

**EMC Backup and Recovery for  
Oracle 11g OLTP**

Enabled by EMC CLARiiON,  
EMC Data Domain, EMC NetWorker,  
and Oracle Recovery Manager  
using Fibre Channel

Reference Architecture



**EMC Global Solutions**

Copyright © 2010 EMC Corporation. All rights reserved.

Published February, 2010

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

Benchmark results are highly dependent upon workload, specific application requirements, and system design and implementation. Relative system performance will vary as a result of these and other factors. Therefore, this workload should not be used as a substitute for a specific customer application benchmark when critical capacity planning and/or product evaluation decisions are contemplated.

All performance data contained in this report was obtained in a rigorously controlled environment. Results obtained in other operating environments may vary significantly.

EMC Corporation does not warrant or represent that a user can or will achieve similar performance expressed in transactions per minute.

No warranty of system performance or price/performance is expressed or implied in this document. Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on [EMC.com](http://EMC.com).

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

Part number: H6835

## Table of Contents

---

Reference architecture overview .....	4
Key components .....	6
Physical architecture.....	10
Validated environment profile .....	11
Hardware and software resources.....	11
Testing and validation.....	13
Conclusion .....	20

---

## Reference architecture overview

---

### Document purpose

This document describes the reference architecture of the EMC Backup and Recovery for Oracle 11g OLTP Enabled by EMC CLARiiON, EMC Data Domain, EMC NetWorker, and Oracle Recovery Manager using Fibre Channel solution.

---

### Solution purpose

EMC Proven™ Solutions are designed to build a replica of some aspects of a customer environment. Information about the creation of the environment and its operation is provided to customers, field personnel, and development organizations in the related Proven Solution Guides.

The solution demonstrates the backup and recovery service (BRS) capabilities of EMC® hardware and software in an enterprise Oracle Database 11g RAC environment. Specifically, this solution provides customers with the best practices for backing up and restoring an Oracle 11g database. This is accomplished by focusing on the advantages and capabilities of CLARiiON® technology, EMC Data Domain®, and EMC NetWorker® software.

---

### The business challenge

Today's IT is being challenged by the business to solve the following pain points around the backup and recovery of the business' critical data:

- Protect the business information as an asset of the business' defined recovery point objective (RPO - amount of data to recover) and recovery time objective (RTO - time to recover)
- Efficient use of both infrastructure and people to support the business
- Difficulties around the backing up of large enterprise-critical systems – multi-terabyte systems

Exponential data growth, changing regulatory requirements, and increasingly complex IT infrastructure all have a major impact on data managers' data protection schemes. Recovery time objectives (RTO) continue to decrease while the precision of the recovery point objective (RPO) increases. In other words, IT managers must be able to recover from a given failure quicker than ever and with less data loss. It is not uncommon for organizations to routinely exceed their backup window or even have a backup window that takes up most of the day. Such long backup operations leave little margin for error and any disruption can place some of the data at risk of loss. Such operations also mean that a guaranteed RPO cannot be met.

Because of the demands generated by data growth and the RTO/RPO requirements in Oracle database environments, it is critical that robust, reliable, and tested backup and recovery processes are in place. Backup and recovery of Oracle databases are a vital part of IT data protection strategies. To meet these backup and recovery challenges enterprises need proven solution architectures that encompass the best of what EMC and Oracle can offer.

---

---

**The technology solution**

This solution describes a backup and recovery environment for an Oracle 11g OLTP database. The database was configured on a two-node RAC using ASM. Backup and recovery was implemented using RMAN, EMC NetWorker, and an EMC Data Domain DD880 appliance.

The backup process was off-loaded to a NetWorker proxy host using a Navisphere® SnapView™ clone to minimize the impact to the production environment. The DD880 appliance enabled an 84 percent saving on the storage capacity required by the backup process. Reference the “Storage saving % after 5 Weeks of backup cycle” graphic in the “Testing and validation” section.

The following table describes the key components and their configuration details within this environment.

Component	Configuration	Software
CLARiiON CX4-960	Four BE 4 Gb FC ports, eight FE 4 Gb FC ports per storage processor, nine DAEs with five 146 GB and 130 x 300 GB disk drives	FLARE® 04.29.000.5.003
Data Domain DD880	Four 4 Gb FC ports - VTL SAN connectivity, two SAS HBAs - disk connectivity, three EOS disk shelves with 48 disks	DDOS 4.7.1.3
Oracle 11g OLTP database system	1 TB Oracle 11g OLTP database on a two-node RAC using ASM	Oracle 11g Database/Cluster/ASM versions 11.1.0.7
NetWorker	NetWorker Management Console, storage nodes, clients	NetWorker 7.6

---

## Key components

---

### Introduction

This section briefly describes the key solutions components. For details on all of the components that make up the reference architecture, refer to the hardware and software sections.

---

### CLARiiON CX4-960

The EMC CLARiiON CX4 model 960 enables you to handle the most data-intensive workloads and large consolidation projects. CLARiiON CX4-960 delivers innovative technologies such as Flash drives, Virtual Provisioning™, 64-bit operating system, and multi-core processors.

The CX4's new flexible I/O module design, UltraFlex™ technology, delivers an easily customizable storage system. Additional connection ports can be added to expand connection paths from servers to the CLARiiON. The CX4-960 can be populated with up to six I/O modules per storage processor.

The CX4-960 also uses a new generation of storage processor CPUs, memory, and PCI bus architecture.

---

### EMC Data Domain DD880

EMC Data Domain deduplication storage systems dramatically reduce the amount of disk storage needed to retain and protect enterprise data. By identifying redundant data as it is being stored, Data Domain systems provide a storage footprint that is five to 30 times smaller, on average, than the original dataset. Backup data can then be efficiently replicated and retrieved over existing networks for streamlined disaster recovery and consolidated tape operations.

The Data Domain DD880 is the industry's highest throughput, most cost-effective and scalable deduplication storage solution for disk backup and network-based disaster recovery (DR).

The high-throughput inline deduplication data rate of the DD880 is enabled by the Data Domain Stream-Informed Segment Layout (SISL) scaling architecture. The level of throughput is achieved by a CPU-centric approach to deduplication, which minimizes the number of disk spindles required.

---

### Navisphere Management Suite

The Navisphere Management Suite of integrated software tools allows you to manage, discover, monitor, and configure EMC CLARiiON systems as well as control all platform replication applications from an easy-to-use, secure, web-based management console.

Navisphere Management Suite enables you to access and manage all CLARiiON advanced software functionality—including EMC Navisphere Quality of Service Manager, Navisphere Analyzer, SnapView, SAN Copy™, and MirrorView™. When used with other EMC storage management software, you gain storage resource, SAN, and replication management functionality—for greater efficiency and control over CLARiiON storage infrastructure.

---

---

**EMC  
PowerPath**

EMC PowerPath® is a server-resident software that enhances performance and application availability. PowerPath works with the storage system to intelligently manage I/O paths, and supports multiple paths to a logical device. In this solution, PowerPath manages I/O paths and provides:

- Automatic failover in the event of a hardware failure. PowerPath automatically detects path failure and redirects I/O to another path.
- Dynamic multipath load balancing. PowerPath intelligently distributes I/O requests to a logical device across all available paths, thus improving I/O performance and reducing management time and downtime by eliminating the need to configure paths statically across logical devices.

PowerPath enables customers to standardize on a single multipathing solution across their entire environment.

---

**EMC NetWorker**

EMC NetWorker software comprises a high-capacity, easy-to-use data storage management solution that protects and helps to manage data across an entire network. NetWorker simplifies the storage management process and reduces the administrative burden by automating and centralizing data storage operations.

**NetWorker Module for Oracle**

NMO provides the capability to integrate database and file system backups, to relieve the burden of backup from the database administrator while allowing the administrator to retain control of the restore process. NMO includes the following features:

- Automatic database storage management through automated scheduling, autochanger support, electronic tape labeling, and tracking.
- Support for backup to a centralized backup server.
- High performance through support for multiple, concurrent high-speed devices, such as digital linear tape (DLT) drives.

Together with the NetWorker server, NMO augments the backup and recovery system provided by the Oracle server and provides a storage management solution that addresses the need for cross-platform support of enterprise applications.

---

**Oracle  
Database 11g  
Enterprise  
Edition**

Oracle Database 11g Enterprise Edition delivers industry-leading performance, scalability, security, and reliability on a choice of clustered or single servers running Windows, Linux, and UNIX. It provides comprehensive features easily managing the most demanding transaction processing, business intelligence, and content management applications.

Oracle Database 11g Enterprise Edition comes with a wide range of options to help grow your business and meet users' performance, security, and availability service level expectations.

## Oracle Database 11g RAC

Oracle Real Application Clusters (RAC) is an optional feature of Oracle Database 11g Enterprise Edition. Oracle RAC supports the transparent deployment of a single database across a cluster of servers, providing fault tolerance from hardware failures or planned outages. If a node in the cluster fails, Oracle continues running on the remaining nodes. If more processing power is needed, new nodes can be added to the cluster providing horizontal scaling.

Oracle RAC supports mainstream business applications of all kinds. This includes Online Transaction Processing (OLTP) and Decision Support System (DSS).

- Decision Support describes the capability of a system to support the formulation of business decisions through complex queries against a database. Decision support applications typically consist of long and often complex read-only queries that access large portions of the database.

## Oracle ASM

Oracle Automatic Storage Management (ASM) is an integrated database filesystem and disk manager. ASM filesystem and volume management capabilities are built into the Oracle database kernel. And it can reduce the complexity of managing the storage for the database.

In addition to providing performance and reliability benefits, ASM can also increase database availability because disks can be added or removed without shutting down the database. ASM automatically rebalances the database files across an ASM diskgroup after disks have been added or removed.

## Oracle ASMLib

ASMLib is a support library for the ASM feature of Oracle Database. It is an add-on module that simplifies the management and discovery of ASM disks. The ASMLib provides an alternative to the standard operating system interface for ASM to identify and access block devices.

ASMLib is composed of the actual ASMLib library, which is loaded by Oracle at Oracle startup, and a kernel driver that is loaded into the OS kernel at system boot. The kernel driver is specific to the OS kernel.

## ASM CMD

The **asmcmd** utility can be used by Oracle database administrators to query and manage their ASM systems. ASM-related information can be retrieved easily for diagnosing and debugging purposes.

## Oracle RMAN

Oracle Recovery Manager (RMAN) is a command-line and Enterprise Manager-based tool for backing up and recovering an Oracle database. It provides block-level corruption detection during backup and restore. RMAN optimizes performance and space consumption during backup with file multiplexing and backup set compression, and integrates with Oracle Secure Backup and third-party media management products for tape backup.

### **Oracle Block Change Tracking**

This database option causes Oracle to track datafile blocks affected by each database update. The tracking information is stored in a block change tracking file. When block change tracking is enabled, RMAN uses the record of changed blocks from the change tracking file to improve incremental backup performance by only reading those blocks known to have changed, instead of reading datafiles in their entirety.

---

### **Swingbench**

Swingbench is a publicly available load generator (and benchmark tool) designed to stress test Oracle databases. Swingbench consists of a load generator, a coordinator, and a cluster overview. The software enables a load to be generated and the transactions/response times to be charted.

Swingbench is provided with four benchmarks:

- OrderEntry – TPC-C-like workload.
- Calling Circle – Telco-based self-service workload.
- Stress Test – Performs simple insert/update/delete/select operations.
- DSS – A DSS workload, based on the Oracle Sales History schema.

The Swingbench workload used in this testing was Order Entry. The Order Entry (PL/SQL) workload models the classic order entry stress test. It has a profile similar to the TPC-C benchmark. It models an online order entry system, with users being required to log in before purchasing goods.

---

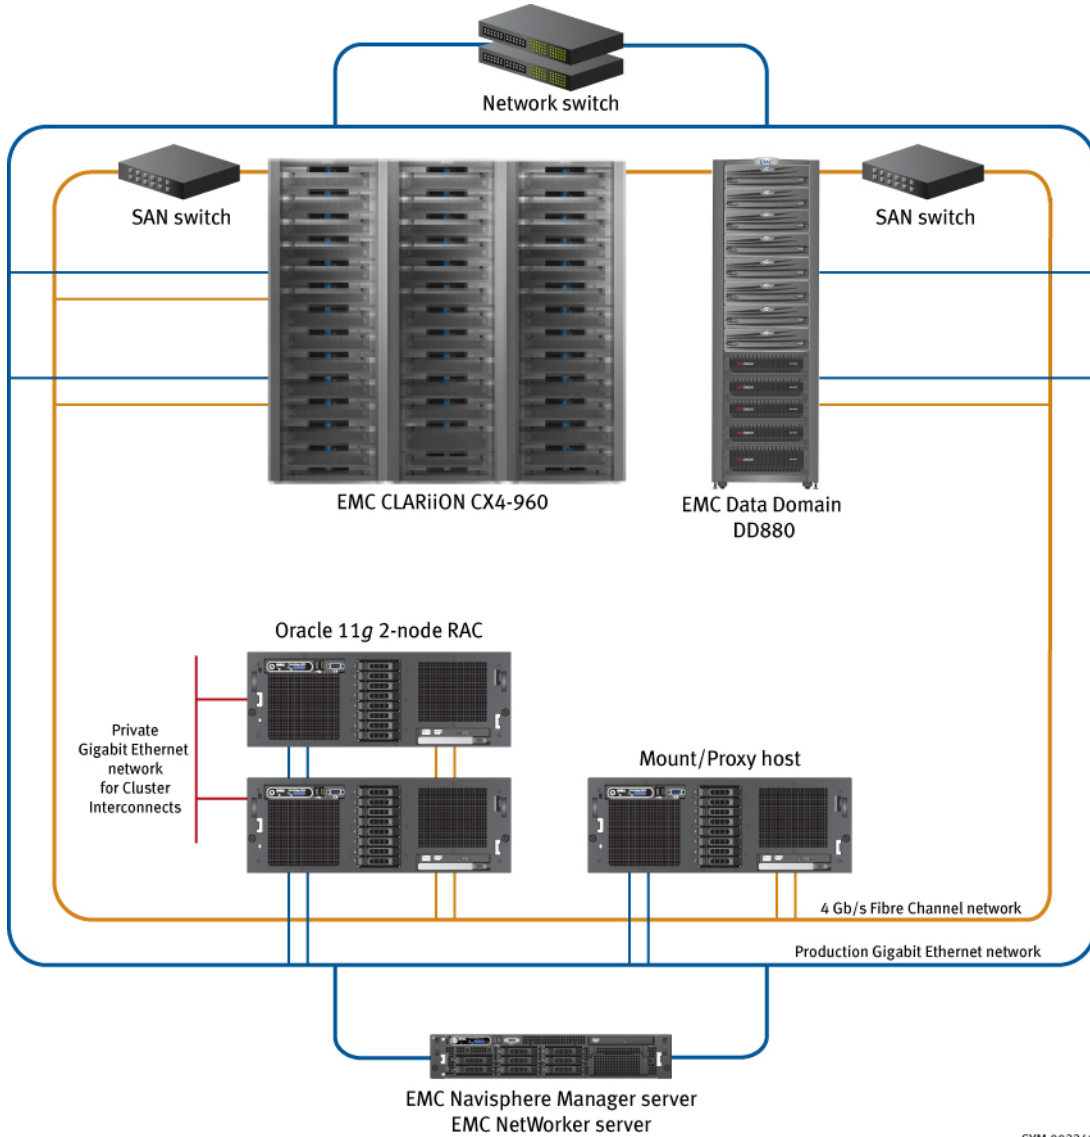
---

## Physical architecture

---

### Architecture diagram

The following image depicts the overall physical architecture of the solution.



SYM-002240

## Validated environment profile

---

**Profile characteristics** The solution was validated with the following environment profile.

Profile characteristic	Value
Database characteristic	OLTP
Benchmark profile	Swingbench OrderEntry - TPC-C-like benchmark
Target Oracle transaction response time	< 10 ms
Read / Write ratio	70 / 30
Database scale	A Swingbench load that keeps the system running within agreed performance limits
Size of databases	1 TB
Number of databases	1
Initial datafile setup	As required by Swingbench
Array drives: size and speed	300 GB 15k rpm

---

## Hardware and software resources

---

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

Equipment	Quantity	Configuration
<b>Storage array</b>	1	CLARiiON CX4-960: <ul style="list-style-type: none"><li>• 9 DAEs</li><li>• 5 x 146 GB FC drives</li><li>• 130 x 300 GB drives</li></ul>
<b>SAN</b>	2	4 Gb capable FC switch, 64 port
<b>Oracle database server</b>	2	Four Quad-Core Xeon E7330 processors, 2.4 GHz, 6 MB, 1066 FSB, 32 GB RAM. Two 73 GB 10k internal disks. Four 4 Gb Emulex LP11002E HBAs.
<b>Proxy server</b>	1	Four Quad-Core Xeon E7330 processors, 2.4 GHz, 6 MB, 1066 FSB, 32 GB RAM. Two 73 GB 10k internal disks. Four 4 Gb Emulex LP11002E HBAs.
<b>Navisphere management server</b> <b>NetWorker server</b>	1	Two Quad-Core processors, 1.86 GHz, 16 GB RAM. Two 4 Gb Emulex LP11002E HBAs.

---

**Software**

The software used to validate the solution is listed below.


<b>Software</b>	<b>Version</b>	<b>Comment</b>
RedHat Linux	5.3	OS for database server
Microsoft Windows	2003 SP2	OS for Navisphere Management Server
Oracle Database/Cluster/ASM	11g Release 1 (11.1.0.7.0)	Database/cluster software/volume management
Oracle ASMLib	2.0	Support library for ASM
Swingbench	2.3	OLTP database benchmark
Orion	10.2	Orion is the Oracle I/O Numbers Calibration Tool designed to simulate Oracle I/O workloads
FLARE operating environment	04.29.000.5.003	FLARE
Navisphere Management Suite		Includes: <ul style="list-style-type: none"><li>• Access Logix™</li><li>• Navisphere Agent</li></ul>
Navisphere Analyzer	6.29.0.6.34	
SnapView	6.29.0.6.34.1	
PowerPath	5.3	Multipathing software
DDOS	4.7.1.3	Data Domain OS
NetWorker	7.6	Backup and recovery suite
NetWorker Module for Oracle	5.0	NetWorker Oracle integration

---

## Testing and validation

### Introduction

Once the disk environment was set up on the CLARiiON CX4-960, the disk configuration was validated using Orion, an Oracle toolset.



The previous graphic shows a high-level view of the CLARiiON RAID group layout for the Oracle ASM disk groups.

Drive Type	Quantity	Purpose	MetaLUN
146 GB	5	FLARE Vault	No
300 GB	5	Redo	No
300 GB	4	Hot Spares	No
300 GB	40	ASM Data Group	Yes
300 GB	10	ASM Temp/Undo Group	No
300 GB	10	FRA Group	No
300 GB	60	Clone Copies	Mixed


## Results and findings

The results show that there is significantly less impact to production nodes by offloading backups to a proxy host. Navisphere SnapView clones were used to take a copy of the database, which was then mounted on the proxy node.


The following charts illustrate:

- Less impact to production users as a penalty was removed from the production nodes.
- SnapView clone replication had reduced impact on the production nodes.
- An 84% saving of disk storage space required for RMAN backup savesets, which yields a reduction ration of 6:1

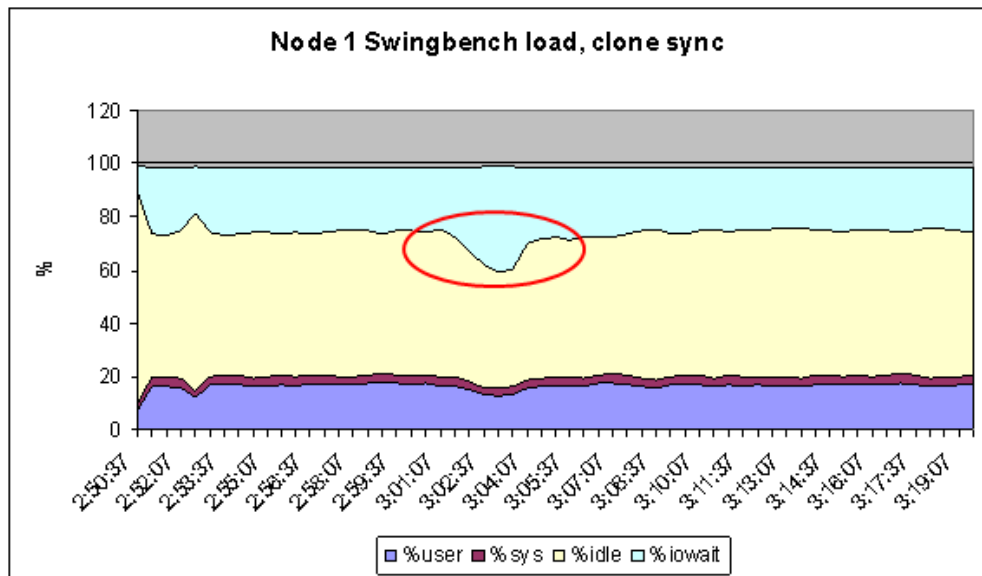
This chart shows one of the RAC nodes while running under a Swingbench simulated production load.




The following chart shows the effect of running a backup concurrently from the production RAC nodes while running the Swingbench production load.



The following chart shows the effect of taking a SnapView clone copy of the production RAC data, while running Swingbench simulated production load against the RAC nodes. The backup is off-loaded to the proxy node, greatly reducing the overall impact of the backup process on production. The clone sync initiated at approximately 3.01 is highlighted. It can be seen that the duration of the impact is greatly reduced.



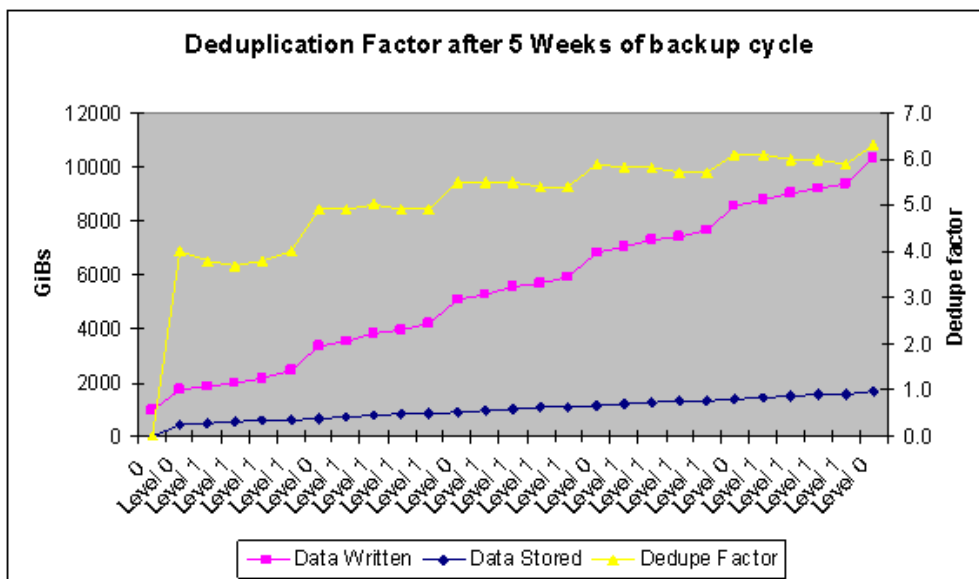
By eliminating redundant data segments, the Data Domain system allows many more backups to be stored and managed than would normally be possible for a traditional storage server. While completely new data has to be written to disk whenever discovered, the variable-length segment deduplication capability of the Data Domain system makes finding identical segments extremely efficient.



The storage saving chart above shows the data written to the DD880 over a five-week period. The backup cycle consisted of an RMAN level 0 (full) backup on the weekends and RMAN level 1 (differential incremental) backups Monday through Thursday.

The "data written" data point on the graph shows that approximately 10 TB was backed up over a five-week period.

The "data stored" data point tracks the unique data stored on the DD880 after inline deduplication. This results in a net saving of 84 percent of storage space required over the five-week period. This is shown by the "% Reduction" data point. This results in a deduplication factor of 6.3:1, as shown in the following chart.



Oracle Block Change Tracking was used to improve incremental backup performance. The database daily change rate is  $\approx 2\%$ . The 2% had a greater effect in real terms due to the growth of the Oracle archive logs, which were also backed up.


**Note**

Archive redo logs are not good deduplication candidates, therefore much better deduplication rates can be achieved if these logs are backed up separately. Deduplication rates in the order of 20x are not uncommon if the archive logs are backed up separately.

**Restore**

Data Domain's Streams Informed Segment Layout (SISL) technology ensures balanced backup and restore speeds.

The following chart shows the time duration of a Level 0 full backup, a Level 1 Incremental differential backup, and a full restore of the database. The backup and restore processes were initiated on the proxy node using eight RMAN channels.



This output shows the restore throughput of the DD880 appliance achieved in this solution. When compared to the subsequent graphic, which shows the throughput achieved by the DD800 appliance during a backup, this further demonstrates the balanced backup and restore speeds.

```

-----
                26      0      0      0      0      0
12/22 11:21:44
Port  UTL      Drive  ops/s  Read KiB/s  Write KiB/s  Soft Errors  Hard Errors
-----
 4a   DD-Library
      0      1      0      0      0      0      0
      1      0      0      0      0      0      0
      2      0      0      0      0      0      0
      7      739    94592    0      0      0      0
      8      0      0      0      0      0      0
      11     0      0      0      0      0      0
      16     0      0      0      0      0      0
      19     0      0      0      0      0      0
      25     0      0      0      0      0      0
 4b   DD-Library
      5      678    86784    0      0      0      0
      6      0      0      0      0      0      0
      14     735    94144    0      0      0      0
      17     0      0      0      0      0      0
      22     0      0      0      0      0      0
      24     0      0      0      0      0      0
 5a   DD-Library
      0      0      0      0      0      0      0
      3      751    96128    0      0      0      0
      4      0      0      0      0      0      0
      12     754    96576    0      0      0      0
      15     0      0      0      0      0      0
      21     0      0      0      0      0      0
      23     0      0      0      0      0      0
 5b   DD-Library
      9      711    91008    0      0      0      0
      10     0      0      0      0      0      0
      13     0      0      0      0      0      0
      18     680    87104    0      0      0      0
      20     0      0      0      0      0      0
      26     0      0      0      0      0      0
-----

```

The previous graphic shows the DD880 appliance throughput during a restore with eight RMAN channels, and a sustained read throughput of over 630 MB/s.

```

-----
12/15 12:39:08
Port  UTL      Drive  ops/s  Read KiB/s  Write KiB/s  Soft Errors  Hard Errors
-----
 4a   DD-Library
      0      1      0      0      0      0      0
      1      0      0      0      0      0      0
      2      0      0      0      0      0      0
      7      0      0      0      0      0      0
      8      772    0      98880    0      0      0
      11     0      0      0      0      0      0
      16     0      0      0      0      0      0
      19     0      0      0      0      0      0
      25     845    0      108224    0      0      0
 4b   DD-Library
      5      0      0      0      0      0      0
      6      788    0      100864    0      0      0
      14     0      0      0      0      0      0
      17     0      0      0      0      0      0
      22     0      0      0      0      0      0
      24     801    0      102528    0      0      0
 5a   DD-Library
      0      0      0      0      0      0      0
      3      0      0      0      0      0      0
      4      770    0      98624    0      0      0
      12     0      0      0      0      0      0
      15     0      0      0      0      0      0
      21     0      0      0      0      0      0
      23     794    0      101632    0      0      0
 5b   DD-Library
      9      0      0      0      0      0      0
      10     838    0      107328    0      0      0
      13     0      0      0      0      0      0
      18     0      0      0      0      0      0
      20     0      0      0      0      0      0
      26     747    0      95680    0      0      0
-----

```

The previous graphic shows the DD880 throughput during a backup with eight RMAN channels, and a sustained write throughput of over 790 MB/s.

Additional detailed information about the testing processes in this solution can be found in the *EMC Backup and Recovery for Oracle 11g OLTP Enabled by EMC CLARiiON, EMC Data Domain, EMC NetWorker, and Oracle Recovery Manager using Fibre Channel Proven Solution Guide*.

---

---

## Conclusion

---

### Summary

This proven solution reference architecture details an Oracle infrastructure design leveraging an EMC CLARiiON CX4-960 array, EMC Data Domain DD880, and EMC NetWorker. Also included are various test results, configuration practices, and recommended specific Oracle storage design layouts that meet both capacity and consolidation requirements. Described in this document are many of the technologies that enable the benefits outlined below.

Traditional hardware compression provides substantial cost savings in Oracle environments. However, in this solution data deduplication before compression has been shown to significantly reduce the amount of data that needs to be stored over an extended period of time. This offers cost savings both from a management standpoint and in the numbers of disks or tapes required by a customer to achieve their long-term backup strategy.

Data deduplication can fundamentally change the way organizations protect backup and nearline data. Deduplication changes the repetitive backup practice of tape, with only unique, new data written to disk. The test results show that, in an environment utilizing RMAN incremental backups, a data deduplication ratio of over 6.3:1, resulting in an 84 percent saving in the storage required to accommodate the backup data, makes it economically practical to retain the savesets for longer periods of time. This reduces the likelihood that a data element must be retrieved from the vault. Both of these factors can significantly improve the RTO.

Although cost savings are generally not the initial reason to consider moving to disk backup and deduplication, financial justification is almost always a prerequisite. With the potential cost savings of disk and deduplication, the justification statement becomes, “we can achieve all of these business benefits and save money.” That is a compelling argument.

The solution meets the business challenges in the following manner:

- **Ability to keep applications up 24x7**
  - Faster backup and restores – meet more aggressive backup windows, and restore your key applications in minutes, not days
  - Reduced backup windows – minimize backup windows to reduce impact on your application and system availability
- **Protect the business information as an asset of the business**
  - Reduced business risk – restore data quickly and accurately with built-in hardware redundancy and RAID protection
  - Reduced backup windows – minimize backup windows to reduce impact on your application and system availability
- **Efficient use of both infrastructure and people to support the business**
  - Improved IT efficiency – save hours of staff time and boost user productivity
  - Correct costs / reduce costs – match infrastructure costs with changing information value via efficient, cost-effective tiered storage

In summary, utilizing the solution components, in particular CLARiiON technology, EMC Data Domain, and EMC NetWorker software, provides customers with the best possible backup solution to prevent both user and business impact. Business can continue as usual, as if there were no backup taking place at all. In customer environments where, more than ever, there is a trend toward 24x7 activity, this is a critical differentiator that EMC can offer.

---

**Next steps**

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

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

---