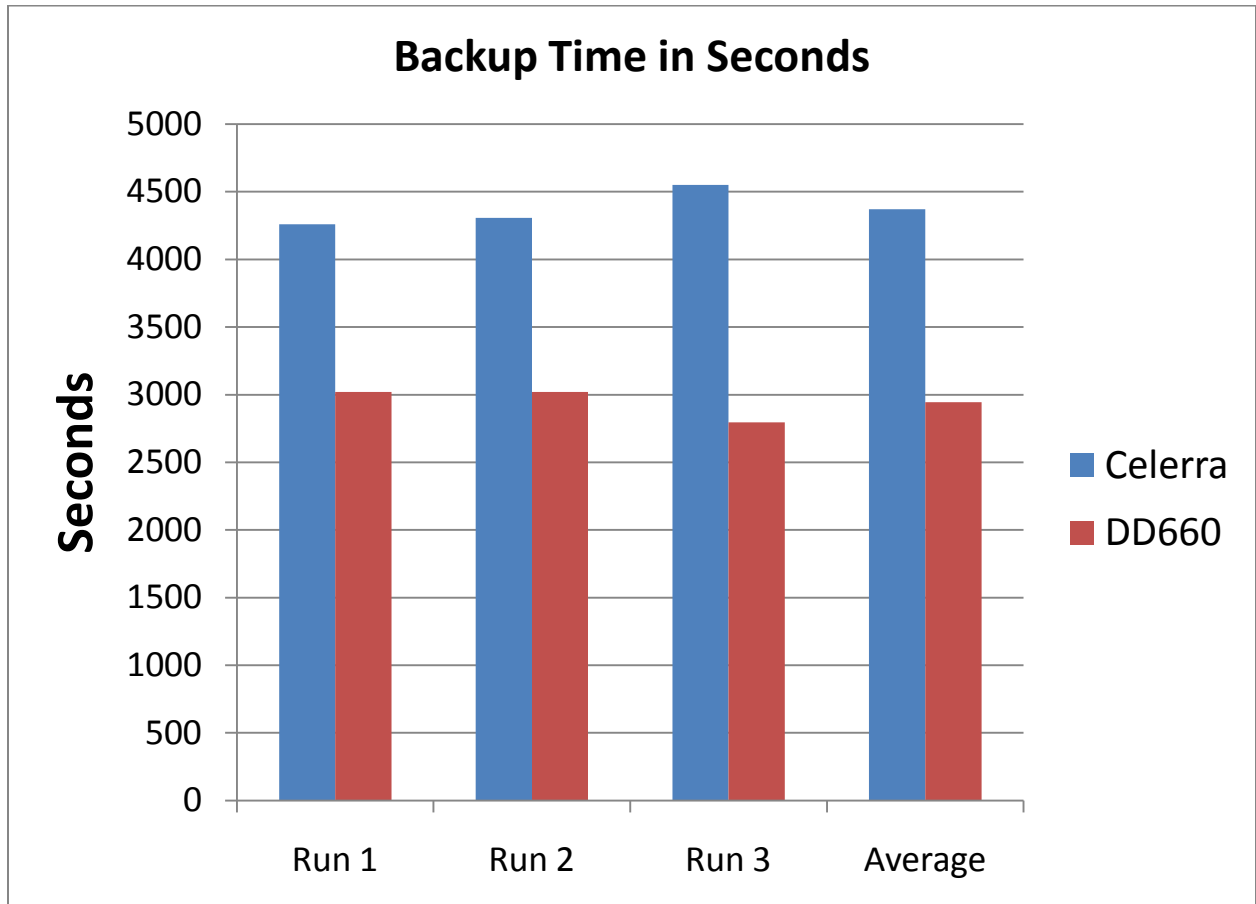


Figure 3. Backup time comparison of DD660 to conventional NFS file system in seconds

As shown in Figure 3, the RMAN backup operation always completes in significantly less time on the Data Domain DD660 than on the normal non-deduplicated NFS file system. On average, the Data Domain backups completed 33 percent faster.

Figure 4 shows the overall result of backup performance testing, demonstrating that the database performance improves with a Data Domain deduplicated NFS file system as the backup target for D2D backup to the FRA.




Figure 4. Backup performance comparison of DD660 to the normal non-deduplicated NFS file system

In Figure 4, the blue line represents the Data Domain DD660 scenarios; the red line represents the normal non-deduplicated NFS scenarios. As shown in this figure, the performance of the Data Domain DD660 during the backup scenarios is consistently higher. In other words, the performance impact of the Data Domain DD660 on the Oracle database server is lower during the RMAN backup operation than the normal non-deduplicated NFS file system. Notice that the performance of Data Domain during the backup operation is always better. In addition, confirming what was shown in Figure 3, the backup operation takes less time on the Data Domain DD660.

The individual views for these performance results are shown below.

Figure 5 shows the performance test result of RMAN backup scenario 2 (performance test on the normal non-deduplicated NFS file system as the FRA), compared to the baseline test.

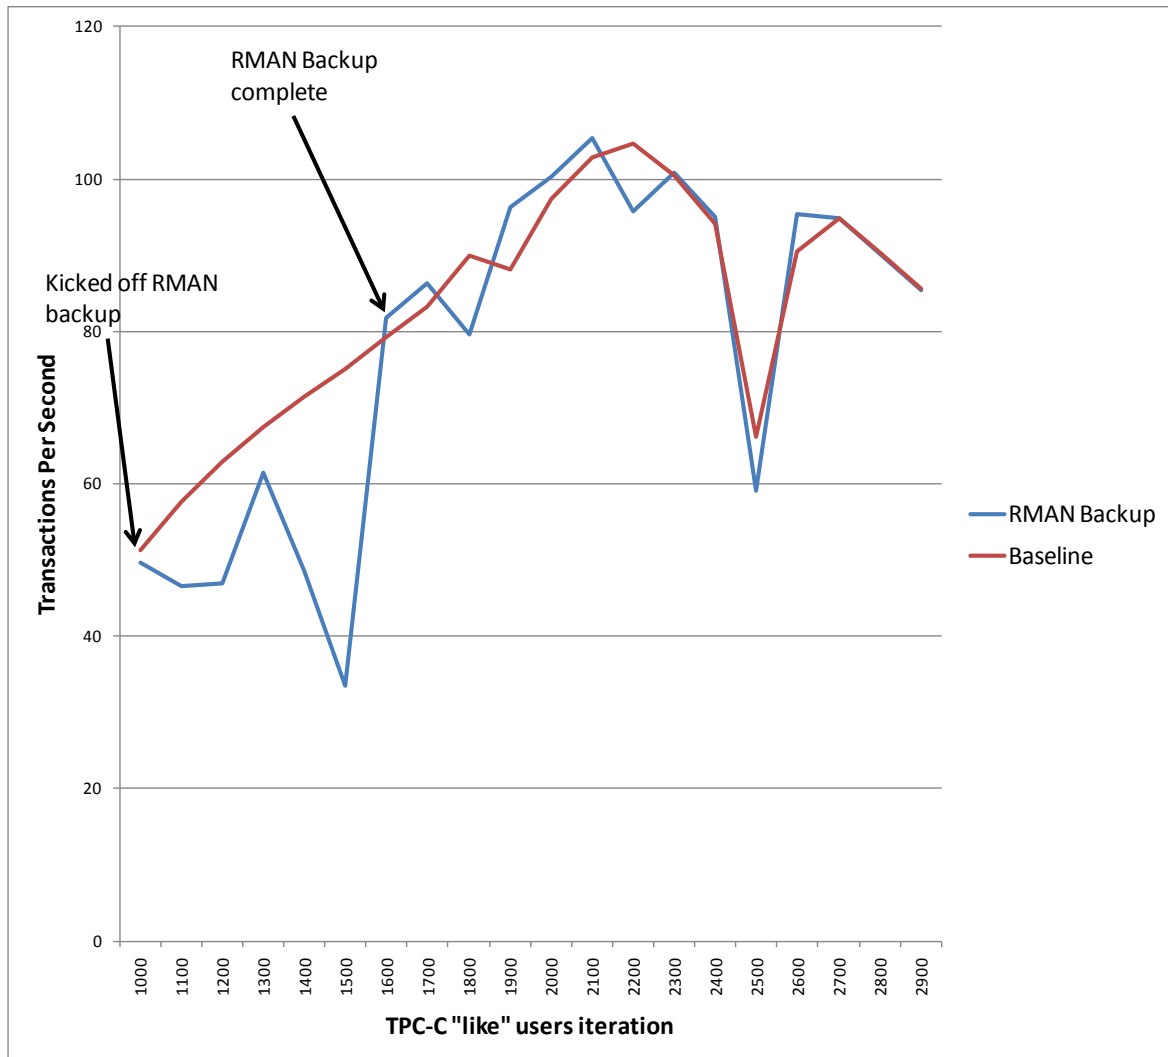


Figure 5. Result of test scenario 2 - Performance test on the normal non-deduplicated NFS file system as the FRA

Figure 6 shows the performance test result of RMAN backup scenario 3 (performance test on the normal non-deduplicated NFS file system as the FRA), compared to the baseline test.

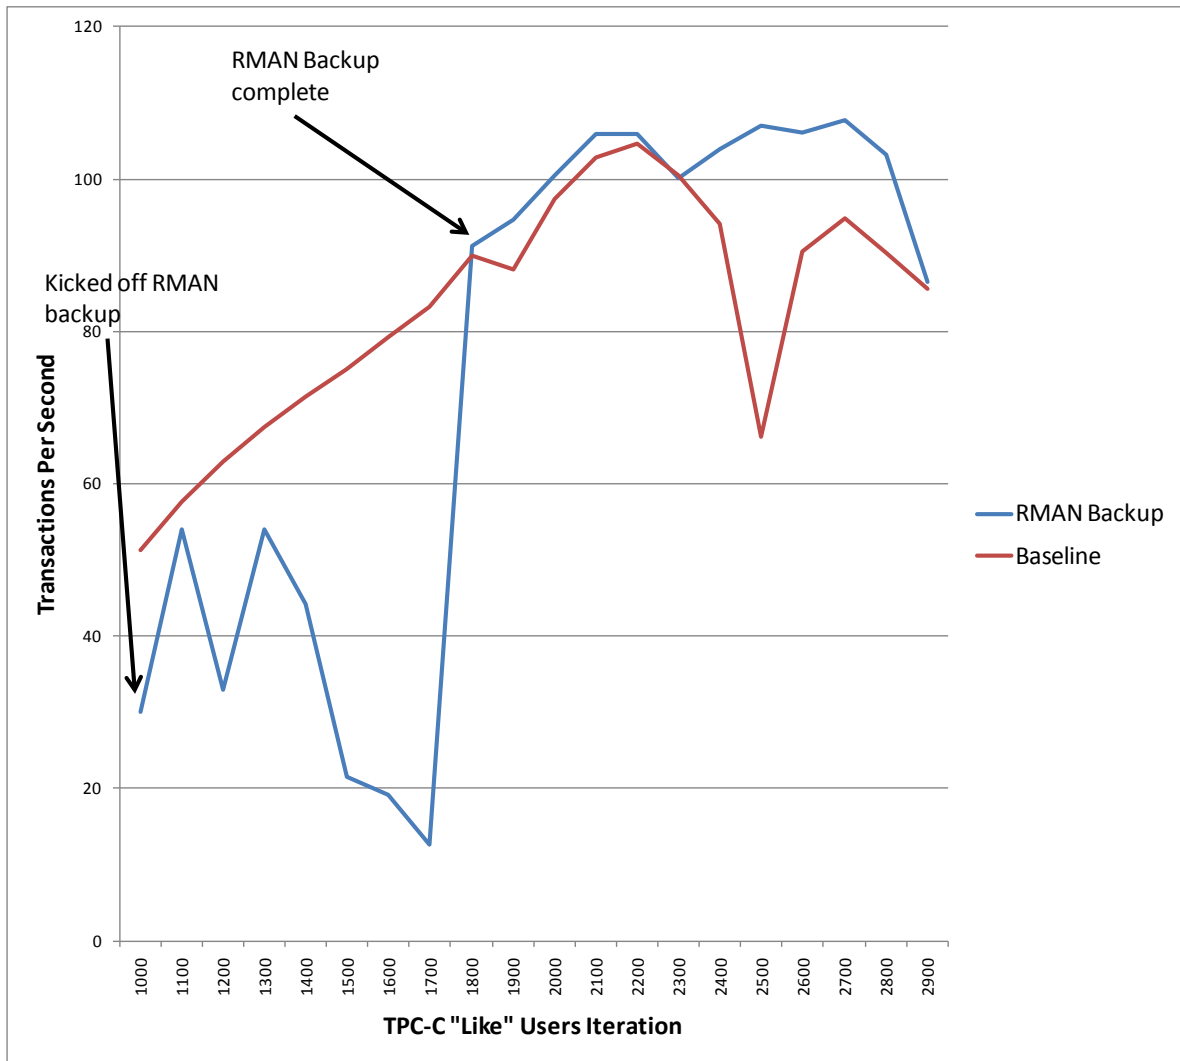


Figure 6. Result of test scenario 3 - Performance test on the normal non-deduplicated NFS file system as the FRA

Figure 7 shows the performance test result of RMAN backup scenario 4 (performance test on the normal non-deduplicated NFS file system as the FRA), compared to the baseline test.




Figure 7. Result of test scenario 4 - Performance test on the normal non-deduplicated NFS file system as the FRA

Figure 8 shows the performance test result of RMAN backup scenario 5 (performance test on Data Domain DD660 as the FRA), compared to the baseline test.




Figure 8. Result of test scenario 5 - Performance test on Data Domain DD660 as the FRA

Figure 9 shows the performance test result of RMAN backup scenario 6 (performance test on Data Domain DD660 as the FRA), compared to the baseline test.




Figure 9. Result of test scenario 6 - Performance test on Data Domain DD660 as the FRA

Figure 10 shows the performance test result of RMAN backup scenario 7 (performance test on Data Domain DD660 as the FRA), compared to the baseline test.




Figure 10. Result of test scenario 7 - Performance test on Data Domain DD660 as the FRA

## Restore and recovery test scenarios

The restore and recovery testing was performed after the backup testing was completed.

**Note** A “restore” operation is the act of placing the database back into the file system from a source backup, and a “recovery” operation is the act of applying transactions to that newly restored database using archived logs in an operation known as media recovery.

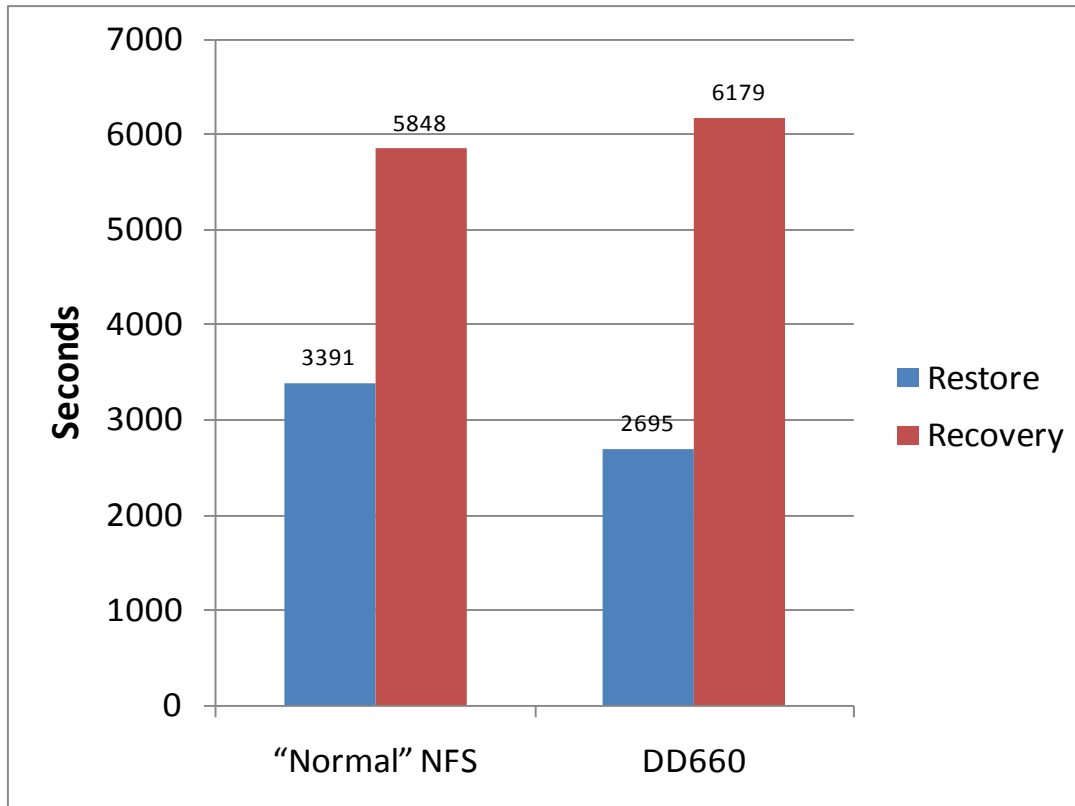
This testing consisted of running timed operations for both the restore and the recovery operations. After destroying the production database (except for the control files), the following scripts were run:

```
time rman <<eof
connect target /
startup mount;
restore database;
exit
eof
```

```
time rman <<eof
connect target /
recover database;
alter database open;
exit
eof
```

## Restore and recovery test results

The result of this testing is shown in Figure 11.



**Figure 11. Restore/recovery performance comparison of Data Domain DD660 to the normal non-deduplicated NFS file system**

As shown in Figure 11, the restore operation was completed significantly faster with the Data Domain DD660 than with the normal non-deduplicated NFS file server. This result proves that Data Domain DD660 provides a significant advantage in terms of MTTR over non-deduplicated file servers.

However, this testing also shows that the recovery time with the Data Domain DD660 is slightly longer. For optimal recovery time, EMC's best practice recommendation is to store archived log files on another file system besides Data Domain. The Data Domain DD660 is excellent for deduplication of redundant data. However, archived logs are inherently unique, and thus minimal, if any, so deduplication can be achieved. Further large sequential I/O, such as that used when applying archived logs, is not optimal on the Data Domain. Therefore, in this case, the normal NFS file server would be a better choice for storing archived logs.

## Conclusion

### Summary

The Data Domain DD660 provides significant performance enhancements for the Oracle DBA when compared to a default-state NFS file server using a similar number of SATA disks. Both backup and restore operations are improved significantly with Data Domain. In addition, the use of NFS provides a completely transparent and simple interface for creating an RMAN backup. The use of Data Domain technology, which exposes an NFS interface, enables this functionality nicely.

The recovery operation is slightly slower with the Data Domain DD660, compared to the default NFS file server, which is also a reasonable result. Therefore, EMC recommends that customers configure the archive dump destination to be on a normal NFS file server and back up those files to a non-Data Domain location as well.

This testing also proves that the Data Domain DD660 works well as a backup target using the FRA as an NFS mount configuration, while performing an RMAN full backup of an Oracle 11g database.

### Integration of EMC technologies and Oracle RMAN

EMC and Data Domain provide tight integration with Oracle RMAN through a variety of protocols and interfaces. Oracle RMAN backups can be managed using EMC NetWorker®, EMC Replication Manager, or simple customer scripts. The use of EMC storage virtualization and replication (such as snapshots) further enhances these capabilities. The use of target-based deduplication, which is the hallmark of the Data Domain technology, is completely transparent to the database server, and provides advantages in terms of file system space savings, backup time, performance impact, and restore and recovery time. This enables the Oracle RMAN backup of larger databases, as the backup window is reduced significantly.

Third-party tools such as Symantec NetBackup can take advantage of these features as well. Target-based deduplication fits into all of these environments transparently without changing the customer's existing processes or scripts.

EMC and Data Domain also support the use of proxy servers to offload the Oracle RMAN process from the production database server. This can be used with Data Domain target-based deduplication as well.

### Next steps

To learn more about this and other solutions, contact an EMC representative or visit [www.EMC.com](http://www.EMC.com).

## References

For documents that provide similar solutions, see the following Reference Architectures:

- [\*EMC Backup and Recovery for Oracle 11g OLTP - Enabled by EMC CLARiiON, EMC Data Domain, EMC NetWorker, and Oracle Recovery Manager using Fibre Channel\*](#)
- [\*EMC Backup and Recovery for Oracle 11g OLTP – Enabled by EMC CLARiiON, EMC Data Domain, EMC NetWorker, and Oracle Recovery Manager using NFS\*](#)