

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

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

Reference Architecture



EMC Global Solutions

Copyright © 2010 EMC Corporation. All rights reserved.

Published May, 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.

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

Part number: H7087

Table of Contents

Reference architecture overview	4
Key components	6
Physical architecture.....	8
Validated environment profile	9
Hardware and software resources.....	9
NetWorker backup with Data Domain.....	11
Conclusion	13

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 NFS solution.

Solution purpose

The purpose of this document is to:

- Demonstrate the dramatic reduction in the amount of disk storage needed to retain and protect enterprise data enabled by Data Domain
- Determine the reduction of backup impact by offloading the backup to a proxy mount host enabled by SnapView™ clones and NetWorker®
- Validate the improvement in recovery time objective (RTO) when the backup schedule utilizes full backups only

In this solution an EMC® CLARiiON® CX4-960 array is used as the primary storage for the Oracle 11g OLTP database. Oracle Real Application Clusters (RAC) is used to deploy a single database across a cluster of servers.

EMC NetWorker software is an enterprise level cross-platform backup and recovery suite used to provide storage management. The EMC Data Domain® DD880 appliance provides cost-effective, long-term onsite retention and protection of data.

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. RTO continues to decrease while the precision of the 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 of an Oracle 11g OLTP database. The database was deployed on a CLARiiON CX4-960 and demonstrates the ease and power of integrating EMC storage with Oracle Automatic Storage Management (ASM). This was tested in a two-node Oracle RAC configuration. Backup and recovery was implemented using Oracle RMAN, SnapView clones, and EMC NetWorker. The backup was deployed over NFS to an EMC Data Domain DD880 deduplication appliance.

The backup process was off-loaded to an EMC NetWorker proxy host using Navisphere® SnapView clones. The replica clone copy of the database was mounted to the proxy node and backups were then executed on the proxy node. This is also referred to as the “clone mount host.”

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

Component	Description	Configuration	Software
Storage array	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
Deduplication appliance	Data Domain DD880	Two 10 GbE optical NICs, two SAS HBAs - disk connectivity, three ES20 disk shelves with 48 disks	DDOS 4.7.1.3
Database	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
Backup manager	EMC NetWorker	NetWorker server, dedicated storage node and clients	NetWorker 7.6, NetWorker Module for Oracle (NMO) 5.0

Key components

Introduction

This section briefly describes the key solution 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 Oracle workloads and large consolidation projects. CLARiiON CX4-960 delivers innovative technologies such as Flash drives, Virtual Provisioning™, a 64-bit operating system, and multi-core processors.

CLARiiON CX4 is designed to work with Oracle ASM to give DBAs the most comprehensive protection for their Oracle database environment, while maintaining the ease-of-use elements offered by ASM.

EMC Data Domain DD880

EMC Data Domain deduplication storage systems dramatically reduce the amount of disk storage needed to retain and protect enterprise data, including Oracle databases. 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. This allows Data Domain appliances to integrate seamlessly into Oracle architectures, maintaining existing backup strategies such as Oracle RMAN with no changes to scripts, backup processes, or system architecture.

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.

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.

EMC NetWorker, together with the NetWorker Module for Oracle (NMO), provides tight integration with Oracle RMAN and seamlessly uses a Data Domain deduplication appliance as an NFS target for RMAN backups.

These elements create a fast, efficient, and nondisruptive backup that offloads the backup burden from the production RAC environment.

**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.

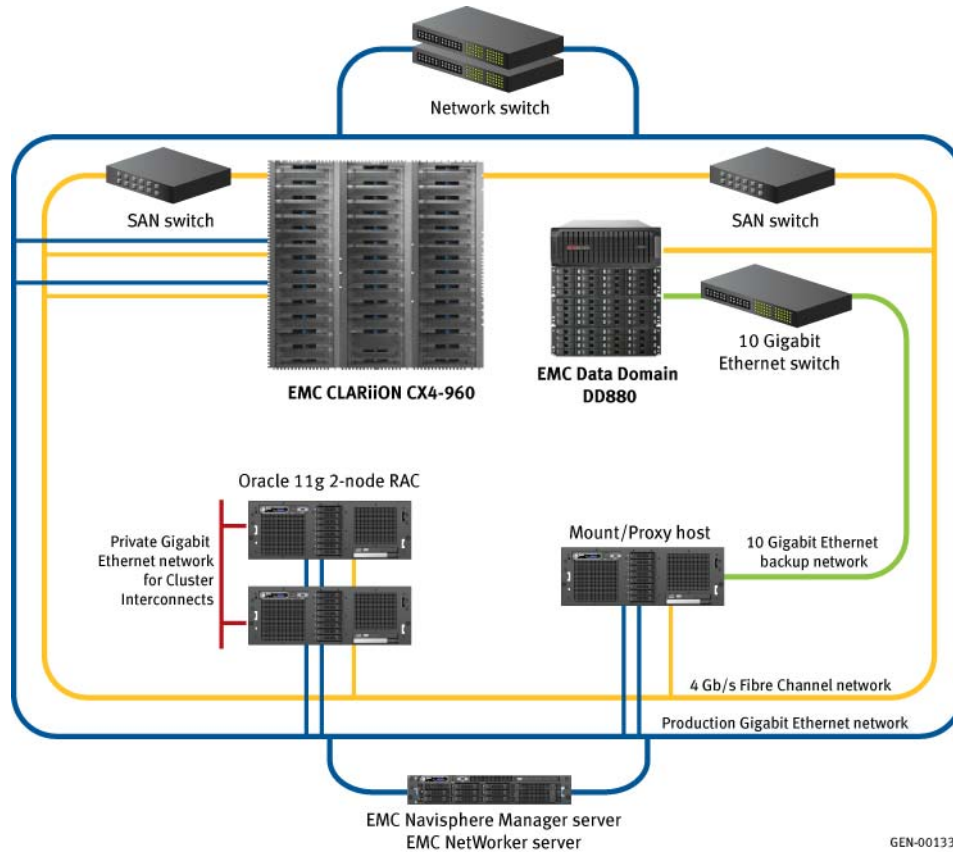
Oracle RMAN

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 backup.

Physical architecture

Architecture diagram

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



Validated environment profile

Profile characteristics EMC Solutions create best practice Oracle environments and these are load tested to simulate real-life database workloads.

Swingbench is a publicly available load generator (and benchmark tool) designed to stress test Oracle databases while operations such as clones, backup, and so on, are conducted.

This solution used Swingbench to simulate a 1 TB Oracle OLTP order entry database with the following characteristics.

Profile characteristic	Value
Database characteristic	OLTP
Benchmark profile	Swingbench OrderEntry - TPC-C-like benchmark
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
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
Storage array	1	CLARiiON CX4-960: <ul style="list-style-type: none"> • 9 DAEs • 5 x 146 GB FC drives • 126 x 300 GB drives • 4 x 300 GB hot spares
SAN	2	4 Gb capable FC switch, 64 port
Oracle database server	2	Four Quad-Core Xeon E7330 processors, 2.4 GHz, 6 MB, 1066 FSB, 32 GB RAM. Two 73 GB 10k internal disks. Four 2 Gb HBAs. Two 10 Gb NICs
Proxy server (mount host)	1	Four Quad-Core Xeon E7330 processors, 2.4 GHz, 6 MB, 1066 FSB, 32 GB RAM. Two 73 GB 10k internal disks. Four 2 Gb HBAs. Two 10 Gb NICs
NetWorker server	1	Two Quad-Core processors, 1.86 GHz, 16 GB RAM. Two 4 Gb HBAs
Network (backup)	1	24 Port 10 GbE switch

Software

The software used to validate the solution is listed below.

Software	Version	Comment
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">• Access Logix™• Navisphere Agent
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
Switch	04.1.00c	10 GbE network
SAN	6.2.0g	SAN

NetWorker backup with Data Domain

Introduction

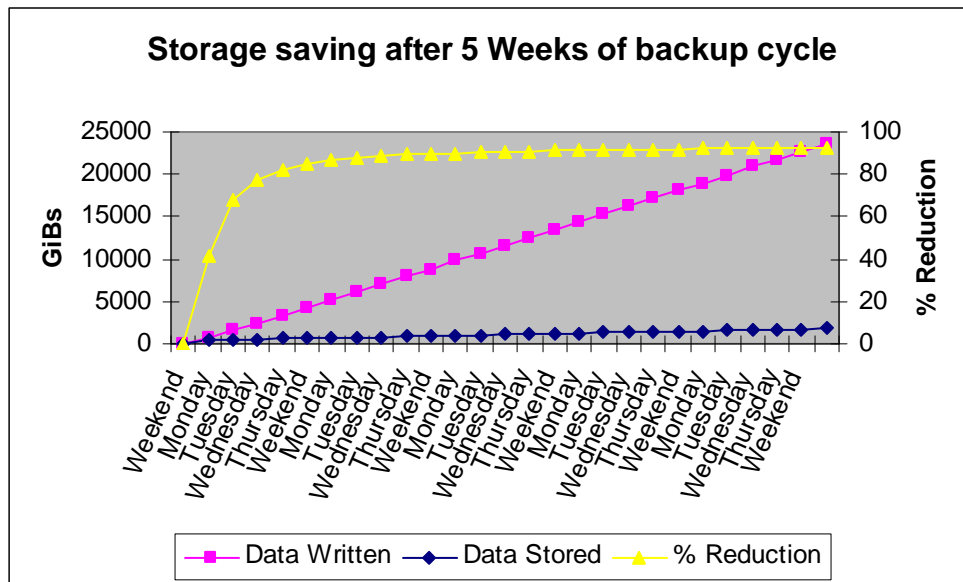
An objective of the testing performed was to validate the reduction of the amount of disk storage needed to retain and protect Oracle enterprise data. To demonstrate the storage saving the backup schedule was run over a number of weeks.

Backup from proxy host

Swingbench was used to generate load against the database instances on the RAC nodes. RMAN Level 0 (full) backups were then executed on the RAC nodes. The performance impact was measured for the duration of the backup. Similar tests were run using SnapView to create a clone copy of the database. The clone copy is a consistent replica of the database that was then mounted to the proxy host. Backups were then executed from the proxy host. The clone operation had much less of an impact on the RAC nodes when compared to running the backup directly from the RAC nodes.

Reduction of disk storage required

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 all RMAN level 0 (full) backups. The "data written" data point on the graph shows that approximately 24 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 92 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 13:1.

Restore

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

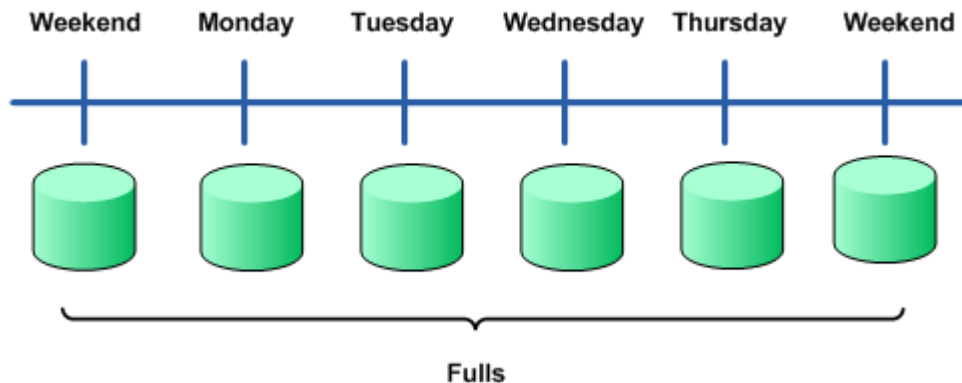
The backup schedule utilizes only full backups every day. This is possible because only unique data is stored on the DD880 appliance. This schedule has the advantage that when a recovery is required for any point in time only a single restore is required as incremental or differential restores are not required. This greatly improves the RTO.

When an incremental backup schedule is implemented, as shown, in the worst case scenario a restore of Thursday's backup is required, then a multi-stage restore operation is necessary.



The full backup from the weekend must first be restored; after this restore is successful then each weekday's incremental backup must be restored one after the other, only Thursday data is restored.

When using a Data Domain deduplication appliance it is possible to implement a full only backup schedule. Therefore only a single restore is required, regardless of where in the schedule the data restore is required.



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. 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 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, therefore the deduplicated backup image does not carry the restore penalty associated with incremental backups because the entire image is always available on the device. This eliminates the need for incremental restores. The test results show that, in an environment utilizing RMAN full backups, a data deduplication ratio of over 13:1, resulting in a 92 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.
