

# Simple, Inexpensive Oracle Recovery Manager Backups on EMC Celerra MPFS

*Applied Technology*

---

**Abstract**

Oracle Recovery Manager (RMAN), included with the Oracle database license, provides a simple and inexpensive backup-to-disk methodology for Oracle. This white paper highlights the benefits that an EMC<sup>®</sup> Celerra<sup>®</sup> server and MPFS provide to RMAN users for uncompromised accessibility and optimal performance.

January 2009

---

---

Copyright © 2009 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](http://EMC.com)

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

Part Number h6048

---

## Table of Contents

<b>Executive summary .....</b>	<b>4</b>
<b>Introduction .....</b>	<b>4</b>
Audience .....	4
Terminology .....	4
<b>Overview.....</b>	<b>6</b>
Network reference architecture.....	6
<b>RMAN: cost-effective backup .....</b>	<b>7</b>
Cost benefits of using RMAN.....	7
<b>Celerra IP storage: uncompromised accessibility .....</b>	<b>8</b>
Cost benefits of using Celerra IP storage .....	8
<b>MPFS: optimal performance .....</b>	<b>9</b>
Cost benefits of using MPFS .....	9
Faster backup speeds .....	10
Improved restore speeds.....	12
Improved EMC Celerra usage.....	14
<b>Testing strategy .....</b>	<b>15</b>
Hardware and software resources .....	16
RAID layout .....	17
File system layout.....	18
<b>Conclusion .....</b>	<b>19</b>
<b>References .....</b>	<b>19</b>

---

## Executive summary

Oracle Recovery Manager (RMAN) is a comprehensive and powerful server-managed tool from Oracle for backing up and restoring hot and cold Oracle databases. As the size and the number of databases continue to grow, the cost and time required for backing up the databases and in turn restoring them become a challenge.

Currently expensive production database servers actively perform backup activities for extended periods of time. The high hardware and licensing costs of production servers make it prudent and efficient to execute RMAN backup operations as efficiently as possible. In addition, they must prevent non-critical database servers such as reporting, test, and development instances, and any process that accesses copies of production data from impacting performance. This white paper describes how the following architecture choices can provide monetary and time savings by accelerating Oracle database RMAN backups using EMC® Celerra® network-attached storage and MPFS:

- RMAN is a feature of the Oracle database and is included as part of the Oracle license. Therefore, it is less expensive than many third-party database backup applications, such as Oracle Secure Backup and Veritas NetBackup. Refer to the [RMAN: cost-effective backup](#) section for more cost advantages of using RMAN.
- Celerra environments offer a central backup repository that is accessible from multiple servers. Refer to the [Celerra IP storage: uncompromised accessibility](#) section for details about the benefits of using a consolidated disk pool for RMAN backups.
- MPFS provides high performance that shortens the backup and recovery windows. Refer to the [MPFS: optimal performance](#) section to see backup and restore testing results that support this solution.

## Introduction

Businesses requirements for recovery time objective (RTO) and recovery point objective (RPO) for database environments have become more stringent. The RMAN utility provides an inexpensive, fast, and convenient way to back up and restore an Oracle database. RMAN requires a media manager to send backups to tape or needs disk storage space to perform backups and restores. Combining the features of RMAN with EMC Celerra and MPFS using SATA disks on a shared file system offers an affordable high performance solution for centralized online backup storage. This white paper highlights the cost and time saving benefits of using Celerra with MPFS for both hot and cold RMAN backups. Additionally, we will show the benefits of restoring from a Celerra shared file system to a separate test or development database server with zero impact to the production Oracle database.

This solution leverages Celerra's industry-leading MPFS feature to enable RMAN to accelerate database backup and recovery with outstanding performance. MPFS acts as a standard NAS file-sharing environment without any application changes on existing infrastructure.

## Audience

The intended audience for this document is IT planners and managers, storage architects, and database, storage, and backup administrators who are involved in IT business and backup strategies or seek a context for the implementation and configuration of optimal RMAN solutions.

## Terminology

- **Backup:** A process that involves making a copy of a database and its contents in another location for safekeeping and for re-creating the database as it was at that point in time.
- **Backup window:** The timeframe that a system is available to perform a backup procedure. Backup procedures can be detrimental to system and network performance since they sometimes require the primary use of the system to be suspended. These effects can be mitigated by arranging a backup window with the users of the system.

- 
- **Channel:** Channel in RMAN is a communication pipeline between a RMAN executable and a target or auxiliary database. The RMAN console sends commands to the database using this channel, and the server session running on the database executes the command on behalf of the RMAN. Multiple channels can be used to achieve a higher level of parallelism.
  - **Cold backup:** Cold (offline) backup is a physical backup. During cold backup the database is shut down and not available to users. All files of the database are copied (image copy). The database does not change during the copy, so the database is in sync upon restore.
  - **DM:** Data Mover. The EMC Celerra NAS-blade.
  - **Fallthrough:** The event when MPFS temporarily reverts back to the NFS protocol to provide continuous data availability, reliability, and protection while block I/O path congestion or unavailability is resolved, seamless to the application requesting the I/O.
  - **FC:** Fibre Channel. A high bandwidth data-transfer mechanism that is designed for high-performance storage systems in a SAN.
  - **Hot backup:** Hot (online) backup is a physical backup much like the cold backup; the difference is that the database remains open and available to users. All files of the database are backed up. There may be changes to the database as the backup takes place, so all log files of changes being made during the backup must be saved as well. When doing a restore, the changes in the log files are reapplied to bring the database back in sync.
  - **MPFS:** Multi-Path File System. A Celerra Network Server feature that allows heterogeneous clients with MPFS client software to concurrently access shared data stored on a Symmetrix<sup>®</sup> or CLARiiON<sup>®</sup> array directly over Fibre Channel or iSCSI.
  - **Oracle Recovery Manager:** Oracle Recovery Manager (RMAN) is an Oracle provided utility for backing up, restoring, and recovering Oracle Databases.
  - **NFS:** Network File System. A protocol developed by Sun Microsystems that allows a computer system to access files over a network as if they were on its local disks.
  - **Restore:** Rebuilding a database from a set of backup files.
  - **RPO:** Recovery Point Objective. The recovery point objective (RPO) is the maximum amount of time allowed where unrecoverable changes may have occurred to your database as defined by your organization; it is often the time between the last available backup and time a disruption could potentially occur.
  - **RTO:** Recovery Time Objective. The recovery time objective (RTO) is the maximum amount of time allowed that a business process or service can be unavailable without unacceptable consequences. For a database outage, this is the amount of time allowed for restoring the database and returning it to full service.
  - **SLA:** Service Level Agreement. The part of a service contract where the level of service is formally defined.
  - **SP:** Storage Processor. An integral component of the storage array (CLARiiON) that controls the transfer of data from a NAS file or application server to the disk drives.
  - **TCO:** Total Cost of Ownership. The financial estimate that is designed to help managers assess direct and indirect costs commonly related to software or hardware.

---

## Overview


This section details the network architecture and hardware and software resources for this solution. The following sections detail the cost benefits using EMC Celerra with MPFS to accelerate backup and restore for RMAN:

- [RMAN: cost-effective backup](#)
- [Celerra IP storage: uncompromised accessibility](#)
- [MPFS: optimal performance](#)

## Network reference architecture

Figure 1 shows the reference architecture of the RMAN solution and how it was configured to back up and restore the database. The solution includes:

- Multiple Oracle database servers, each with a gigabit Ethernet NIC connected to the EMC Celerra NS40FC, NS120, or NS480 and a dual-port Fibre Channel host bus adapter (HBA) connected to each of the storage processors integrated into the Celerra. In the sample testbed, one server was used as a production server and the other was used as a testing and development (test/development) server.
- One EMC Celerra NS40FC, NS120, or NS480 that acts as the target destination for the backup data.
- A CLARiiON or Symmetrix array that holds the production Oracle file system. In our sample testbed, a Linux file system was striped across four CLARiiON RAID 5 4+1 LUNs on both servers for the production database and for the test/development database copy.



**Figure 1 RMAN to Celerra with MPFS reference architecture**

---

## RMAN: cost-effective backup

RMAN is a command line and Oracle Enterprise Manager-based tool, and is the Oracle preferred method for efficiently backing up and recovering your Oracle database. RMAN is designed to work intimately with the Oracle database by providing block-level corruption detection during backup and restore operations. Backing up Oracle databases over tape is very time-consuming and error-prone. By using disk as the backup media, you can accelerate RMAN backup and restore operations. A highly available Celerra shared file system is a perfect complement for RMAN disk backups. EMC MPFS provides a high performance, robust, feature-rich, and economical shared file system. The X-Blade in an MPFS environment handles metadata processing, client authentication, file sharing, and maintaining data consistency using notifications and shared file system journaling and other EMC Celerra advanced features.

The RMAN backup methodology ties in to the complete disaster recovery / business continuity (DR/BC) strategy. The Celerra file system is the primary storage for the online backups, and the Celerra can be backed up to tape with NDMP or other backup methods without the expense of Oracle database backup-to-tape licenses. RMAN still provides catalog information for its backups, and the centralized file-based backup method to EMC Celerra allows for flexibility from there.

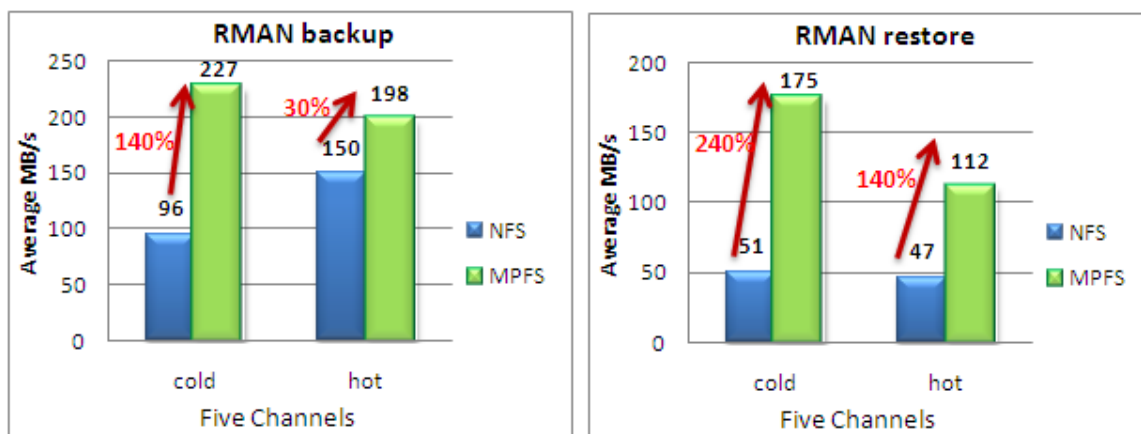


Figure 2 Example RMAN backup and restore speeds over MPFS and legacy NFS

### Cost benefits of using RMAN

RMAN is the Oracle-recommended method to back up an Oracle database. Customers are comfortable with this approach from a support and best practices position. Oracle makes RMAN mandatory if backing up an Oracle database that uses ASM and RMAN backup to disk is preferred over tape backup because it offers a significantly quicker data restore and recovery.

Customers with existing Network Data Management Protocol (NDMP) licenses can use this functionality to write database backups to tape at zero additional cost. NDMP is already owned by many existing Celerra customers. In this scenario:

- The tools and mechanisms are familiar to database administrators.
- Additional tools or media manager software is not required.

Companies can eliminate the cost of tape media management software needed to transfer RMAN backups to tape. Depending on the vendor, this can vary from \$5,000 to \$30,000 for software licenses and set fees for the number of tape devices used.

---

## **Celerra IP storage: uncompromised accessibility**

NAS simplifies storage management and improves the reliability, performance, and efficiency of the network. Since the Celerra is a self-contained, intelligent device that attaches directly to an existing LAN, any file system that is located and managed on the Celerra can supply data to clients using industry-standard file sharing protocols. Therefore, this allows true data sharing among heterogeneous network clients.

### ***Cost benefits of using Celerra IP storage***

The high-availability and resilient nature of Celerra IP storage solutions may eliminate the need for tape backup altogether, thus saving time, cost, and speeding up backup and recovery.

In addition to providing a central backup repository that is easily accessible from other systems, using shared backups to deploy databases clones for testing or development environments does not disrupt the production database CPU. Here are some other advantages to using consolidated disk pool:

- Faster restore times. Using a shared backup to restore a database takes less time than restoring from low-speed tape.
- Better utilization of disk space. Instead of allocating 30-40 percent of free space on multiple disparate direct attached storage (DAS) devices, this approach centralizes the pool and makes it easier to manage and scale according to the demand.
- Ease of management. When archive logs and other data are backed up and consolidated on a single Celerra, this centralized data management provides huge time and cost savings.
- Automation of backup. A standardized methodology for backup on a centralized share lends itself to automation of backup and archive to tape. One location, one process.
- High availability. Consolidated backups on Celerra IP storage are protected and can be replicated asynchronously across wide geographical areas. This offers peace of mind and security in the knowledge that the backup system will always be operational and available.

---

## MPFS: optimal performance

EMC Celerra with MPFS technology allows customers to experience higher performance, shorter backup windows, and faster restore times while maximizing the investment in hardware infrastructure. Using EMC Celerra with MPFS technology can reduce backup and restore windows by up to 120 percent.

To effectively demonstrate the time savings gained by using MPFS to accelerate NAS backup with RMAN, we performed a total of 40 tests on both MPFS as well as legacy NFS protocols. Hot and cold RMAN backups were performed on the production server and stored on a shared file system; these backups were restored to the test and development server.

In addition to running hot and cold backup and restore tests on MPFS and NFS, RMAN backup was performed scaling the backup and restore channels from one to five. We ran a total of five tests for each dimension of testing. Adding multiple threads significantly improves performance. [Table 1](#) lists the dimensions for this storage configuration.

**Table 1** Testing dimensions of Oracle RMAN performance

Type of test	Protocol	Mode of operation	Number of RMAN channels
Backup	MPFS	Cold	1, 2, 3, 4, and 5
	MPFS	Hot with 500 users	1, 2, 3, 4, and 5
	NFS	Cold	1, 2, 3, 4, and 5
	NFS	Hot with 500 users	1, 2, 3, 4, and 5
Restore	MPFS	Cold	1, 2, 3, 4, and 5
	MPFS	Hot with 500 users	1, 2, 3, 4, and 5
	NFS	Cold	1, 2, 3, 4, and 5
	NFS	Hot with 500 users	1, 2, 3, 4, and 5

---

**Note:** A TPC-C workload in QWEST Benchmark Factory was used to generate a minimal 500-user load on the production database for the hot backup testing to simulate a real-world production database with active users.

---

Refer to the [Testing strategy](#) section for additional details about the backup and restore testing. The figures in the [Cost benefits of using MPFS](#) section reveal that using MPFS greatly reduces backup and restore times.

### ***Cost benefits of using MPFS***

Testing shows that MPFS outperforms legacy NFS in several categories. The benefits of MPFS with RMAN on Celerra backups include:

- Significantly faster backup speeds, producing shorter backup windows
- Greatly improved restore speeds, producing quicker recoveries and decreased downtime
- Increased Celerra scalability by offloading I/O processing

## Faster backup speeds

Backups on MPFS improve performance up to 140 percent over legacy NFS. As we scale the number of backup channels, the backup window reduces in both hot and cold backups. The scalability on MPFS is better over legacy NFS. The performance can be improved by using multiple threads.

Using MPFS provides significant improvement over NFS for both cold and hot backups as well as restores.

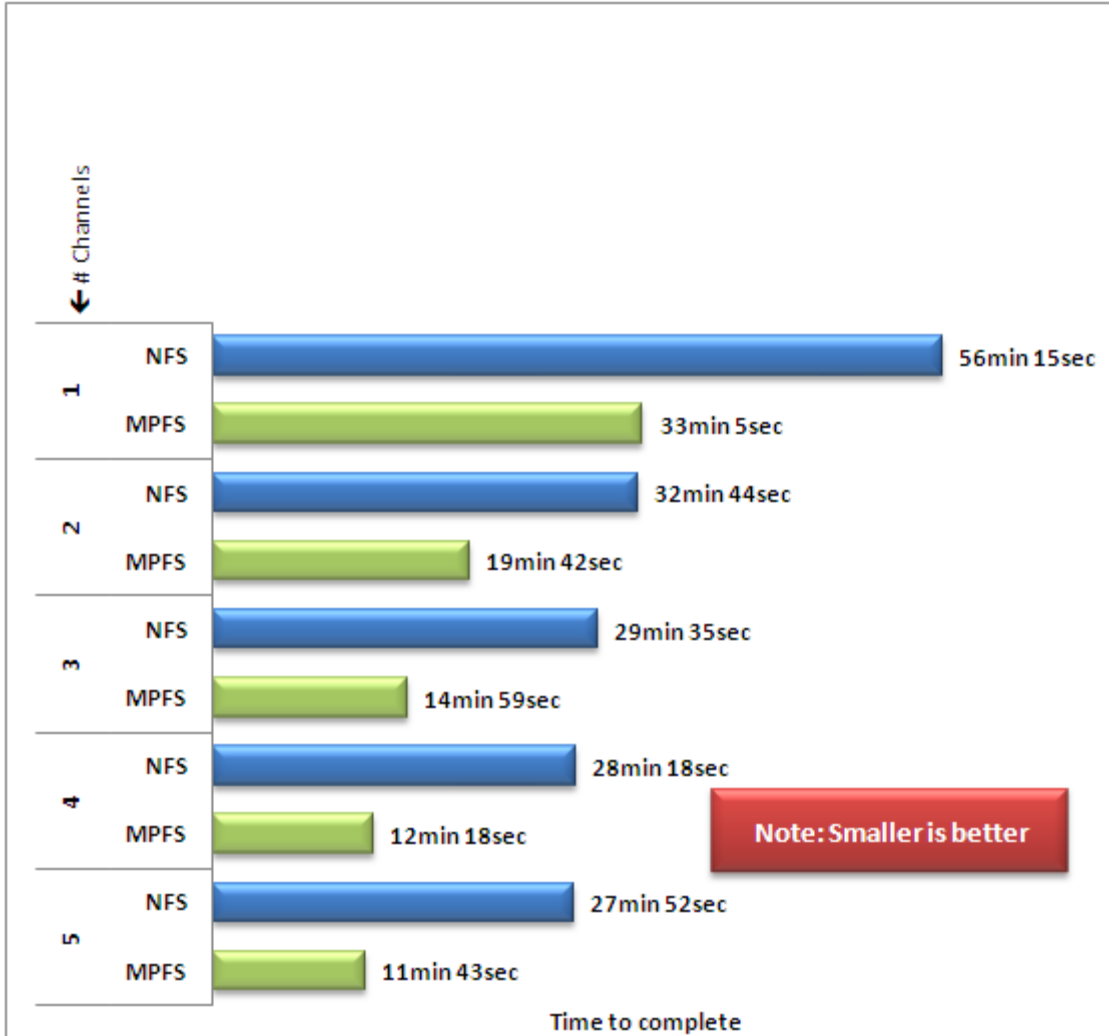


Figure 3 RMAN cold backup with multiple channels

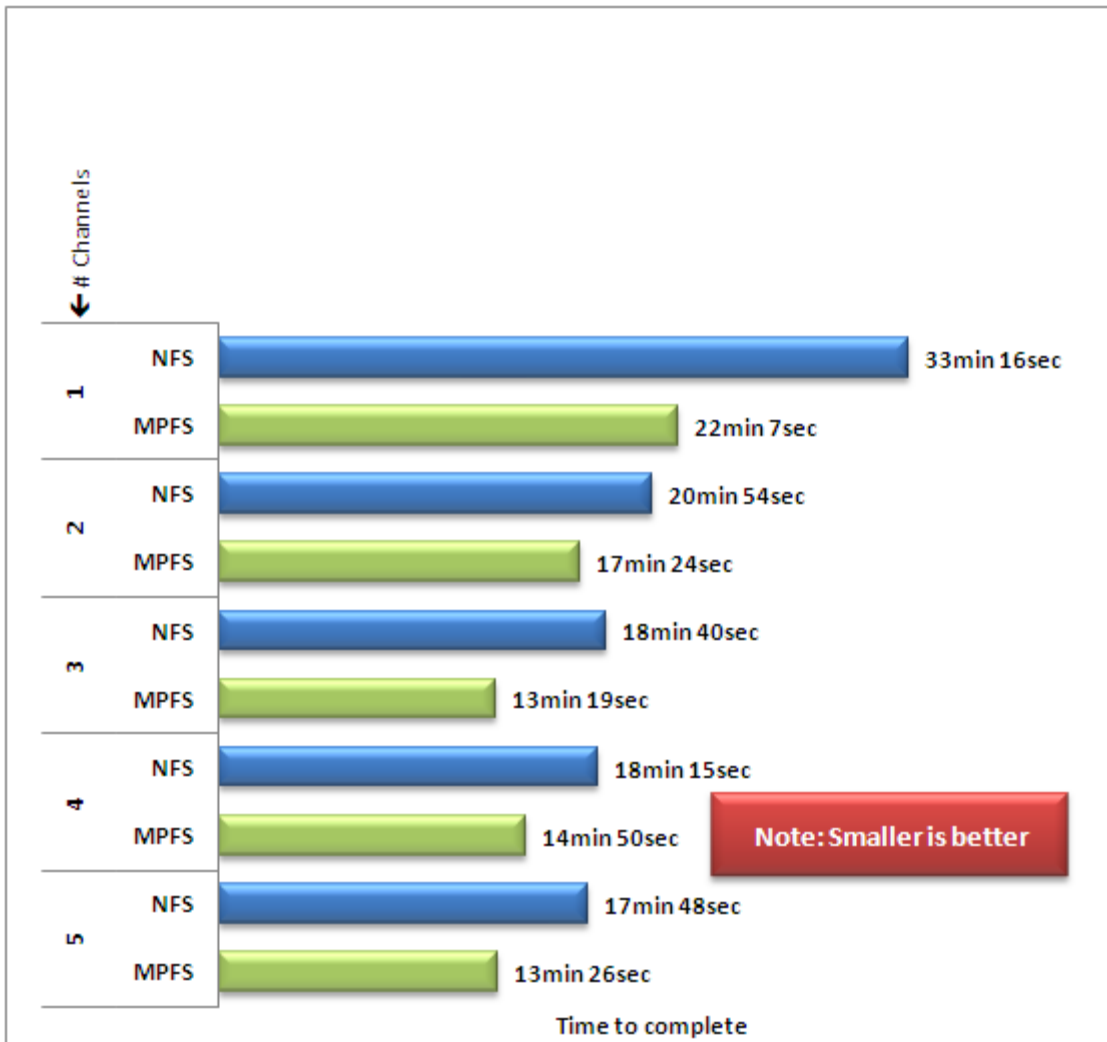


Figure 4 RMAN hot backup window with multiple channels

## Improved restore speeds

Restore on MPFS improves performance up to 240 percent over legacy NFS. Database restore performance over MPFS improved 240 percent over legacy NFS using multiple channels. The benefit to organizations is that faster recovery time reduces the application downtime, and saves money.

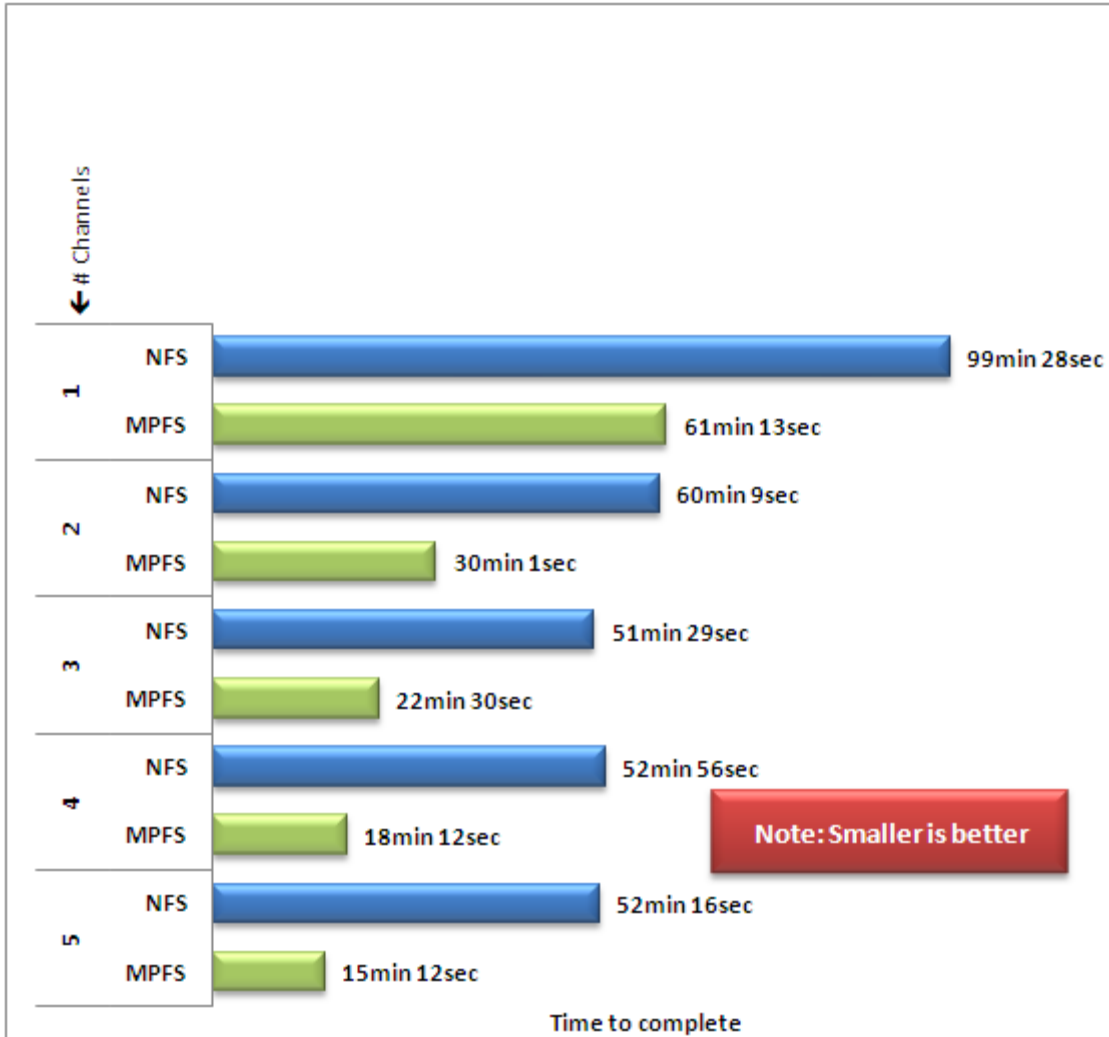


Figure 5 RMAN restore speed from cold backup with multiple channels

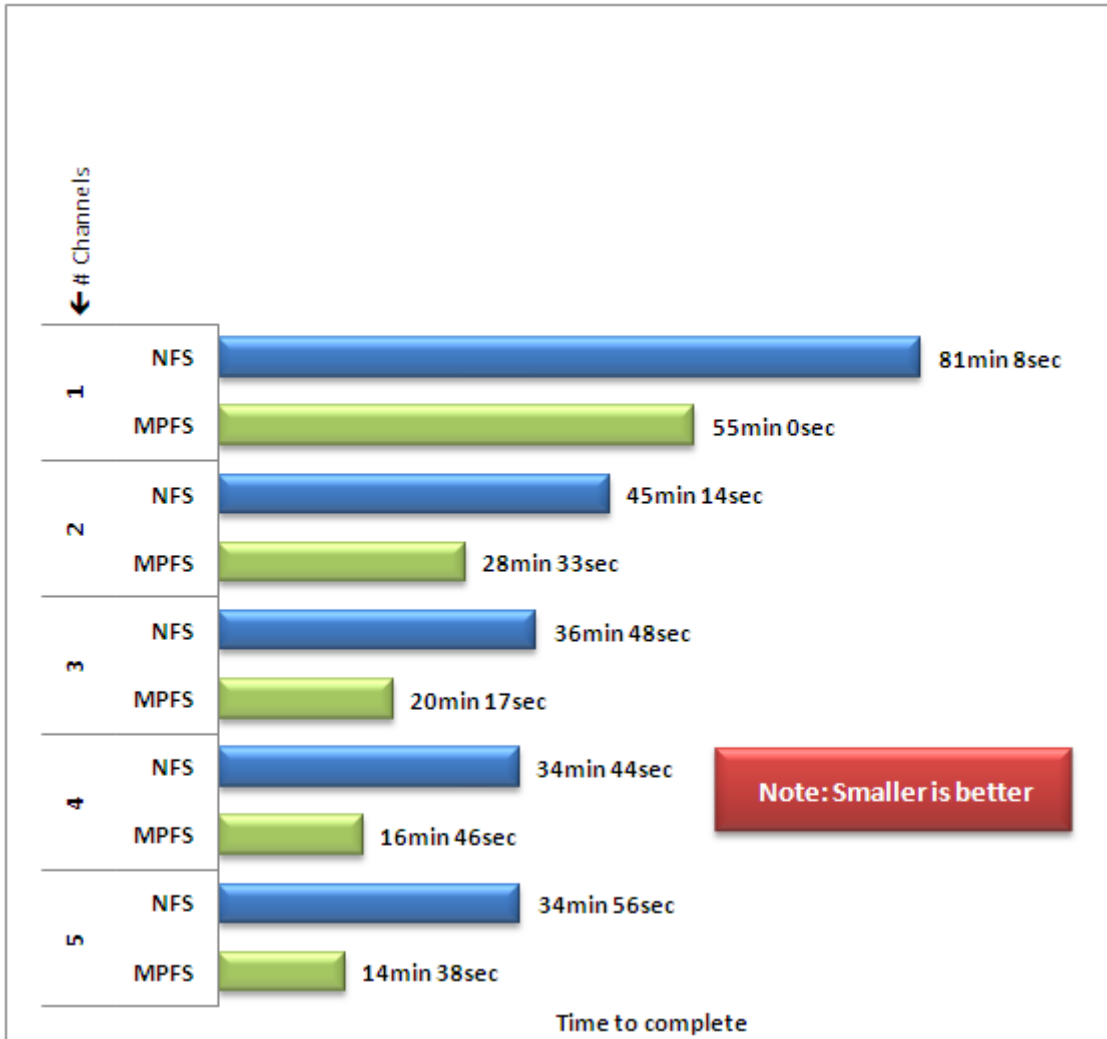


Figure 6 RMAN restore speed from hot backup with multiple channels

## Improved EMC Celerra usage

MPFS improves performance and scalability by offloading data processing and parallelizing storage array access. This reduces the Celerra Data Mover workload. The Data Mover CPU utilization and network throughput with MPFS are orders of magnitude less for RMAN backup and restore operations compared to the legacy NFS environment.




Figure 7 Data Mover percentage utilization

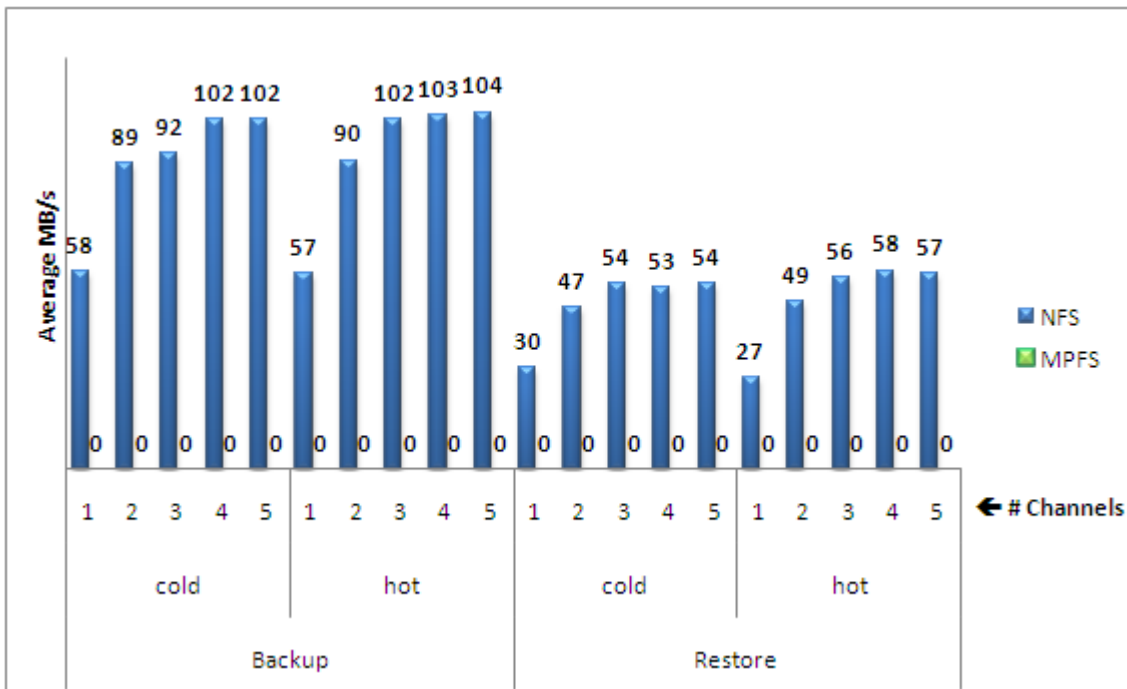


Figure 8 Data Mover throughput in megabytes per second

---

## Testing strategy

This section provides details about the core tests completed on this storage configuration and lists the following configuration details:

- [Hardware and software resources](#)
- [RAID layout](#)
- [File system layout](#)

In the production environment, we used 1,000 warehouses with fifty 3 GB database files. We then mounted the shared file system on both the production and test and development servers. The following testing dimensions, performed in the order listed, were the foundation of our testing.

1. Performed RMAN cold backup (database in the “mount” state) with RMAN on Oracle data files on the production server with the production database files stored on CLARiiON LUNs. RMAN wrote its backups to the Celerra shared file system.

---

**Note:** Use the **Backup as copy database** command to obtain a block-level copy of all database, system, sysaux, and undo table space and user files to the shared file system.

---

2. Restored the Oracle data files from the shared file system to the test/development server. Then, we reset the Oracle logs and opened the database in the test /development server for service.
3. Started a load simulation of 500 users against the production database.
4. Performed RMAN hot backup from the production server to the shared file system. Backup of the database included the archive logs this time, since it is a hot backup.

---

**Note:** Backup of the Oracle database to a shared file system using legacy NFS and MPFS requires the following mount options:

```
rw,hard,nointr,rsize=32768,wspace=32768,tcp,noac,nfsvers=3
```

In addition, administrators must enable the DirectIO flag in MPFS for RMAN to work within the MPFS environment (defined as `globDirectIO=1` in `/etc/mpfs.conf`).

---

5. Restored the database and archive logs from the shared file system, recovered the database by committing the archive logs, reset the logs, and opened the database on the test/development server for service.

---

**Note:** When a database is restored from cold backup it is in sync state. Therefore, you do not need to recover the database by applying the archivelog.

When the database is restored from the hot backup, the changes made to database during the backup need to be applied using the archive logs. In hot backup mode, recovery of database is required.

---

---

## Hardware and software resources

Table 2 lists the hardware resources for this tested storage configuration:

**Table 2 Hardware resources**

Hardware	Quantity	Configuration
EMC Celerra NS40FC	One	<ul style="list-style-type: none"><li>Two X-blades, one active and one standby</li><li>One Control Station</li><li>120 750 GB 7.2k rpm SATA HDDs</li><li>Six 300 GB 15k rpm 4 Gb FC HDDs</li></ul>
EMC CLARiiON CX3-40	One	<ul style="list-style-type: none"><li>Two DAEs</li><li>30 300 GB 15k rpm 4 Gb FC HDDs</li></ul>
2U Intel server	Two	<ul style="list-style-type: none"><li>Two 3.0 GHz dual-core Xeon processors, 64-bit</li><li>16 GB RAM</li><li>Two Onboard Intel NICs</li><li>QLogic QLA2464 FC HBA (dual port)</li></ul>
1U Intel server	One	<ul style="list-style-type: none"><li>Two 3.0 GHz Xeon processors</li><li>4 GB RAM</li></ul>
Gigabit Ethernet switch	One	<ul style="list-style-type: none"><li>Cisco Catalyst 6509</li></ul>
4 Gb/s Fibre Channel switch	One	<ul style="list-style-type: none"><li>Cisco MDS 9513</li></ul>

Table 3 lists the software resources for this tested storage configuration:

**Table 3 Software resources**

Software	Quantity	Configuration
DART release 5.6.37.6	One	Installed on NS40FC
Red Hat Enterprise Linux 4 Update 5 64-bit Kernel version 2.6	Two	Installed on the 2U servers
MPFS Linux client 5.0.20.7	Two	Installed on the 2U servers
Oracle 10g 64-bit	Two	Installed Oracle on the servers
Oracle 10.0.2.3 Patch	Two	Installed the Oracle Patches on servers
QLogic driver 8.01.07.15	Two	Installed for the HBAs
Microsoft Windows 2003 server	One	Installed on the 1U server
Oracle Client	One	Installed on the 1U server
Qwest Benchmark factory	One	Installed on the 1U server
Access Logix™	One	Installed on NS40FC SPs

## RAID layout

Figure 9 shows the RAID layout on an NS40FC tested for RMAN backup speeds over both MPFS and NFS. The array has six FC spindles and 120 SATA spindles. The Celerra Control LUNs (CCL) and FLARE<sup>®</sup> were on the FC spindles — a total of 24 RAID groups were created with SATA spindles and two LUNs were bound per RAID group. We alternated LUN ownership to statically distribute the load across SP A and SP B.

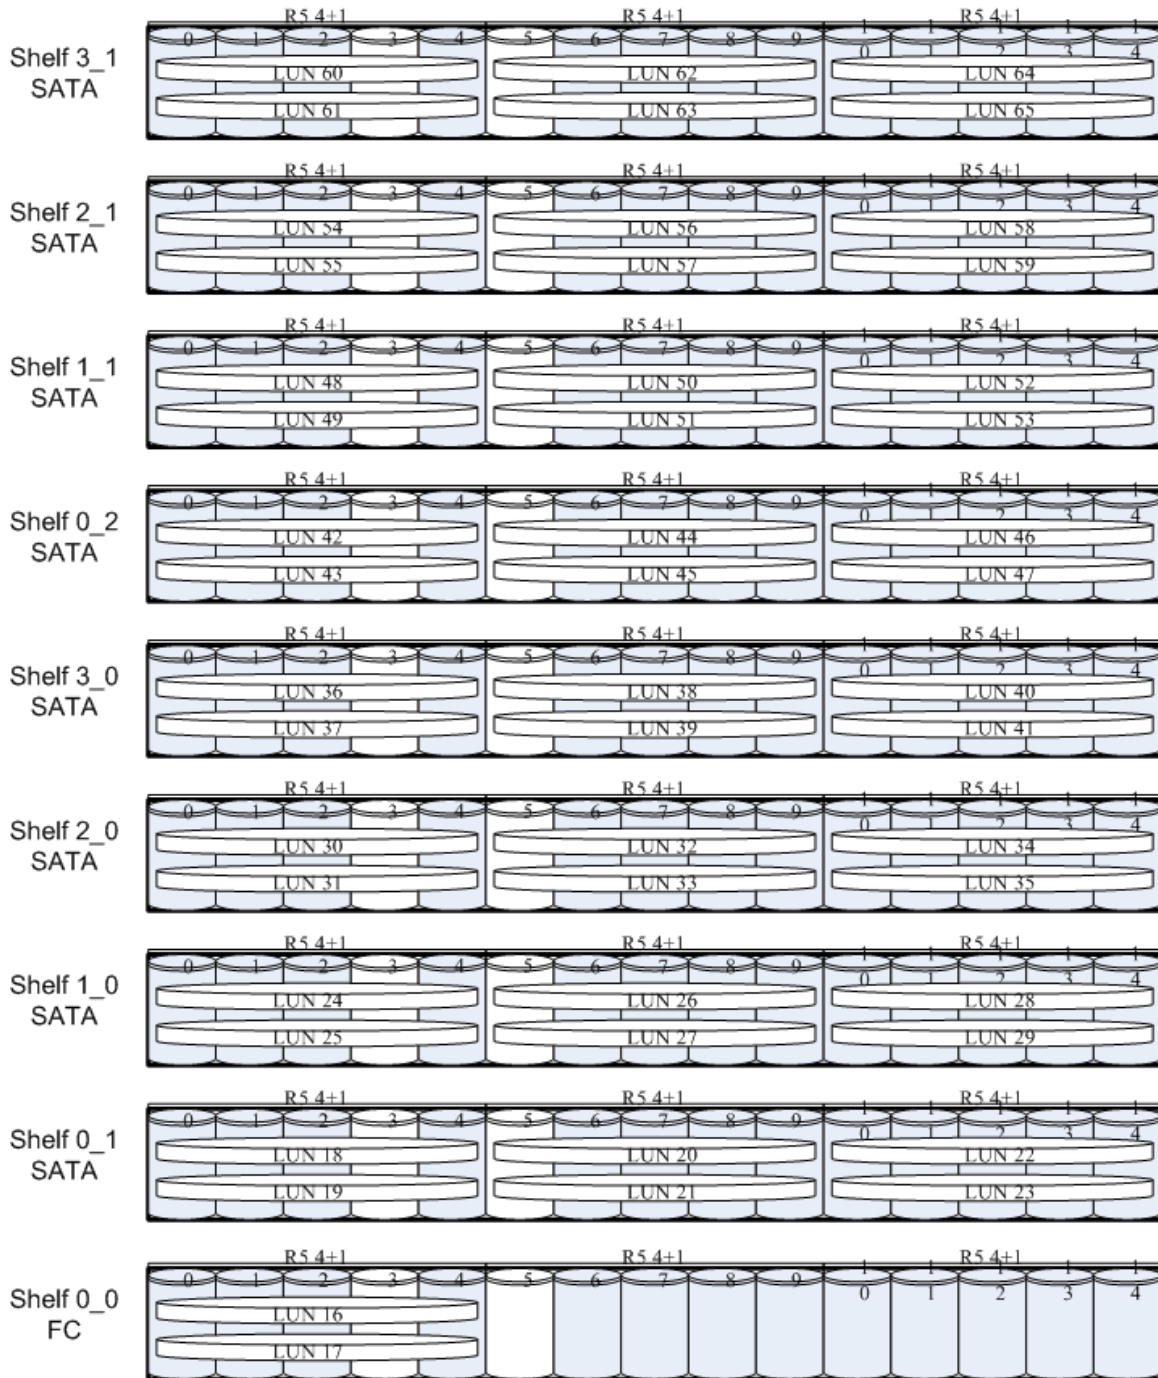



Figure 9 RAID layout in the tested environment

---

## File system layout

Figure 10 shows two striped volumes that were striped across 24 LUNs in a way that balanced them across the storage processors. Those stripes were concatenated into the Meta1 metavolume. That metavolume was used when defining a new user-defined storage pool. A single file system of 16 TB, called Lr1, was created from the user-defined storage pool.



**Figure 10** File system volume layout

---

## Conclusion

Celerra IP storage solutions offer significant management, backup, and recovery advantages. There are also significant advantages in terms of cost, performance, and availability. Empirical data from the validation tests support observations already made by Celerra customers running similar configuration in live, Oracle production environments. MPFS performs significantly better than NFS for RMAN backup and restore with both hot and cold backups. The cold backup on MPFS shows the highest levels of throughput, and the backup window is the smallest of all the tested methods. Restoring the database from the cold backup, using MPFS, speeds recovery time. MPFS also provides performance benefits in hot backup mode, when database applications must remain available during the backup windows. Celerra and MPFS are excellent choices for reducing RMAN backup and restore windows and saving your organization money.

EMC offers customers a wide variety of benefits with RMAN using Celerra with MPFS:

- Reduced TCO by utilizing the no charge (disk-based backup) version of RMAN
- Management and cost benefits offered by Celerra storage consolidation
- Significantly faster backup speeds produce shorter backup windows
- Improved restore speeds produce quicker recoveries and decreased downtime
- Stronger SLA with shorter backup maintenance windows

## References

The following documents, located on Powerlink, provide additional, relevant information about MPFS and RMAN backups. Access to these documents is based on your login credentials. If you do not have access to the following content, contact your EMC representative

- *EMC Celerra MPFS over FC and iSCSI v5.0 Linux Client Product Guide.*
- *EMC Celerra MPFS Integrated Configurations Quick Start Guide*
- *EMC Solutions for Backup to Disk Solutions Overview*