

# EMC Data Domain Boost for Oracle Recovery Manager (RMAN)

A Technical Review

## Abstract

EMC delivers Database Administrators (DBAs) complete control of Oracle backup, recovery, and offsite disaster recovery with advanced integration between EMC Data Domain Boost™ (DD Boost™) and Oracle Recovery Manager® (RMAN). This white paper describes the key considerations to effectively leverage DD Boost for faster, more efficient Oracle backup and recovery. Integration of Oracle RMAN and DD Boost gives the Oracle DBA to have total visibility into and control of backup and disaster recovery processes, without involving the backup administrator.

June 2012

Copyright © 2012 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).

Oracle and Oracle Recovery Manager (RMAN) are registered trademarks or trademarks of Oracle, Inc. in the United States and/or other jurisdictions. All other trademarks used herein are the property of their respective owners.

Part Number h10683

## Contents

<b>Executive Summary</b> .....	<b>4</b>
The Solution .....	4
<b>Introduction</b> .....	<b>5</b>
Audience .....	5
<b>Background</b> .....	<b>5</b>
Oracle Backup and Recovery .....	5
EMC Data Domain Deduplication Systems .....	5
<b>EMC Data Domain Boost – Advanced Integration with Oracle RMAN</b> .....	<b>6</b>
Distributing the Deduplication Process .....	7
RMAN Managed File Replication .....	8
Advanced Load Balancing and Failover.....	9
Features and Benefits Summary .....	10
DD Boost for RMAN Configurations and Capabilities .....	10
<b>Oracle Backup and Recovery Flexibility</b> .....	<b>11</b>
EMC NetWorker and EMC Avamar .....	11
<b>Centralizing Management of Backups</b> .....	<b>12</b>
EMC Data Protection Advisor for Oracle RMAN and Data Domain .....	12
<b>Conclusion</b> .....	<b>13</b>
<b>Appendix</b> .....	<b>14</b>
Best Practice – Optimizing Data Deduplication .....	14
Additional Resources .....	15

## List of Figures

Figure 1: DD Boost for Oracle RMAN Solution Overview .....	7
Figure 2: RMAN Backup to Data Domain without DD Boost .....	7
Figure 3: RMAN Backup to Data Domain with DD Boost .....	8
Figure 4: DD Boost - Managed File Replication Workflow .....	9

## Executive Summary

Businesses and government entities worldwide run Oracle databases to support their mission-critical applications. Not only are these applications the least tolerant of downtime, they often generate high rates of data growth. Oracle applications typically run the business from order entry to payroll and accounts receivable, as well as manufacturing assembly lines. Maintaining service level agreements (SLAs) for Oracle databases is vital to corporate health and sustainability. However, exponential data growth and shrinking backup windows are putting SLAs at risk.

Many IT organizations perform full Oracle backups on a nightly basis – often together with an incremental backup strategy throughout the day. To meet their backup and recovery requirements, most businesses store these backups for 30 days or more. Unfortunately this leads to rapid growth in backup storage requirements, which has kept some users stuck with legacy tape systems as the default solution for database backup. However, this reliance on tape can limit the number of backups that can be performed, impacting recovery point objectives (RPOs).

In addition, DBA's are constantly challenged to improve recovery time objectives (RTOs). Recovering Oracle from full and incremental backups, then rolling the archive / redo logs forward is time consuming and complex. However, restoring the database in the shortest possible time is essential to business operations.

## The Solution

EMC® Data Domain® deduplication storage systems continue to revolutionize disk backup, archiving, and disaster recovery with high-speed, inline deduplication. For nearly a decade, Data Domain systems have brought autonomy and control to Oracle database administrators (DBAs) by supporting Oracle Recovery Manager® (RMAN) backups that write directly to the Data Domain system via a CIFS share or NFS mount without the use of a backup application. The high-speed, inline, deduplication capabilities of Data Domain systems have enabled DBAs to deal with the exponential data growth, and the reliability issues and operational inefficiencies of magnetic tape operations. Data Domain systems provide operationally superior and cost effective disaster recovery.

EMC Data Domain Boost™ integration with Oracle Recovery Manager provides the industry's first solution that gives database administrator complete control of backup and disaster recovery processes. DD Boost for RMAN provides a faster, more efficient Oracle backup and recovery solution.

## Introduction

The purpose of this white paper is to exemplify how DD Boost provides unique and powerful integration between Oracle RMAN and Data Domain systems. The goal is to better serve database administrators in their challenge of keeping the company running with this business critical application.

## Audience

This white paper is intended for Oracle database administrators, technical consultants, partners and members of the EMC and partner professional services community who are looking for faster, more efficient Oracle backup and recovery with complete control over Oracle disaster recovery procedures.

## Background

### Oracle Backup and Recovery

Oracle Recovery Manager® (RMAN) is a native utility from Oracle that is designed for online (hot) backup and recovery of Oracle database files. RMAN is built into Oracle and does not require separate licensing or installation. RMAN metadata is stored in the Oracle control file of the database being backed up and optionally in a recovery catalog database within Oracle. This cataloged metadata can be browsed during a restore. RMAN can be managed from Oracle Enterprise Manager or from the Oracle command line interface. Oracle RMAN is a pre-requisite to efficiently backup and recover Oracle databases.

Oracle RMAN Block Change Tracking (BCT) optimizes incremental backups by utilizing a tracking file to keep a record of blocks that have changed within each datafile since the last backup. The tracking file is read during an incremental backup to avoid having to read each data file individually to identify which blocks have changed.

EMC Data Domain Boost™ for RMAN is compatible with Oracle BCT. With DD Boost for RMAN configured and BCT enabled, and an incremental backup is run, the RMAN script will direct the backup to read the tracking file. Only unique data segments will be copied to the Data Domain system, without the need to scan the complete file system.

### EMC Data Domain Deduplication Systems

Data Domain deduplication storage systems offer a cost-effective alternative to tape that allows users to enjoy the retention and recovery benefits of inline deduplication, as well as network-efficient replication over the wide area network (WAN) for disaster recovery (DR).

Data Domain systems reduce the amount of disk storage needed to retain and protect data by 10 to 30 times. Data on disk is available online and onsite for longer retention periods, and restores become fast and reliable. Storing only unique data on disk also means that data can be cost-effectively replicated over existing networks to

remote sites for DR. With the industry's fastest deduplication storage controller, Data Domain systems allow more backups to complete faster while putting less pressure on limited backup windows.

All Data Domain systems are built as the data store of last resort, which is enabled by the EMC Data Domain Data Invulnerability Architecture – end-to-end data verification, continuous fault detection and self-healing, and other resiliency features transparent to the application.

Data Domain Boost extends the Data Domain Data Invulnerability Architecture by generating checksums on the Oracle server before RMAN sends data to the Data Domain system. The Data Domain system receiving the data computes new checksums on the incoming data and compares them to the computed values from the backup application, ensuring inline verification of data.

For more information on Data Domain Data Invulnerability Architecture, click here: <http://www.emc.com/collateral/software/white-papers/h7219-data-domain-data-invul-arch-wp.pdf>

## EMC Data Domain Boost – Advanced Integration with Oracle RMAN

EMC Data Domain Boost™ for Oracle Recovery Manager is an industry first to deliver complete control of Oracle backup and disaster recovery, enabling DBAs to have confidence in self-administered recovery from the local or the DR site. DD Boost for RMAN enables Oracle database administrators (DBAs) to manage their backup, operational recovery, and disaster recovery processes without dependence on a backup administrator.

In addition, DD Boost for RMAN accelerates backup performance (up to 50 percent faster than NFS), enables more backups with existing resources and simplifies administration.

Implementing DD Boost for RMAN simply requires installing the DD Boost plug-in on the Oracle server, and then the DD Boost plug-in interfaces between the Oracle Media Management Layer (MML) API (also known as the Simple Backup to Tape or “SBT” API) and DD Boost (see Figure 1 below). The Oracle MML API allows backup applications to interface with Oracle RMAN.




Figure 1: DD Boost for Oracle RMAN Solution Overview

### Distributing the Deduplication Process

Prior to DD Boost, Oracle Recovery Manager would send all data, unique or redundant, to a Data Domain system for deduplication processing, as shown in figure 2 below. With DD Boost, Oracle servers send only unique data segments to the Data Domain system.




Figure 2: RMAN Backup to Data Domain without DD Boost

DD Boost for RMAN significantly increases performance by distributing parts of the deduplication process to the DD Boost plug-in on the Oracle server, as shown in Figure 3.




Figure 3: RMAN Backup to Data Domain with DD Boost

With DD Boost, Oracle servers send only unique data segments to the Data Domain system. By sending only unique data, DD Boost decreases the load on the Oracle server since sending data is significantly more CPU and memory intensive than executing parts of the deduplication process.

In addition to performance improvements, the amount of data transferred over the LAN is greatly reduced. This optimized efficiency enables substantially more backups with existing Oracle database server(s) and Ethernet networks.

### RMAN Managed File Replication

EMC Data Domain Replicator software provides network-efficient, automated, and encrypted replication for disaster recovery, copying only unique and compressed data segments to a remote Data Domain system. With DD Boost for RMAN, DBAs gain full visibility and control of these disaster recovery processes via DD Boost managed file replication.

Configuring DD Boost managed file replication is simple and starts with setting up DD Replicator software using the Data Domain Enterprise Manager console and CLI. Then, to setup RMAN to manage Data Domain replication, simply define a backup.cmd file in RMAN to enable the remote Data Domain system as the target for replication. As each RMAN file write completes on the primary Data Domain system during backup, replication automatically begins to send only unique and compressed data segments to the remote Data Domain system. Remote backup copies are immediately visible to the Oracle server and available for recovery, providing faster "time-to-DR" readiness.

Figure 4 below, depicts the workflow that is transparently executed with managed file replication, from the Oracle server.




Figure 4: DD Boost - Managed File Replication Workflow

### DD Boost – Managed File Replication Steps

1. RMAN performs the backup through the DD Boost plug-in resulting in:
  - 50 percent faster performance
  - Up to 99 percent less LAN bandwidth requirements
  - Between 20 to 40 percent less CPU resource load on the Oracle server
2. When each RMAN file write completes on the local Data Domain system, network-efficient replication automatically and transparently begins to copy the unique and compressed file segments to the remote Data Domain system
3. When the RMAN backup to the local Data Domain system as well as replication to the remote Data Domain system are complete, the RMAN catalog is updated with information about both the local and remote backup copies

In addition, with the EMC Data Domain Encryption software option, deduplicated data can be encrypted inline before landing on the Data Domain system. Similarly, DD Replicator can encrypt and decrypt data in-flight when replicating between Data Domain systems. These encryption options can be enabled separately or concurrently to achieve different security goals.

### Advanced Load Balancing and Failover

DD Boost provides application-aware load balancing and failover of multiple 1GbE or 10GbE links. This feature provides automatic distribution of backup and restore jobs and dynamic load balancing on multiple ports on the Data Domain system.

The DD Boost plug-in negotiates with the Data Domain system for an interface and distributes the load based on the number of jobs scheduled on a particular interface. In addition, DD Boost provides automatic link failover by transparently moving jobs on failed links to healthy links, which speeds backups by increasing reliability and fault tolerance.

In addition, DD Boost reduces complexity by eliminating the need to hard mount a Data Domain system through NFS. Due to this simplification, DD Boost automatic link failover eliminates the risk of hard mount failures bringing down the entire database.

## Features and Benefits Summary

Feature	Benefit
Distributed Deduplication Process	<ul style="list-style-type: none"><li>• 50 percent faster performance versus NFS</li><li>• Up to 99 percent less LAN bandwidth required</li><li>• 20 to 40 percent CPU utilization improvement on Oracle server</li><li>• Improve RPO and RTO service level agreements</li></ul>
Oracle Managed Replication	<ul style="list-style-type: none"><li>• Oracle database administrator control of replication process</li><li>• Network-efficient replication copies only unique, compressed data segments to the remote Data Domain system for faster “time-to-DR” readiness</li><li>• Both local and remote backup copies are visible in the Oracle RMAN catalog</li><li>• Restores can be initiated using backup filesets from either the local or remote Data Domain system</li></ul>
Advanced Load Balancing and Link Failover	<ul style="list-style-type: none"><li>• Application-aware aggregation of multiple 1GbE or 10GbE links<ul style="list-style-type: none"><li>▪ Backup and restore jobs load are automatically distributed on multiple ports on the Data Domain system</li><li>▪ In-flight jobs on failed ports on a Data Domain system are transparently moved to healthy links</li><li>▪ Compatible with switch assisted aggregation</li></ul></li></ul>
Ease of Use	<ul style="list-style-type: none"><li>• DBA-controlled simplified management<ul style="list-style-type: none"><li>▪ No dependency on backup applications</li><li>▪ Eliminates mount point failures and uses standard scripting</li></ul></li></ul>
Rich Oracle Environment Support	<ul style="list-style-type: none"><li>• See the online support matrix for details at <a href="https://my.datadomain.com/US/en/search.jsp">https://my.datadomain.com/US/en/search.jsp</a></li></ul>

## DD Boost for RMAN Configurations and Capabilities

DD Boost for Oracle RMAN supports a number of commonly deployed configurations and capabilities:

1. Backup and recovery of Oracle 10g and 11g databases (Oracle RAC, Oracle database appliance and Exadata configurations)
  - Full backup and Incremental backup with Oracle Block Change Tracking
  - Frequent backup, faster restores
  - Short and long-term retention of database backups

2. Backup and recovery of multiple Oracle databases with different SID's within an instance
3. Backup and recovery of multiple distinct Oracle databases using a centralized RMAN backup server (typically seen in service provider deployments)
4. Backup replication for disaster recovery for Oracle databases
5. Refresh test and development instances from local or remote backup copies

## Oracle Backup and Recovery Flexibility

In addition to DD Boost for RMAN, EMC offers a variety of other Oracle backup solutions using a Data Domain system:

- DD Boost for EMC Avamar<sup>®</sup>, EMC NetWorker<sup>®</sup>, Symantec NetBackup<sup>®</sup> and Backup Exec<sup>®</sup>
- Direct backups via NFS, dNFS and CIFS

EMC backup and recovery solutions allow enterprises to choose the best approach for their specific environment and requirements.

### EMC NetWorker and EMC Avamar

Some corporate policies may dictate that the backup administrators must retain control of all backup and recovery processes. In this case, DD Boost for NetWorker or Avamar provides benefits that are similar to those listed above, but allow the backup administrators to retain the control.

#### EMC NetWorker

DD Boost for NetWorker distributes parts of the deduplication process to either a NetWorker storage node or directly to an Oracle server via the EMC NetWorker Module for Databases and Applications (NMDA).

With DD Boost, backup administrators can gain faster, more efficient Oracle backup and manage Data Domain replication for catalog consistent disaster recovery.

#### EMC Avamar

Avamar enables fast, efficient backup and recovery for Oracle data by reducing the size of backup data at the client before it is transferred across the network and stored. Avamar integrates with Data Domain via DD Boost software to allow users to enjoy the Data Domain system's scale and performance as well as the simplicity and network efficiency enabled by Avamar. DD boost enables Avamar clients to send specific data types, including Oracle, that are better suited to high-speed inline deduplication, directly to Data Domain systems.

## Centralizing Management of Backups

Giving database administrators control of the backup and recovery process does not mean that IT directors and backup administrators lose visibility into Oracle backup, recovery and disaster recovery operations. EMC Data Protection Advisor (DPA) provides a single, comprehensive view of the entire backup infrastructure through automated data collection, analysis, alerting and reporting. DPA integrates directly with Data Domain system and the RMAN control file or catalog database to report on all RMAN backup and recovery operations.

Centralized management and reporting provide the operational control and transparency needed for Oracle backup and recovery and assures backup and recovery operations, predictable data recovery.

### EMC Data Protection Advisor for Oracle RMAN and Data Domain

DPA gathers detailed information about Oracle backup jobs with Oracle RMAN. When an RMAN backup runs, it writes detailed information about the backup to either the control file or to a designated recovery catalog within the environment. Then the DPA Oracle RMAN module gathers information directly from the database control file or from the centralized recovery catalogue database.




Figure 5: EMC Data Protection Advisor – Data Protection Management Software

In addition, DPA provides centralized rollup reporting of multiple Data Domain systems, as well as historical and projected consumption and performance reporting. DPA compiles and reports specific information about the Data Domain environment, including hardware status, throughput, deduplication ratio, as well as disk utilization and performance. DPA tracks utilization over time and generates predictive trend alerts such as disk space consumption. In addition, DPA monitors and provides status reports on environment components, including network interfaces, CPU condition, disk health, etc. When thresholds are exceeded, DPA sends an alert to the administrator. DPA enables centralized reporting while maintaining DBA control of Oracle backup, recovery, and disaster recovery processes.

## Conclusion

EMC is first to deliver complete control of Oracle backup and disaster recovery to Oracle database administrators. The advanced integration between DD Boost and RMAN provides a faster, more efficient database backup solution.

For years, Data Domain deduplication storage systems have empowered Oracle database administrators to manage their own backup and recovery processes without the assistance of backup administrators.

Data Domain Boost has enhanced this capability with advanced integration with Oracle Recovery Manager. By distributing parts of the deduplication process to the Oracle database server, DBAs gain faster, more efficient backup and recovery. In addition, DD Boost provides the Oracle DBA with complete control of disaster recovery operations with full RMAN catalog consistency. Finally, advanced load balancing and failover increases fault tolerance and simplifies management.

Data Domain Boost for Oracle Recovery Manager gives DBAs confidence and control; keeping the business running, no matter what.

EMC provides flexibility with DD Boost, giving you the option to backup Oracle with leading backup applications like NetWorker and Avamar or from RMAN directly via NFS or CIFS. Finally, Data Protection Advisor provides a single pane for complete visibility, control, and reporting to ensure SLAs are met.

## Appendix

### Best Practice – Optimizing Data Deduplication

There are certain parameters that are configured within RMAN that can adversely affect data deduplication performance results and should be set according to the information listed in table 2 below.

**Table 2: Data Domain – Best Practice Settings**

Option	Setting	Explanation / Benefit
Compression	Disabled	<ul style="list-style-type: none"><li>All supported versions of RMAN can apply a binary compression algorithm (BZIP2) to the backup set. When using a Data Domain system as the backup target, RMAN lossless compression should not be used, as a pre-compression of the backup streams will randomize the data patterns and defeat deduplication.</li></ul>
Encryption	Disabled	<ul style="list-style-type: none"><li>Encrypted files are by definition, unique. The encryption software that is part of the RMAN will create unique files, on-the-fly for each backup. When using a Data Domain system as the backup target, RMAN encryption should not be used, as completely unique data defeats deduplication benefits.</li><li>DD Encryption software is available for Data Domain systems to provide for encryption of data-at-rest and data-in-flight</li></ul>
Multiplexing / FILESPERSET	1	<ul style="list-style-type: none"><li>FILESPERSET is the option Oracle RMAN uses to control multiplexing of the datafiles into the backup set, controlling how many datafiles are written to a particular file within the backup set. When using a Data Domain system as the backup target, this should be is set to 1, or data deduplication will be adversely affected.</li></ul>

## Additional Resources

EMC Data Domain Boost for Oracle Recovery Manager (RMAN) Solutions Overview:  
<http://www.emc.com/collateral/software/solution-overview/h9687-data-domain-boost-oracle-rman-so.pdf>

EMC Data Domain Deduplication Storage Systems: <http://www.emc.com/backup-and-recovery/data-domain/data-domain.htm>

EMC Data Domain Boost Data Sheet: <http://www.emc.com/collateral/software/data-sheet/h7034-datadomain-boost-sw-ds.pdf>

EMC Data Domain Data Invulnerability Architecture: Enhancing Data Integrity and Recoverability: <http://www.emc.com/collateral/software/white-papers/h7219-data-domain-data-invul-arch-wp.pdf>

EMC NetWorker Application Modules – Data Sheet:  
<http://www.emc.com/collateral/software/data-sheet/h2479-nw-app-ds.pdf>

IDC Study – Worldwide Purpose Built Backup Appliances:  
<http://www.emc.com/collateral/analyst-reports/idc-worldwide-purpose-built-backup-appliance-2011-2015.pdf>