

# Dell EMC Oracle RMAN Agent

Version 4.5

## Administration Guide

302-004-149

REV 04

Copyright © 2012-2018 Dell Inc. or its subsidiaries. All rights reserved.

Published September 2018

Dell 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." DELL 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 DELL SOFTWARE DESCRIBED IN THIS PUBLICATION REQUIRES AN APPLICABLE SOFTWARE LICENSE.

Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners.  
Published in the USA.

Dell EMC  
Hopkinton, Massachusetts 01748-9103  
1-508-435-1000 In North America 1-866-464-7381  
[www.DellEMC.com](http://www.DellEMC.com)

# CONTENTS

<b>Figures</b>		<b>7</b>
<b>Tables</b>		<b>9</b>
<b>Preface</b>		<b>11</b>
<b>Chapter 1</b>	<b>Introducing the Oracle RMAN Agent</b>	<b>15</b>
	Overview of Oracle RMAN agent.....	16
	Supported configurations.....	17
<b>Chapter 2</b>	<b>DD Boost Features</b>	<b>19</b>
	Distributed segment processing.....	20
	Advanced load balancing and link failover.....	22
	Configuration restrictions.....	24
	DD Boost-over-Fibre Channel transport.....	24
	DD Boost-over-Fibre Channel path management.....	26
	Initial path selection.....	27
	Client path failover.....	27
	Dynamic re-balancing.....	28
	Queue-depth constraints.....	28
	Encrypted managed file replication.....	29
	DD Boost and high availability.....	30
	DD Boost, HA, and failover.....	30
	Partial HA configurations .....	30
	Low-bandwidth optimization.....	31
	Client access validation.....	31
	In-flight encryption.....	31
	Marker support .....	31
<b>Chapter 3</b>	<b>Configuring the Data Domain System</b>	<b>33</b>
	Oracle RMAN agent and DD OS upgrades.....	34
	Firewalls and ports.....	34
	Enabling DD Boost on a Data Domain system.....	34
	Enabling Oracle optimized deduplication.....	35
	Configuring the DD Boost server.....	36
	Creating storage units.....	36
	Configuring distributed segment processing.....	36
	Configuring advanced load balancing and link failover.....	37
	Enabling low-bandwidth optimization .....	38
	Enabling encrypted file replication.....	39
	Configuring client access validation.....	39
	Configuring DD Boost-over-FC service.....	41
	Sizing DD Boost-over-FC device-set.....	43
	Sizing calculation.....	44
	Configuring the SCSI device driver for AIX clients.....	46

	Configuring the SCSI generic device driver for Solaris clients.....	46
<b>Chapter 4</b>	<b>Configuring the Oracle RMAN Agent</b>	<b>49</b>
	Installing and configuring the Oracle RMAN agent on AIX, HP-UX, Linux, and Solaris systems.....	50
	Installing the Oracle RMAN agent on database servers.....	50
	Installing the Oracle RMAN agent on Linux and UNIX.....	51
	Uninstalling the Oracle RMAN agent on database servers.....	54
	Uninstalling the Oracle RMAN agent on Linux and UNIX.....	56
	Configuring the software on Linux and UNIX systems.....	56
	Installing and configuring the Oracle RMAN agent on Windows systems....	57
	Installing the Oracle RMAN agent on database servers.....	57
	Installing the Oracle RMAN agent on Windows systems.....	57
	Uninstalling the Oracle RMAN agent on Windows systems.....	60
	Configuring the software on Windows systems.....	61
	Registering each Data Domain system.....	61
	Performing backups of a clustered environment to Data Domain.....	65
	Data Domain replication.....	65
	Performing managed file replication (MFR).....	66
	Restoring from MTree replicated backups.....	69
	Specifying subdirectories for Oracle RMAN agent backups.....	70
	Performing backups and restores of Oracle CDBs and PDBs.....	73
	Performing DD Boost compressed restores.....	74
	Display and deletion of backups.....	74
	Using the ddutil command to display and delete backups.....	75
	Using the ddutil command to display backup information.....	78
	Using the ddutil command to delete backups.....	79
	Using the ddutil command to prune backups based on age.....	80
	Configuring the use of Data Domain Cloud Tier for data movement to the cloud.....	81
	Setting up the data movement to the cloud.....	81
	Performing the data recall from the cloud.....	82
	RMAN crosscheck and deletion of backup data in the cloud.....	83
	Usage limits on Data Domain resources.....	83
	Oracle RMAN agent operations with Data Domain usage limits on capacity.....	84
	Oracle RMAN agent operations with Data Domain usage limits on streams.....	84
	Configuring usage limits on Data Domain resources.....	85
	Estimating the Data Domain resource usage on Oracle systems....	86
<b>Chapter 5</b>	<b>Data Domain System Administration</b>	<b>89</b>
	Restricting DD Boost access to specific database servers.....	90
	File replication version compatibility.....	91
	Modifying an interface group.....	91
	Removing the advanced load balancing and link failover configuration.....	92
<b>Chapter 6</b>	<b>Troubleshooting</b>	<b>93</b>
	Troubleshooting overview.....	94
	Investigating problems.....	94
	Oracle limitations.....	94
	Reporting problems.....	94
	Data Domain system settings for file replication.....	95

	Lockbox troubleshooting information.....	95
	Failure to load the media management library.....	95
<b>Chapter 7</b>	<b>Using DD Boost Commands</b>	<b>97</b>
	DD Boost command summary.....	98
	The ddboost command options.....	98
	access.....	98
	clients.....	99
	destroy.....	99
	disable.....	99
	enable.....	100
	fc.....	100
	file-replication.....	100
	ifgroup.....	104
	option.....	106
	reset.....	107
	set user-name.....	108
	show.....	108
	status.....	109
	storage-unit.....	110
<b>Appendix A</b>	<b>Installation and Configuration Checklist</b>	<b>113</b>
	Configure the Data Domain system.....	114
	Download and install the Oracle RMAN agent.....	114
	Configure RMAN for the environment.....	114
	Test the installation by using sbttest (optional) .....	115
	Backup examples.....	115
<b>Appendix B</b>	<b>Performing a Backup with Oracle Enterprise Manager</b>	<b>117</b>
	Example screens.....	118
<b>Appendix C</b>	<b>Installing and Configuring the Oracle RMAN Agent</b>	<b>121</b>
	Installation and configuration examples.....	122

## CONTENTS

# FIGURES

1	Oracle RMAN agent configuration.....	16
2	High-level data flow with distributed segment processing disabled.....	21
3	High-level data flow with distributed segment processing enabled.....	22
4	Example of advanced load balancing and link failover configuration.....	23
5	DD Boost-over-IP transport.....	24
6	SCSI commands between media server and Data Domain system.....	25
7	Fibre Channel communication path.....	26
8	DD Boost-over-FC path management scenario.....	26
9	Login window.....	118
10	Database instance page.....	118
11	Processing page.....	118
12	Schedule Customized Backup: Options page.....	119
13	Schedule Customized Backup: Schedule page.....	119
14	Schedule Customized Backup: Review page.....	119
15	Execution page.....	120
16	Oracle RMAN agent package and its contents.....	122
17	Oracle RMAN agent installation.....	123
18	Oracle RMAN agent installed files.....	124
19	Registration of the Data Domain user.....	125
20	Oracle RMAN agent lockbox files.....	125
21	Oracle RMAN agent channel configuration and backup.....	126

## FIGURES

# TABLES

1	Revision history.....	11
2	Style conventions.....	13
3	Environment variables for installation or upgrade of Oracle RMAN agent.....	51
4	Software package names on UNIX and Linux.....	52
5	Product software components on Linux or UNIX.....	53
6	Environment variables for uninstallation of Oracle RMAN agent.....	55
7	Product software components on Windows.....	60
8	The ddutil command options for backup display and deletion.....	76
9	Summary of DD Boost commands.....	98
10	Inbound and outbound file transfer.....	102
11	Inbound and outbound data transfer.....	103
12	Output definitions for DD Boost histogram.....	108



# Preface

As part of an effort to improve product lines, periodic revisions of software and hardware are released. Therefore, all versions of the software or hardware currently in use might not support some functions that are described in this document. The product release notes provide the most up-to-date information on product features.

If a product does not function correctly or does not function as described in this document, contact a technical support professional.

---

## Note

This document was accurate at publication time. To ensure that you are using the latest version of this document, go to the Support website at <https://support.emc.com>.

---

## Purpose

This document describes how to install, configure, and use the Oracle RMAN agent version 4.5 software.

In earlier versions of the product, the Oracle RMAN agent was known as Data Domain Boost (DD Boost) for Oracle Recovery Manager (RMAN) or the RMAN plug-in. In version 4.0, the product name changed to Oracle RMAN agent.

## Audience

This guide is for system administrators who are familiar with Oracle backup applications and general backup administration.

## Revision history

The following table presents the revision history of this document.

**Table 1** Revision history

Revision	Date	Description
04	September 21, 2018	Added the following topic in Chapter 3, "Configuring the Data Domain System": <ul style="list-style-type: none"><li>Enabling Oracle optimized deduplication—Added this topic that recommends you enable the Oracle optimized deduplication at the MTree level.</li></ul>
03	June 20, 2018	Updated the following topic in Chapter 4, "Configuring the Oracle RMAN Agent": <ul style="list-style-type: none"><li>Restoring from multiple backup copies—Updated the paragraph that provides an example with "three copies of a backup".</li></ul>
02	November 14, 2017	Updated the following topics in Chapter 4, "Configuring the Oracle RMAN Agent": <ul style="list-style-type: none"><li>Data Domain replication—Added this topic about the support of managed file replication and MTree replication.</li></ul>

**Table 1** Revision history (continued)

Revision	Date	Description
		<ul style="list-style-type: none"> <li>Performing managed file replication (MFR)—Updated the heading and first two paragraphs and repositioned this topic under "Data Domain replication".</li> <li>Restoring from multiple backup copies—Updated this topic and repositioned it under "Performing managed file replication (MFR)".</li> <li>Restoring from MTree replicated backups—Added this topic about how to enable and perform restores of the MTree replicated backups.</li> <li>Performing backups and restores of Oracle CDBs and PDBs—Added this topic about the support of operations with Oracle container databases (CDBs) and pluggable databases (PDBs).</li> </ul>
01	September 25, 2017	Initial release of this document for Oracle RMAN agent 4.5.

**Related documentation**

The following publications provide additional information:

- *Data Domain Operating System Release Notes*
- *Data Domain Operating System Initial Configuration Guide*
- *Data Domain Operating System Administration Guide*
- *Data Domain Operating System Command Reference Guide*
- *Data Domain Expansion Shelf Hardware Guide*
- The Data Domain system installation and setup guide for each of the supported platforms (for example DD890, DD670, and so on).

**RMAN Documentation**

11g Release 2: From the Oracle Database Documentation Library at [www.oracle.com/pls/db112/homepage](http://www.oracle.com/pls/db112/homepage), navigate to the Database Administration folder and select Backup and Recovery.

- *Oracle Database Backup and Recovery User's Guide*
- *Oracle Database Backup and Recovery Reference*

10g Release 2: From the Oracle Database Library at [www.oracle.com/pls/db102/homepage](http://www.oracle.com/pls/db102/homepage), select the Administration tab and locate the Backup and Recovery heading.

- *Oracle Database Backup and Recovery Reference*
- *Oracle Database Backup and Recovery Advanced User's Guide*

**Special notice conventions that are used in this document**

The following conventions are used for special notices:

**NOTICE**

Identifies content that warns of potential business or data loss.

---

**Note**

Contains information that is incidental, but not essential, to the topic.

---

**Typographical conventions**

The following type style conventions are used in this document:

**Table 2** Style conventions

<b>Bold</b>	Used for interface elements that a user specifically selects or clicks, for example, names of buttons, fields, tab names, and menu paths. Also used for the name of a dialog box, page, pane, screen area with title, table label, and window.
<i>Italic</i>	Used for full titles of publications that are referenced in text.
Monospace	Used for: <ul style="list-style-type: none"> <li>• System code</li> <li>• System output, such as an error message or script</li> <li>• Pathnames, file names, file name extensions, prompts, and syntax</li> <li>• Commands and options</li> </ul>
<i>Monospace italic</i>	Used for variables.
<b>Monospace bold</b>	Used for user input.
[ ]	Square brackets enclose optional values.
	Vertical line indicates alternate selections. The vertical line means or for the alternate selections.
{ }	Braces enclose content that the user must specify, such as x, y, or z.
...	Ellipses indicate non-essential information that is omitted from the example.

---

You can use the following resources to find more information about this product, obtain support, and provide feedback.

**Where to find product documentation**

- <https://support.emc.com>
- <https://community.emc.com>

**Where to get support**

The Support website at <https://support.emc.com> provides access to licensing information, product documentation, advisories, and downloads, as well as how-to and troubleshooting information. This information may enable you to resolve a product issue before you contact Support.

To access a product specific Support page:

1. Go to <https://support.emc.com/products>.
2. In the **Find a Product by Name** box, type a product name, and then select the product from the list that appears.
3. Click .
4. (Optional) To add the product to **My Saved Products**, in the product specific page, click **Add to My Saved Products**.

### Knowledgebase

The Knowledgebase contains applicable solutions that you can search for by solution number, for example, 123456, or by keyword.

To search the Knowledgebase:

1. Go to <https://support.emc.com>.
2. Click **Advanced Search**.  
The screen refreshes and filter options appear.
3. In the **Search Support or Find Service Request by Number** box, type a solution number or keywords.
4. (Optional) To limit the search to specific products, type a product name in the **Scope by product** box, and then select the product from the list that appears.
5. In the **Scope by resource** list box, select **Knowledgebase**.  
The **Knowledgebase Advanced Search** panel appears.
6. (Optional) Specify other filters or advanced options.
7. Click .

### Live chat

To participate in a live interactive chat with a support agent:

1. Go to <https://support.emc.com>.
2. Click **Chat with Support**.

### Service requests

To obtain in-depth help from Licensing, submit a service request. To submit a service request:

1. Go to <https://support.emc.com>.
2. Click **Create a Service Request**.

---

### Note

To create a service request, you must have a valid support agreement. Contact a sales representative for details about obtaining a valid support agreement or with questions about an account. If you know the service request number, then directly enter the service request number in the `Service Request` field to get the valid details.

---

To review an open service request:

1. Go to <https://support.emc.com>.
2. Click **Manage service requests**.

### Online communities

Go to the Community Network at <https://community.emc.com> for peer contacts, conversations, and content on product support and solutions. Interactively engage online with customers, partners, and certified professionals for all products.

### How to provide feedback

Feedback helps to improve the accuracy, organization, and overall quality of publications. You can send feedback to [DPAD.Doc.Feedback@emc.com](mailto:DPAD.Doc.Feedback@emc.com).

# CHAPTER 1

## Introducing the Oracle RMAN Agent

This chapter contains the following topics:

- [Overview of Oracle RMAN agent](#)..... 16
- [Supported configurations](#)..... 17

## Overview of Oracle RMAN agent

The Oracle RMAN agent enables database servers to communicate with Data Domain systems in an optimized way, without the need to use a backup application.

**Note**

In earlier versions of the product, the Oracle RMAN agent was known as Data Domain Boost (DD Boost) for Oracle Recovery Manager (RMAN) or the RMAN plug-in. In version 4.0, the product name changed to Oracle RMAN agent.

The use of the Oracle RMAN agent improves performance while reducing the amount of data that is transferred over the LAN. In the context of Oracle RMAN, the software has two components:

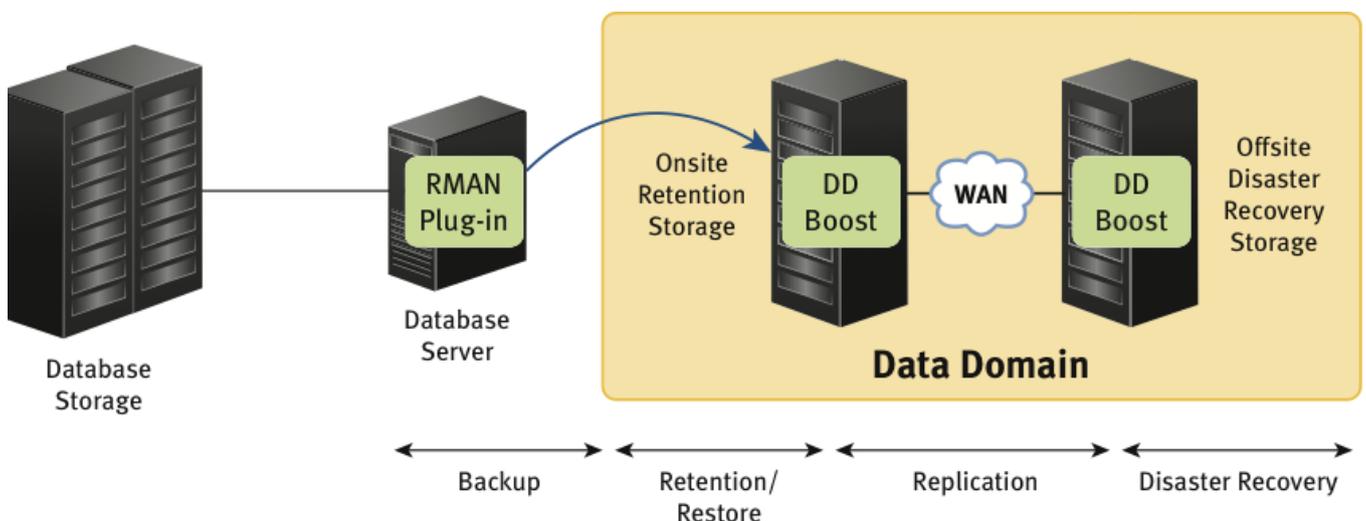
- An RMAN agent that you install on each database server. This RMAN agent includes the DD Boost libraries for communicating with the DD Boost server that runs on the Data Domain system.
- The DD Boost server that runs on the Data Domain systems.

RMAN sets policies that control when backups and replications occur. Administrators manage backup, replication, and restore from a single console and can use all the features of DD Boost, including WAN-efficient replicator software. RMAN manages all the files (collections of data) in the catalog, even those created by the Data Domain system.

The Data Domain system exposes pre-made disk volumes called storage units to a DD Boost enabled database server. Multiple database servers, each with the Oracle RMAN agent, can use the same storage unit on a Data Domain system as a storage server. Each database server can run a different operating system, provided that it is supported by Data Domain.

The following figure shows an example configuration of the Oracle RMAN agent.

**Figure 1** Oracle RMAN agent configuration



## Supported configurations

Data Domain supports DD Boost on the following Data Domain systems:

- Data Domain systems DD1xx, DD5xx, DD6xx, DD8xx, DD9xx
- Data Domain Extended Retention (formerly Archiver) systems

The Oracle RMAN agent version must be compatible with the Data Domain system and RMAN configurations. Data Domain does not support combinations other than those detailed in the online compatibility guide, which is available at <http://compatibilityguide.emc.com:8080/CompGuideApp/>.



# CHAPTER 2

## DD Boost Features

This chapter contains the following topics:

- [Distributed segment processing](#) ..... 20
- [Advanced load balancing and link failover](#) ..... 22
- [DD Boost-over-Fibre Channel transport](#) ..... 24
- [DD Boost-over-Fibre Channel path management](#) ..... 26
- [Encrypted managed file replication](#) ..... 29
- [DD Boost and high availability](#) ..... 30
- [Low-bandwidth optimization](#) ..... 31
- [Client access validation](#) ..... 31
- [In-flight encryption](#) ..... 31
- [Marker support](#) ..... 31

## Distributed segment processing

Distributed segment processing enables parts of the deduplication process to be performed by the Oracle RMAN agent, which prevents the sending of duplicate data to the Data Domain system that is configured as a storage server. The option to distribute the processing is enabled or disabled on the Data Domain system. The Oracle RMAN agent negotiates with the Data Domain system for the current setting of the option and behaves appropriately. The Oracle RMAN agent offers two modes of operation for sending backup data to a Data Domain system, one with distributed segment processing enabled and the other with distributed segment processing disabled.

Distributed segment processing provides the following benefits:

- Higher throughput is possible because the Oracle RMAN agent sends only unique data to a Data Domain system, instead of sending all the data over the LAN. Throughput improvements depend on the redundant nature of the data being backed up, the overall workload on the database server, and the capability of the database server. In general, greater throughput is attained with higher redundancy, greater database server workload, and database server capability.
- The network bandwidth requirements are significantly reduced because only the unique data is sent over the LAN to the Data Domain systems.
- Recovery from failed backups can potentially proceed much faster. If a large backup fails in the middle or towards the end, when it is restarted by RMAN, the data already sent to the Data Domain system does not have to be resent. This results in faster completion of the backup on retry.

The *Data Domain Operating System Administration Guide* provides more information about local compression and its configuration.

You manage distributed segment processing through the `ddboost option` commands. Data Domain recommends that you use distributed segment processing if your network connection is 1 Gb Ethernet. [Configuring distributed segment processing](#) on page 36 provides details.

### Distributed segment processing disabled

In this mode, the Oracle RMAN agent sends the data directly to the Data Domain system over the LAN. The Data Domain system then segments, deduplicates, and compresses the data before it is written to the disk. The Oracle RMAN agent runs on the database server.

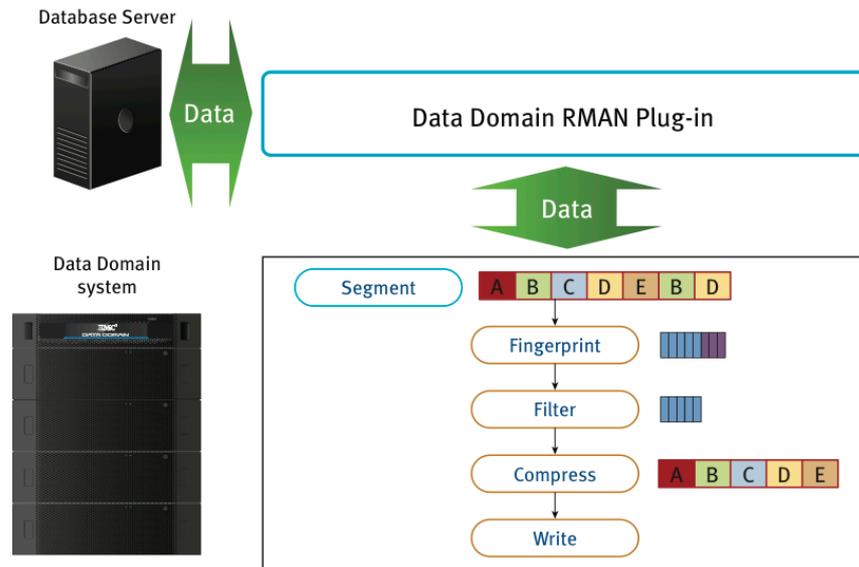
---

#### Note

Distributed segment processing cannot be disabled on an Extended Retention Data Domain system.

---

The following figure shows a high-level data flow scenario for a single-system Data Domain system with distributed segment processing disabled.

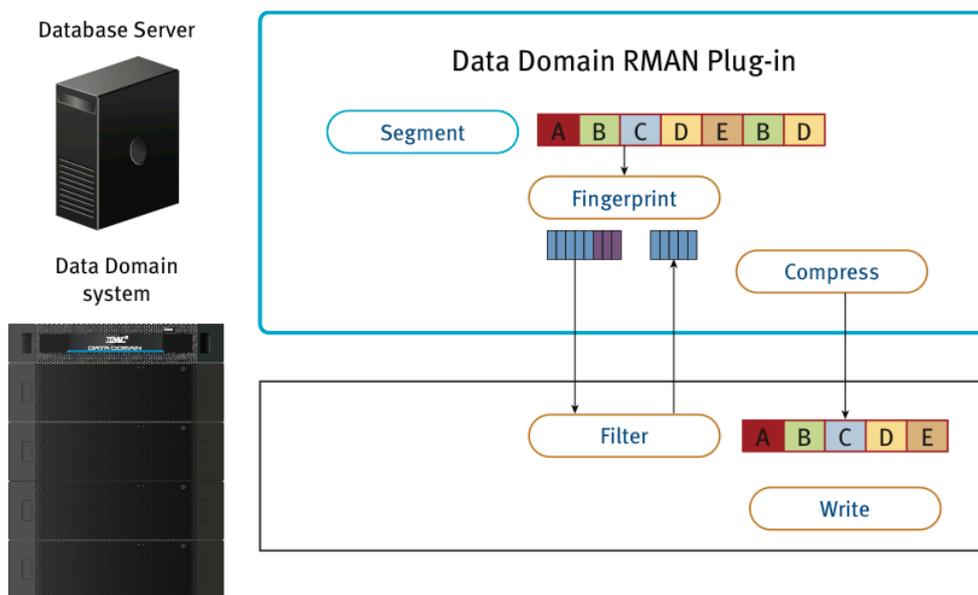
**Figure 2** High-level data flow with distributed segment processing disabled**Distributed segment processing enabled**

In this mode, the deduplication process is distributed between the Oracle RMAN agent and the Data Domain system. The Oracle RMAN agent runs on the database server. Parts of the deduplication process are run on the Oracle RMAN agent so that only unique data is sent to a Data Domain system over the LAN.

The Oracle RMAN agent segments the data, computes IDs for the segments, checks with the Data Domain system for duplicate segments, compresses unique segments that are not found on the Data Domain system, and sends the compressed data to the Data Domain system. The Data Domain system then writes the unique data to disk.

The Oracle RMAN agent's handling of data is transparent to RMAN, and the agent does not store any data on the database server.

The following figure shows a high-level data flow scenario for a single-system Data Domain system with distributed segment processing enabled.

**Figure 3** High-level data flow with distributed segment processing enabled

## Advanced load balancing and link failover

The advanced load balancing and link failover feature enables the combination of multiple Ethernet links into a group. Only one of the interfaces on the Data Domain system is registered with RMAN. The Oracle RMAN agent negotiates with the Data Domain system on the interface registered with RMAN to obtain an interface to send the data. The load balancing provides higher physical throughput to the Data Domain system compared to configuring the interfaces into a virtual interface using Ethernet level aggregation.

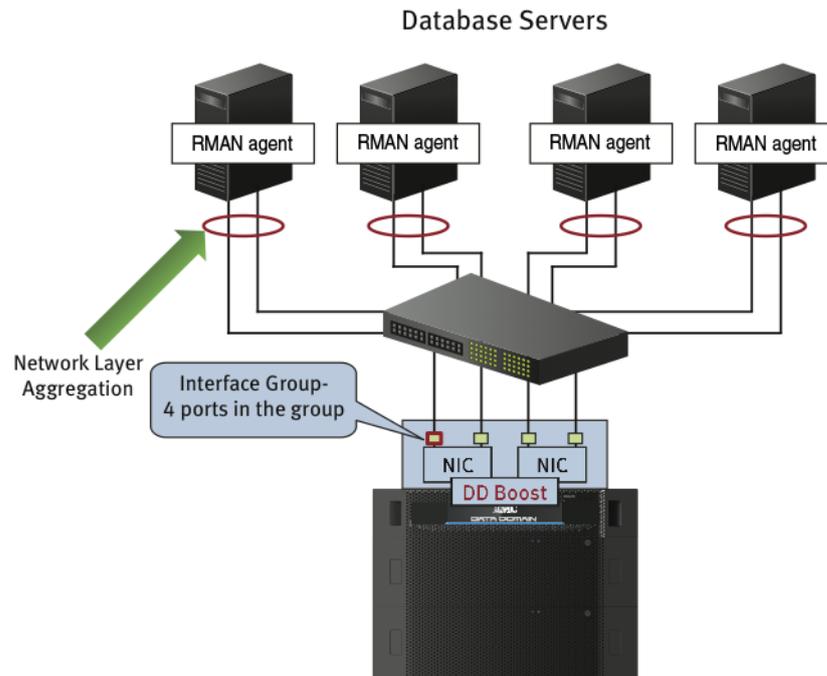
The Data Domain system load balances the connections coming in from multiple database servers on all the interfaces in the group. The load balancing is transparent to RMAN. Because advanced load balancing and link failover works at the software layer of the Oracle RMAN agent, it is seamless to the underlying network connectivity and supports both physical and virtual interfaces.

The data transfer is load balanced based on the number of connections outstanding on the interfaces. Only connections for backup and restore jobs are load balanced.

### Note

- The file replication connection between the Data Domain systems is not part of the load balancing. A single IP address is used for the target Data Domain system.
- It is recommended that you exclude one interface from the interface group (ifgroup) and reserve it for the file replication path between the source and target Data Domain systems. Every Oracle RMAN agent must be able to connect to every interface that is a member of the interface group on the Data Domain system.

The following figure shows an example high-level configuration diagram. The interface marked in red is registered with RMAN.

**Figure 4** Example of advanced load balancing and link failover configuration

The advanced load balancing and link failover feature can be used with other network layer aggregation and failover technologies. The sample configuration in the preceding figure shows one such possibility. The links connecting the database servers and the switch that connects to Data Domain system are put in an aggregated failover mode. This configuration provides end-to-end network failover functionality. Any of the available aggregation technologies can be used between the database server and the switch.

The advanced load balancing and link failover functionality also works with other network layer functionality on the Data Domain systems, including VLAN tagging and IP aliasing. This functionality enables additional flexibility in segregating traffic into multiple virtual networks, all of which run over the same physical links on the Data Domain system.

---

#### Note

The *Data Domain Operating System Administration Guide* provides more information about how to configure VLAN tagging and IP aliasing on a Data Domain system.

---

Advanced load balancing and link failover provides the following benefits:

- Eliminates the need to register multiple storage servers (one for each interface) with RMAN, which can potentially simplify installation management.
- If one of the interfaces in the group goes down while the Data Domain system is still operational, the subsequent incoming backup jobs are routed to the available interfaces.
- The backup and restore jobs are automatically load balanced on multiple interfaces in the group, which can potentially result in higher utilization of the links.
- All in-flight jobs to the failed interface are failed over transparently to healthy operational links. From the point of view of RMAN, the jobs continue uninterrupted.

[Configuring advanced load balancing and link failover](#) on page 37 provides more details.

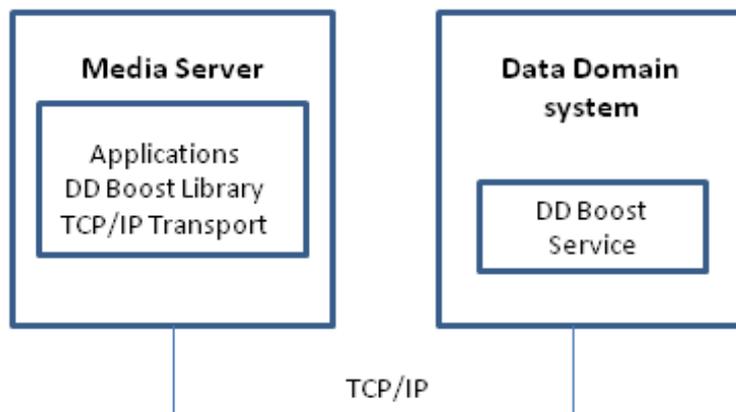
## Configuration restrictions

- Interfaces can be added only to the group by using an IP address.
- Although the advanced load balancing and link failover feature works with mixed 1 GbE interfaces and 10 GbE interfaces in a group, this is not a recommended setup. It is recommended to have interfaces with the same link speed in a group.
- Prior to DD OS 5.2, only one interface group was supported on a Data Domain system. This meant that, to connect more than one database server, a switch was needed in the middle. DD OS 5.2 and later support multiple interface groups, so this restriction no longer applies.

## DD Boost-over-Fibre Channel transport

Prior to DD OS release 5.3, all communication between the DD Boost library and any Data Domain system was performed by using IP networking. The application specified the Data Domain system by using its hostname or IP address, and the DD Boost Library used TCP/IP connections to transmit requests to, and receive responses from, the Data Domain system, as shown in the following figure.

**Figure 5** DD Boost-over-IP transport



DD OS release 5.3 introduced Fibre Channel as an alternative transport mechanism for communication between the DD Boost library and the Data Domain system.

Certain installations prefer or require the use of Fibre Channel as the media for data transfer between media servers and storage systems. The DD Boost-over-Fibre Channel transport (DD Boost-over-FC) enables such installations to access the features provided by the DD Boost technology.

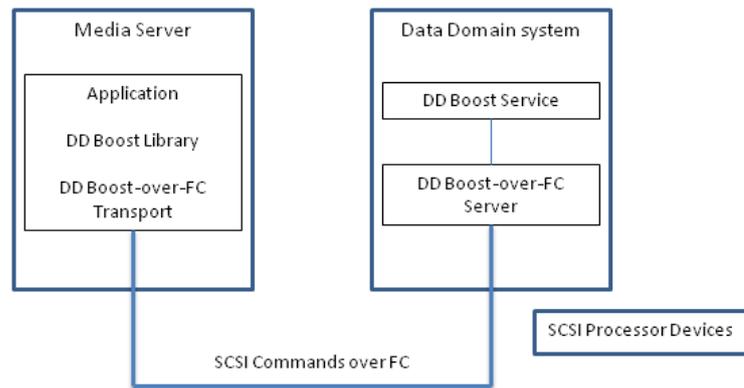
Although Fibre Channel is specified as a general-purpose data transport mechanism, in practice Fibre Channel is used solely as a transport for SCSI device access. In standard host operating systems, such as Windows and Linux, Fibre Channel hardware and drivers reside solely within the SCSI protocol stacks. Therefore, the DD Boost-over-FC transport must use SCSI commands for all communication.

In the DD Boost-over-FC solution, the Data Domain system advertises one or more SCSI devices of type Processor. The media server operating system discovers these devices, and makes them available to applications through a generic SCSI mechanism (Linux: SCSI Generic driver; Windows: SCSI Pass-Through Interface).

To request access to a Data Domain system by using the DD Boost-over-FC transport, the application specifies the Data Domain system by using the special string DFC-*<dfc-server-name>*, where *dfc-server-name* is the DD Boost-over-FC server name configured for the Data Domain system. The DD Boost-over-FC transport logic within the DD Boost library examines the set of generic SCSI devices available on the media server and uses SCSI commands to identify a catalog of devices, which are paths to the specified Data Domain system.

Referencing this catalog, the DD Boost-over-FC transport logic issues SCSI commands to the identified generic SCSI devices, to transfer DD Boost protocol requests and responses between the library and the Data Domain system.

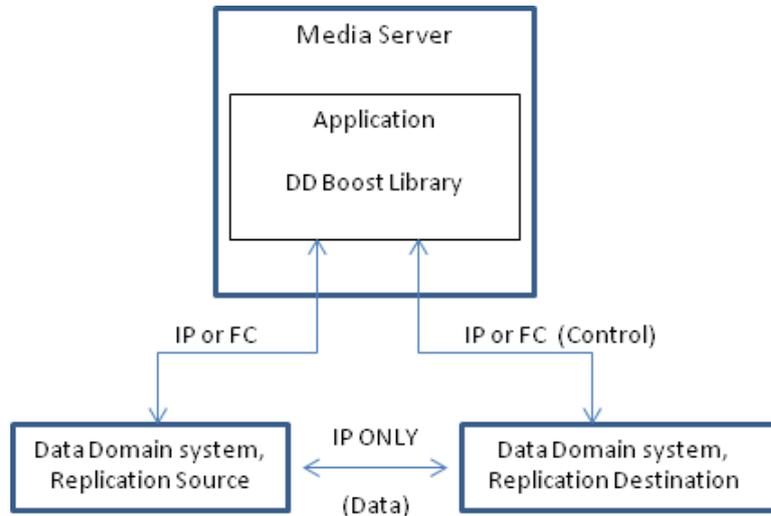
**Figure 6** SCSI commands between media server and Data Domain system



Most DD Boost features are independent of the transport. One notable exception is the DD Boost-over-IP advanced load balancing and link failover feature, and its associated ifgroups. This feature is specific to the IP transport. For the DD Boost-over-FC transport, load balancing and link-level high availability is achieved through different means.

The DD Boost-over-FC communication path applies only between the media server/DD Boost library and the Data Domain system, and does not apply to communication between two Data Domain systems. For example, the managed file replication section of this document describes efficient replication of data from one Data Domain system to another, possibly over a WAN. As shown in the following figure, such communication is always over a TCP/IP network, regardless of the communication path between the media server and the Data Domain systems.

**Figure 7** Fibre Channel communication path



## DD Boost-over-Fibre Channel path management

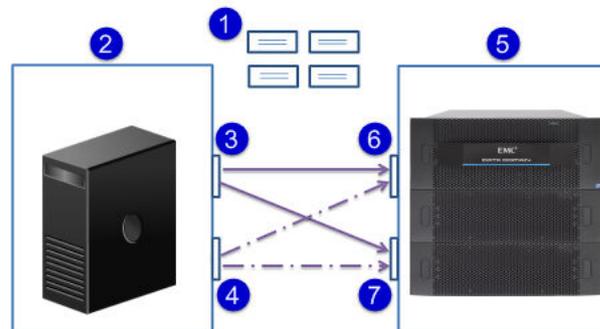
The Dynamic Interface Group (DIG)-based mechanism is based on Ethernet interfaces and is not applicable to the Fibre Channel transport. Instead, a different path mechanism is provided for the DD Boost-over-FC solution.

The Data Domain system advertises one or more SCSI processor-type devices to the backup server, over one or more physical paths. The operating system discovers all devices through all available paths, and creates a generic SCSI device for each discovered device and path.

For example, consider the following case:

- Backup server has 2 initiator HBA ports (A and B)
- Data Domain System has 2 FC target endpoints (C and D)
- Fibre Channel Fabric zoning is configured such that both initiator HBA ports can access both FC target endpoints
- Data Domain system is configured with a SCSI target access group that contains:
  - Both FC target endpoints on the Data Domain System
  - Both initiator HBA ports
  - 4 devices (0, 1, 2, and 3)

**Figure 8** DD Boost-over-FC path management scenario



1. Four devices
2. Backup Server
3. HBA Initiator A
4. HBA Initiator B
5. Data Domain System
6. Fibre Channel Endpoint C
7. Fibre Channel Endpoint D

In this case, the backup server operating system may discover up to 16 generic SCSI devices, one for each combination of initiator, target endpoint, and device number:

- /dev/sg11: (A, C, 0)
- /dev/sg12: (A, C, 1)
- /dev/sg13: (A, C, 2)
- /dev/sg14: (A, C, 3)
- /dev/sg15: (A, D, 0)
- /dev/sg16: (A, D, 1)
- /dev/sg17: (A, D, 2)
- /dev/sg18: (A, D, 3)
- /dev/sg19: (B, C, 0)
- /dev/sg20: (B, C, 1)
- /dev/sg21: (B, C, 2)
- /dev/sg22: (B, C, 3)
- /dev/sg23: (B, D, 0)
- /dev/sg24: (B, D, 1)
- /dev/sg25: (B, D, 2)
- /dev/sg26: (B, D, 3)

When the application requests that the DD Boost library establish a connection to the server, the DD Boost-over-FC transport logic within the DD Boost library uses SCSI requests to build a catalog of these 16 generic SCSI devices, which are paths to access the DD Boost-over-FC service on the required Data Domain system. As part of establishing the connection to the server, the DD Boost-over-FC transport logic provides to the server this catalog of paths.

## Initial path selection

The server maintains statistics on the DD Boost-over-FC traffic over the various target endpoints and known initiators. During the connection setup procedure, path management logic in the server consults these statistics, and selects the path to be used for this connection, based upon the following criteria:

- For queue-depth constrained clients (see below), evenly distribute the connections across different paths.
- Choose the least-busy target endpoint.
- Choose the least-busy initiator from among paths to the selected target endpoint.

## Client path failover

The client may start using a different path because it is directed to do so by the server dynamic re-balancing logic. However, the client may also decide, on its own, to start

using a different available path. This happens if the client receives errors when using the connection's current path.

For example, assume the path catalog for a connection consists of eight paths:

```

/dev/sg21: (A, C, 0)
/dev/sg22: (A, C, 1)
/dev/sg23: (A, D, 0)
/dev/sg24: (A, D, 1)
/dev/sg25: (B, C, 0)
/dev/sg26: (B, C, 1)
/dev/sg27: (B, D, 0)
/dev/sg28: (B, D, 1)

```

The server selects the (A, C, 0) path during initial path selection. The DFC transport logic in the DD Boost library starts sending and receiving data for the connection, using SCSI commands to `/dev/sg21`.

Later, the link from target endpoint C to its switch becomes unavailable, due to cable pull or some hardware failure. Any subsequent SCSI request submitted by the DFC transport logic to `/dev/sg21` fails with an error code indicating that the SCSI request could not be delivered to the device.

In this case, the DFC transport logic looks in the catalog of devices, for a path with a different physical component; that is, a different combination of initiator and target endpoint. The SCSI request is retried on the selected path, and the process is repeated until a path is discovered over which the SCSI request can be successfully completed.

## Dynamic re-balancing

The server periodically performs dynamic re-balancing. This involves consulting the statistics to look for situations where:

- For queue-depth constrained clients (see below), connections are distributed unequally across available paths.
- Workload across target endpoints is out of balance.
- Workload across initiators is out of balance.

If such a situation is discovered, the server can mark one or more connections for server-directed path migration. This is achieved by having the server request, during a future data transfer operation, that the DD Boost library start using a different available path from the catalog for subsequent operations.

## Queue-depth constraints

For the purposes of the DD Boost-over-FC solution, the specific SCSI device over which a request is received is irrelevant. All SCSI devices are identical, destination objects for SCSI commands as required by the SCSI protocol. When processing a SCSI request, the server logic gives no consideration to the specific device on which the SCSI request arrived.

Why bother to allow for more than one device? Because certain client-side operating systems impose a restriction on the number of outstanding I/O requests that can be conducted simultaneously over a given generic SCSI device. For example, the Windows SCSI Pass-Through Interface mechanism will only conduct one SCSI request at a time through each of its generic SCSI devices. This impacts the performance of the DD Boost-over FC solution, if multiple connections (such as backup jobs) try to use the same generic SCSI device.

Additionally, the Data Domain system also imposes a limit on the number of outstanding I/O requests per advertised SCSI device. For performance reasons with larger workloads, multiple SCSI devices may need to be advertised on the Data Domain system.

The term *queue-depth* describes the system-imposed limit on the number of simultaneous SCSI requests on a single device. Client systems, such as Windows, whose queue depth is so low as to impact performance are considered *queue-depth constrained*.

Refer to [Sizing DD Boost-over-FC device-set](#) on page 43 for guidance on how many devices to configure based on the workload, type of Data Domain system, and whether or not the client system is queue-depth constrained.

## Encrypted managed file replication

By default, file replication jobs are set up between two Data Domain systems without encryption after being authenticated by using the preconfigured DD Boost username and password. If the encrypted file replication feature is enabled, the session between the source and destination Data Domain systems is encrypted by using Secure Sockets Layer (SSL), which ensures that all the image data and metadata is sent encrypted over the WAN.

Enabling this option on the Data Domain system is transparent to RMAN. When RMAN requests the Data Domain system to perform a file replication job, the source and destination systems negotiate automatically to perform encryption without the knowledge of RMAN. Encrypted file replication uses the ADH-AES256-SHA cipher suite. There is no ability to configure a different suite in the DD OS.

Encrypted file replication is available to a Data Domain system with an installed Replicator license, and applies to all the file replication jobs on that system. Both the source and the destination Data Domain systems that are participating in file replication jobs must have this option enabled. Otherwise, replication fails.

Encrypted file replication can be used with the encryption of data-at-rest feature that is available on the DD OS with the optional Encryption license. When encrypted file replication is used with the encryption of data-at-rest feature, the encrypted backup image data is encrypted again by using SSL for sending over a WAN.

---

### Note

- Both the source and the destination Data Domain systems must run DD OS 5.0 or later to use this feature.
- Enabling this feature does not require restarting the file system on a Data Domain system.
- The low-bandwidth optimization and the encryption options can be used with each other.

---

[Enabling encrypted file replication](#) on page 39 and the *Data Domain Operating System Administration Guide* provide more information about this topic.

## DD Boost and high availability

Starting with DD OS 5.7.1, Data Domain systems with DD Boost can accommodate high availability (HA) configurations.

During normal operations, DD Boost on the active node sends to the standby node any DD Boost data and state information necessary to continue the DD Boost operations on the standby node if a failure should occur.

---

### Note

DD Boost currently supports only active-standby configurations.

---

DD Boost performs periodic operations to force the user data to disk on the server. DD Boost on the client buffers all the user data between these periodic synchronize-to-disk operations so that if a DD server fails, the data can be resent.

This method also applies to virtual writes. You can mix standard write operations with synthetic write operations.

## DD Boost, HA, and failover

When a Data Domain system with HA enabled fails, recovery occurs in less than ten minutes. Once the failed system recovers, DD Boost recovery begins and applications that use DD Boost automatically recover without failing or receiving an error. DD Boost recovery might take longer than ten minutes because DD Boost recovery cannot start until the failover of the DD system is complete.

No changes are necessary to allow applications to take advantage of DD Boost HA capabilities. With DD Boost 3.2.1 and DD OS 5.7.1 on HA configurations, applications automatically recover if a failover occurs. No action is required from the application.

## Partial HA configurations

Managed file replication (MFR) is supported between any two Data Domain systems that run compatible versions of DD OS, regardless of whether one or both of the DD systems is enabled for HA.

MFR between two HA systems succeeds in the event of failure of either system because both systems support HA. An MFR in progress recovers seamlessly if either the source HA system or the destination HA system fails.

MFR between an HA system and a non-HA system succeeds if the HA system fails, but does not succeed if the non-HA system fails.

### MFR to HA enabled systems

A single-node Data Domain system that runs DD OS 5.7 or later and performs MFR to an HA system recovers seamlessly if the HA system fails. The MFR does not recover seamlessly if the single-node Data Domain source system fails.

### MFR from HA enabled systems

An MFR from an HA system to a single-node Data Domain system that runs DD OS 5.7 or later recovers seamlessly if the source HA system fails. However, the MFR does not recover seamlessly if the single-node Data Domain destination system fails.

In all cases that involve partial HA configurations, the non-HA system must run DD OS 5.7 to enable an MFR to continue seamlessly when a failure occurs. In partial HA

configurations where the non-HA system runs a DD OS version earlier than 5.7, the MFR does not recover seamlessly from a failure of either system.

In all cases, the application must use DD HA Boost 3.2.1 libraries to enable the seamless recovery of the MFR.

## Low-bandwidth optimization

The low-bandwidth Replicator option reduces the WAN bandwidth utilization. This option is useful if file replication is performed over a low-bandwidth network (WAN) link. This option provides additional compression during data transfer and is recommended only for file replication jobs that occur over WAN links that have fewer than 6 Mb/s of available bandwidth.

The low-bandwidth optimization option is available to Data Domain systems with an installed Replicator license. The option is enabled on a Data Domain system and applies to all the file replication jobs on that system.

Enabling this option on Data Domain system is transparent to RMAN. When RMAN requests a Data Domain system to perform a file replication job, the source and destination systems automatically perform the additional compression without involving RMAN.

Each Data Domain system that participates in managed file replication must have this option enabled.

[Enabling low-bandwidth optimization](#) on page 38 and the *Data Domain Operating System Administration Guide* provide more information about this topic.

## Client access validation

Configuring client access validation for DD Boost limits the access to the Data Domain system for DD Boost clients by requiring DD Boost authentication per connection for:

- The initial connection to the Data Domain system
- Each restart of DD Boost (Enable/Disable)
- Each file system restart
- Each Data Domain system reboot

The list of clients can be updated at any time without a restart requirement, which eliminates the access validation impact on jobs in progress.

## In-flight encryption

In-flight encryption enables applications to encrypt in-flight backup or restore data over the LAN from the Data Domain system. This feature was introduced to offer a more secure data transport capability.

When configured, the client is able to use TLS to encrypt the session between the client and the Data Domain system. The specific cipher suite used is either ADH-AES256-SHA, if the HIGH encryption option is selected, or ADH-AES128-SHA, if the MEDIUM encryption option is selected.

## Marker support

The DD Boost library supports application-specific tape markers. To enable the marker support, the application must specify the DDP\_O\_MARKER flag during the file open.

If the flag is specified, the marker settings on the Data Domain system determine the handling of markers. The default is auto marker handling. If any marker is encountered in the first 512 KB of a write stream, support is enabled for that stream. Users can toggle the marker settings on a Data Domain system. The *Data Domain Operating System Administration Guide* provides details.

# CHAPTER 3

## Configuring the Data Domain System

This chapter contains the following topics:

- [Oracle RMAN agent and DD OS upgrades](#).....34
- [Firewalls and ports](#)..... 34
- [Enabling DD Boost on a Data Domain system](#)..... 34
- [Enabling Oracle optimized deduplication](#)..... 35
- [Configuring the DD Boost server](#).....36
- [Configuring DD Boost-over-FC service](#).....41

## Oracle RMAN agent and DD OS upgrades

The Oracle RMAN agent and DD OS maintenance compatibility is defined in the online compatibility guide, which is available at <http://compatibilityguide.emc.com:8080/CompGuideApp/>. Consult that publication before you upgrade either the Oracle RMAN agent or the DD OS.

## Firewalls and ports

The Data Domain system as it is initially configured does not work through a firewall (a database server to a Data Domain system, or from one Data Domain system to another). If you need the Data Domain system to work in the presence of a firewall, contact your network support provider.

The following ports must be open in a firewall to enable the DD Boost backups and file replication:

- TCP 2049 (NFS)
- TCP 2051 (Replication)
- TCP 111 (NFS portmapper)
- TCP *xxx* (Select a port for NFS `mountd`. The default `MOUNTD` port is 2052.)

## Enabling DD Boost on a Data Domain system

Every Data Domain system that is enabled for DD Boost must have a unique name. You can use the Data Domain system's DNS name, which is always unique.

### Procedure

1. On the Data Domain system, log in as an administrative user.
2. Verify that the file system is enabled and running:

```
# fileysys status
```

```
The file system is enabled and running.
```

If disabled, enable the file system:

```
# fileysys enable
```

3. Verify that the DD Boost license is enabled:

```
# license show
```

```
Feature licenses:
```

##	License Key	Feature
1	ABCD-EFGH-IJKL-MNOP	DDBOOST

If disabled, enter the license key that is provided:

```
# license add license-key
```

```
License "ABCE-BCDA-CDAB-DABC" added.
```

#### 4. Establish the DD Boost username and password for the Data Domain system.

---

##### Note

The username and password are case-sensitive, and must match the username and password provided in [Registering each Data Domain system](#) on page 61.

---

```
# user add username password password
# ddbboost set user-name username
```

---

##### Note

If the username is changed or deleted, the change in access rights does not affect any current operations. For example, deleting the current clients from the DD Boost access list by using the `ddbboost access del clients` command does not stop a file backup that is already in progress. All operations in progress continue. The change in access rights does not cause existing operations to fail. It is recommended to change access rights only when no RMAN operations are in progress. The `ddbboost disable` command can be used to prevent operations while access is changed. Disabling DD Boost terminates any existing operations, so be careful that no RMAN operations are in progress when you disable DD Boost. After changing access, `ddbboost enable` re-enables DD Boost and enables RMAN operations that are then subject to the changed access rights.

---

#### 5. Enable DD Boost:

```
# ddbboost enable
```

```
DD Boost enabled
```

## Enabling Oracle optimized deduplication

For Oracle RMAN agent operations, ensure that you have enabled the Oracle optimized deduplication feature at the MTree level in the Data Domain system by running the following command:

```
# mtree option set app-optimized-compression oracle1 mtree
<mtree_path>
```

In the following command output, the MTrees with the value of oracle1 have the Oracle optimized deduplication enabled:

```
# mtree option show
```

Name	Option	Value
/data/coll/db2west	app-optimized-compression	global (none)
/data/coll/nfswest	app-optimized-compression	oracle1
/data/coll/oracleoldf	app-optimized-compression	global (none)
/data/coll/oraclevlab	app-optimized-compression	oracle1
/data/coll/pool-test	app-optimized-compression	global (none)
/data/coll/tdeoracle	app-optimized-compression	oracle1

Refer to the online white paper, [Dell EMC Data Domain Oracle Optimized Deduplication](#), for more details on when to enable the Oracle optimized deduplication feature.

## Configuring the DD Boost server

Use the tasks in this section to configure the DD Boost server.

### Creating storage units

Create one or more storage units on each Data Domain system that will be used with RMAN. Each storage unit name on a single Data Domain system must be unique. However, you can use the same storage unit name on more than one Data Domain system. You must provide the storage unit name when you perform the RMAN backup and restore operations.

Storage unit names are case-sensitive.

[Configuring the software on Linux and UNIX systems](#) on page 56 and [Configuring the software on Windows systems](#) on page 61 provide more information.

#### Procedure

1. On the Data Domain system, enter the following command:

```
# ddbboost storage-unit create storage_unit_name
```

2. Repeat step 1 for each storage unit that you want to create. Create at least one storage unit on each Data Domain system that you will use with RMAN. You can share a storage unit on a Data Domain system with more than one RMAN client system.

## Configuring distributed segment processing

The distributed segment processing option is configured on the Data Domain system and applies to all the database servers and the Oracle RMAN agents that are installed on them.

If distributed segment processing is disabled on the Data Domain system, then the data flow from the database servers is as shown in [Figure 2](#) on page 21. If distributed segment processing is enabled, then the data flow is as shown in [Figure 3](#) on page 22.

To configure the option, run the following command:

```
# ddbboost option set distributed-segment-processing {enabled | disabled}
```

---

#### Note

Enabling or disabling the distributed segment processing option does not require a restart of the Data Domain file system.

Distributed segment processing is supported with the Oracle RMAN agent 1.0 or later communicating with a Data Domain system that runs DD OS 5.0 or later.

Distributed segment processing is enabled by default on a system initially installed with DD OS 5.2. If a system is upgraded from DD OS 5.0.x or 5.1.x to DD OS 5.2, distributed segment processing is left in its previous state.

## Configuring advanced load balancing and link failover

If an interface group is configured, when the Data Domain system receives data from the database server clients, the data transfer is load balanced and distributed as separate jobs on the private network, providing higher throughput, especially for customers who use multiple 1 GbE connections. [Advanced load balancing and link failover](#) on page 22 provides details.

Create an interface group on the Data Domain system by adding existing interfaces to the group and registering the Data Domain system with RMAN, as described in the following steps.

#### Procedure

1. Add the interfaces into the group (the interfaces must already have been created with the `net` command):

```
# ddbboost ifgroup default add interface 192.168.1.1
# ddbboost ifgroup default add interface 192.168.1.2
# ddbboost ifgroup default add interface 192.168.1.3
# ddbboost ifgroup default add interface 192.168.1.4
```

---

#### Note

This example assumes that no additional named interface groups have been created and uses the default interface group. To see help for the `net` command, enter the command with no arguments at the DD OS prompt.

2. Select one interface on the Data Domain system to register with RMAN. It is recommended that you create a failover aggregated interface and register that interface with RMAN.

---

#### Note

It is not mandatory to use an interface in the ifgroup to register with RMAN. An interface that is not part of the ifgroup can also be used to register with RMAN. It is recommended that the interface be registered with a resolvable name using DNS or any other name resolution mechanism.

## 3. Enable the feature on the Data Domain system:

```
# ddbboost ifgroup enable
```

Verify the configuration:

```
# ddbboost ifgroup show config interfaces
```

Group Name	Status	Interface
default	enabled	192.168.1.1
default	enabled	192.168.1.2
default	enabled	192.168.1.3
default	enabled	192.168.1.4

After the interface group is set up, you can add or delete interfaces from the group. [Modifying an interface group](#) on page 91 provides details.

---

#### Note

You can manage the advanced load balancing and link failover through either of the following methods:

- Run the `ddbboost ifgroup` command, as described in [ifgroup](#) on page 104.
  - Use the Data Domain System Manager (DD System Manager) **Data Management > DD Boost** page, as described in the *Data Domain Operating System Administration Guide*.
- 

## Enabling low-bandwidth optimization

To enable the low-bandwidth option, run the following command:

```
# ddbboost file-replication option set low-bw-optim enabled
```

```
Low bandwidth optimization for file-replication set to "enabled".
```

---

#### Note

Enabling or disabling the low-bandwidth optimization option does not require a restart of the Data Domain file system. However, after low-bandwidth optimization is enabled, a full cleaning cycle on the Data Domain system is needed for effective operation.

You can monitor and manage the low-bandwidth optimization through either of the following methods:

- Run the `ddbboost file-replication` command, as described in [file-replication](#) on page 100.
- Use the DD System Manager **Data Management > DD Boost** page, as described in the *Data Domain Operating System Administration Guide*.

No configuration changes are required on the database server as this feature is transparent to RMAN.

---

**Note**

- Enabling this feature takes additional resources of CPU and memory on the Data Domain system, so it is recommended that this option be used only when file replication is being done over low-bandwidth networks with less than 6 Mbps aggregate bandwidth.
  - The *Data Domain Operating System Administration Guide* provides more information about this feature.
- 

## Enabling encrypted file replication

To enable the encrypted file replication option, run the following command:

```
# ddbboost file-replication option set encryption enabled
```

The output indicates that the encryption you requested was enabled.

[file-replication](#) on page 100 provides more details about how to monitor and manage the encryption option for file replication.

Turning on this feature takes additional resources of CPU and memory on the Data Domain system.

---

**Note**

Enabling or disabling the encrypted file replication option does not require a restart of the Data Domain file system. For more information on this feature, see the discussion about this topic in the *Data Domain Operating System Administration Guide*.

---

Encrypted file replication must be enabled on both of the Data Domain systems involved in the replication. If either system does not have encryption file replication enabled, encryption will not be used during replication between the two Data Domain systems.

## Configuring client access validation

Configuring the client access control for DD Boost limits the access to the Data Domain system for DD Boost clients and removes the dependency on the DNS. By default, if no clients are added to the clients list when DD Boost is enabled, all clients are automatically included in the clients list. By default, a \* wildcard is used.

To restrict access, remove the \* wildcard from the list and then add the new clients.

The backup server client list may contain both fully qualified domain names or short names. The backup host's fully qualified domain name needs to be correctly configured for reverse lookup in DNS.

To delete all the clients from the DD Boost clients list, run the following command:

```
# ddbboost clients delete client-list
```

Optionally, to delete all the clients previously added and reset the DD Boost clients list, run the following command:

```
# ddbboost client reset
```

Clients can be added as both fully qualified domain names and short names. To add clients to the DD Boost clients list, run the following command:

```
# ddbboost clients add client-list [encryption-strength {medium | high} authentication-mode {one-way | two-way | anonymous}]
```

For example:

```
# ddbboost clients add ddbboost-dl.emc.com ddbboost-dlddbboost-dl.emc.com
```

```
: Addedddbboost-dl : Added
```

To view the DD Boost clients list, run the following command:

```
# ddbboost clients show config
```

Client	Encryption Strength	Authentication Mode
*	none	none
*.corp.emc.com	medium	anonymous
rtp-ost-ms02.domain	high	anonymous
rtp-ost-ms02.domain.com	high	anonymous

During access validation, the following search order is used to restrict access:

- Wild card \* followed by partial, for example, \*.emc.com followed by \*.com
- Perfect match of sent client name, for example, ddbboost-dl.emc.com

If the search does not find a matching entry for the client, the client will be denied access.

## Enabling in-flight encryption

To enable in-flight encryption for backup and restore operations over a LAN, run the following command:

```
# ddbboost clients add <client-list> [encryption-strength {medium | high} authentication-mode {one-way | two-way | anonymous}]
```

This command can enable encryption for a single client or for a set of clients. The authentication-mode option is used to configure the minimum authentication requirement. A client trying to connect by using a weaker authentication setting will be blocked. Both one-way and two-way authentication require the client to be knowledgeable about certificates.

### One-Way Authentication

The DD Boost client requests authentication from the Data Domain server, and the Data Domain server sends the appropriate certificate to the DD Boost client.

The DD Boost client verifies the certificate. The communication channel between the DD Boost client and the Data Domain server is encrypted.

### Two-Way Authentication

The DD Boost client requests authentication from the Data Domain server using the server certificate. The Data Domain server also requests authentication from the DD Boost client using the client certificate. After authentication through an SSL handshake, the communication channel between the DD Boost client and the Data Domain server is encrypted.

### Anonymous Authentication

No certificates are exchanged, but information is exchanged. After the SSL handshake, the communication channel between the DD Boost client and the Data Domain server is encrypted.

---

#### Note

This option does not apply to DD Boost-over-Fibre Channel (FC). If both IP and FC are in use, encryption can be enabled on IP connections.

---

## Configuring DD Boost-over-FC service

To enable the DD Boost-over-FC service, install the supported Fibre Channel Target HBAs into the system. The *Data Domain Operating System Command Reference Guide* and *Administration Guide* provide information about `scsitarget` as a related command that may be helpful in managing the SCSI target subsystem.

The following additional configuration steps are also required.

### Procedure

1. Enable the DD Boost-over-FC service:

```
# ddbboost option set fc enabled
```

2. Optionally set the DFC-server-name:

```
# ddbboost fc dfc-server-name set <server-name>
```

Or accept the default, which is the base hostname of the Data Domain system. A valid `dfc-server-name` consists of one or more of the following characters:

- lower-case letters (“a”–“z”)
- upper-case letters (“A”–“Z”)
- digits (“0”–“9”)
- underscore (“\_”)
- dash (“-”)

---

**Note**

- The dot or period character (“.”) is not valid within a `dfc-server-name`; this precludes using the fully-qualified domain name of a Data Domain system as its `dfc-server-name`.
  - Similar to IP hostnames, the `dfc-server-name` is not case-sensitive. Multiple Data Domain systems accessible by the same clients using DDBoost-over-FC should be configured without case-sensitive `dfc-server-name`.
- 

**3. Create a SCSI target access group:**

```
# ddbost fc group create <group-name>
```

For example:

```
# ddbost fc group create lab_group
```

**4. To display the available list of `scsitarget` endpoint:**

```
# scsitarget endpoint show list
```

Endpoint	System Address	Transport	Enabled	Status
endpoint-fc-0	6a	FibreChannel	Yes	Online
endpoint-fc-1	6b	FibreChannel	Yes	Online

**5. Configure the device set of the SCSI target access group for one of the endpoint:**

```
# ddbost fc group modify <group-name> device-set count count
endpoint endpoint-list
```

For example:

```
# ddbost fc group modify lab_group device-set count 8
endpoint 6a
```

**6. Add initiators to the SCSI target access group:**

```
# ddbost fc group add group-name initiator initiator-spec
```

For example:

```
# ddbost fc group add lab_group initiator
"initiator-15,initiator-16"
```

## Sizing DD Boost-over-FC device-set

The Data Domain system advertises one or more "DFC devices" of type Processor, which the DD Boost library uses to communicate with the DD Boost-over-FC service. On the Data Domain system, access to these DFC devices is granted to one or more initiators by adding the initiators to a `ddboost-type scsitarget` access group:

```
# ddboost fc group add lab_group initiator
"initiator-15,initiator-16"
```

The number of DFC devices advertised to the initiator is controlled by configuring the device-set of the `scsitarget` access group:

```
# ddboost fc group modify lab_group device-set count 4
```

The maximum number of supported DFC devices per Data Domain system is 64. You can have the same devices in multiple groups, but each group is limited to 64 devices.

---

### Note

AIX DDdfc drivers support 128 devices. However, if you use the `disk` option with the `ddboost fc add` command, this limitation is removed.

---

Because the DFC client sees each path to the Data Domain system as a separate device, more paths and more DFC devices mean better performance for constrained clients such as AIX, Windows, and Solaris.

So, how many DFC devices should be advertised to initiators on a given backup server? The answer depends upon several factors:

1. Is the backup server queue-depth constrained?

Windows platforms are considered "queue-depth constrained," because the Windows SCSI Pass-Through Interface mechanism will only conduct 1 SCSI request at a time through each of its generic SCSI devices. This impacts the performance of the DD Boost-over FC solution, if multiple connections (for example, backup jobs) are trying to use the same generic SCSI device. So, for Windows platforms running more than one job, it is useful to advertise multiple DFC devices.

Contrast this with the behavior of the Linux SCSI Generic driver, which imposes no such restriction. Linux is not considered "queue-depth constrained," so it is sufficient to simply advertise one DFC device to initiators on Linux systems.

2. Number of physical paths between backup server and Data Domain system

For each advertised DFC device, the backup server operating system will create  $n$  generic SCSI devices, one for each physical path through which the backup server OS can access the device.

For example, if:

- Backup server has 2 initiator HBA ports (A and B)
- Data Domain System has 2 FC target endpoints (C and D)
- Fibre Channel Fabric zoning is configured such that both initiator HBA ports can access both FC target endpoints

then the backup server OS will see each device through four physical paths:

A -> C  
A -> D  
B -> C  
B -> D

and will create 4 generic SCSI devices for each advertised DFC device.

For a Windows backup server (with its queue-depth=1 limitation), this allows up to 4 simultaneous SCSI requests to the Data Domain system, even with only one DFC device advertised.

## Sizing calculation

The following calculation may be used to determine the number of DFC devices to advertise on the Data Domain system and to the initiators on a specified media server. It is recommended that the same number of DFC devices be advertised to all initiators on the same media server.

### On the Data Domain system

The Data Domain system imposes a limit on the number of simultaneous requests to a single DFC SCSI device. Because of this limit, the number of devices advertised needs to be tuned depending on the maximum number of simultaneous jobs to the system at any given time. In general, the larger the number of jobs expected from media servers using DD Boost over FC, the higher the number of devices advertised.

Let J be the maximum number of simultaneous jobs running using DFC, to the Data Domain system at any given time.

Let C be the maximum number of connections per job:

- 3 for DD Extended Retention systems
- 1 for other types of Data Domain systems

Calculate:

- Maximum simultaneous connections to the Data Domain system, using DFC, from all backup servers:
  - $S = J * C$
  - DFC device count  $D = \text{minimum}(64, 2 * (S/128))$ , round up
  - All DFC access groups must be configured with "D" devices.

### Example

Assume:

- 8 backup servers, DD Extended Retention systems, each server running a maximum of 30 jobs at any given time.
- Here,  $J = 8 * 30 = 240$ ,  $C = 3$  (DD Extended Retention systems),  $S = J * C = 720$ ,  $D = 2 * 720 / 128 = 11.25$ , round up to 12.
- As a result, all DFC groups on the Data Domain system must be configured with 12 devices.

### Linux backup servers

The number of DFC devices advertised on the Data Domain system using the calculations listed under [On the Data Domain system](#) on page 44 is sufficient for Linux backup servers. No additional configuration is required. Linux backup servers are not queue-depth constrained, so many connections can share the same DFC generic SCSI device with no performance impact.

### Windows backup servers

The Data Domain server path management logic spreads out connections across available logical paths (Initiator, Target Endpoint, DFC Device). Enough DFC devices must be configured so that each connection uses its own generic SCSI device (logical path) on the backup server, with a maximum DFC device count of 64.

Let  $X$  be the number of DFC devices configured on the Data Domain system (from [On the Data Domain system](#) on page 44).

Let  $P$  be the number of physical paths between backup server and Data Domain system.

Let  $J$  be the maximum number of simultaneous jobs.

Let  $C$  be the maximum number of connections per job:

- 3 for DD Extended Retention systems
- 1 for other types of Data Domain systems

Calculate, assuming the following:

- 4 physical paths between the backup server and Data Domain system, 30 maximum jobs, DD Extended Retention system
- In this case,  $X = 25$ ,  $P = 4$ ,  $J = 30$ , and  $C = 3$
- Maximum simultaneous connections from backup server  $S = (J * C) = 90$
- DFC device count  $D = (90/4, 25) = 25$

As a result, the Data Domain system should be configured to advertise 25 devices to each initiator on the backup server.

Assume:

- 2 physical paths between the backup server and Data Domain system, 50 maximum jobs, single Data Domain system
- In this case,  $X=18$ ,  $P = 2$ ,  $J = 40$ ,  $C = 1$
- Maximum simultaneous connections from backup server  $S = (J * C) = 40$
- DFC device count  $D = \text{maximum}(40/2, 18) = 20$

As a result, the Data Domain system should be configured to advertise 20 devices to each initiator on the backup server.

Because the value of  $D$  (20) is greater than the value of  $X$  (18), it is sufficient to configure two devices only for the DFC access group with Windows clients.

### AIX backup servers

For AIX, the proprietary device entries are exclusively locked on a per-process basis—one and only one process can use a device entry. Calculations are based on application instance usage. If an application spawns multiple processes, each process exclusively locks at least one device entry. Multi-threaded applications lock one device per thread. For these reasons, you should configure the Data Domain system to advertise as many DFC devices as possible (up to the maximum of 128). A `Device Busy` error can result if there are not enough devices accessible to the AIX clients.

---

#### Note

If you are using the proprietary device driver, the total number of streams in a policy should not exceed the number of AIX DFC devices available; otherwise, the backup job might fail.

---

### Solaris backup servers

For Solaris, device entries are exclusively locked on a per-process basis—one and only one process can use a device entry. Calculations are based on application instance

usage. If an application spawns multiple processes, each process exclusively locks at least one device entry. Multi-threaded applications lock one device per thread. For these reasons, configure the Data Domain system to advertise as many DFC devices as possible to avoid "in use" errors from the sgen device driver. A `Device Busy` error can result if there are not enough devices accessible to the Solaris clients.

The number of sgen devices is the number of Fibre Channel ports accessible to the Solaris instance times the number of different paths to the Data Domain system endpoints times the number of LUNs in the access group.

## Configuring the SCSI device driver for AIX clients

DD Boost-over-FC is supported on clients that run the supported AIX versions on IBM Power hardware. DFC for AIX uses the SCSI generic device driver, which is included in the AIX installation.

By default, the Oracle RMAN agent uses the SCSI generic device driver, which does not require the installation of any drivers.

To check for the SCSI generic device driver, you can run the `lsdev` command and review the command output. The command output for this device driver includes the device names as `hdisk*` and the type as `Other FC SCSI Disk Drive`. For example:

```
lsdev

hdisk1    Available 05-00-01    Other FC SCSI Disk Drive
hdisk2    Available 05-00-01    Other FC SCSI Disk Drive
hdisk3    Available 05-00-01    Other FC SCSI Disk Drive
```

---

### Note

To discover the DD Boost devices that are added to the DD Boost FC group in Data Domain system, run the `cfgmgr` command on the client. You might need to run the `cfgmgr` command if the `lsdev` command cannot show all the devices.

---

To complete the configuration and integration of the device driver on an AIX system, perform the following steps.

1. Follow the required instructions in the *Data Domain Fibre Channel Configuration and Integration with Data Domain Boost for Enterprise Applications Database Application Agent Technical Notes*. Follow the instructions on configuration of an AIX DFC setup without DFC drivers by adding disk type devices. The document is available on the Support website at <https://support.emc.com>.
2. To ensure that the DBA can perform backup and restore operations as a non-root user, follow the configuration instructions in the section "AIX with SCSI generic disk driver" in the following Knowledgebase article:

*Fibre Channel Devices with Products using DD Boost in Linux/UNIX Environment*  
(Article Number 000182275)

The article is available on the Support website at <https://support.emc.com>.

## Configuring the SCSI generic device driver for Solaris clients

DD Boost-over-FC is supported on clients that run Solaris 10 and 11 on SPARC and x86 hardware. DFC for Solaris uses the SCSI generic device driver (sgen), which is

included in the Solaris installation. Use the following procedure to ensure that sgen successfully identifies the processor devices at startup.

### Procedure

1. Add the following line in the forceload section of `/etc/system`:

```
forceload: drv/sgen
```

This step should resolve issues with sgen not properly loading during startup and keep the sgen driver loaded.

2. To check for existing usage of sgen, run the command `grep sgen /etc/driver_aliases`.

---

#### Note

The existence of a `/dev/scsi`, `/dev/scsi/processor`, or `/dev/scsi/*` directory does not necessarily mean that sgen is currently configured. There could be dangling files.

---

3. If there is no existing use of sgen, or if sgen is used only for `"scsiclass,03"`, perform the following steps:
  - a. Run the command `rem_drv sgen`.
  - b. Run the command `add_drv -m '* 0600 root sys' -i '"scsiclass,03"' sgen`.

---

#### Note

It is critical that you use single and double quotes exactly as shown.

---

This command should return to the prompt with no errors or warnings. Check connectivity to the Data Domain system. There should be at least one file in `/dev/scsi/processor`.

- c. To confirm at least one entry for three configuration files, run the command `grep sgen /etc/minor_perm /etc/name_to_major and /etc/driver_aliases`.

Example results of this command are as follows:

```
/etc/minor_perm:sgen * 0600 root sys
/etc/name_to_major:sgen 151
/etc/driver_aliases:sgen "scsiclass,03"
```

---

#### Note

The `name_to_major` number will likely be different than this example.

---

4. If the sgen device is already used for other devices, perform the following steps:
  - a. Run the command `rem_drv sgen`.
  - b. Run the command `add_drv -m '* 0600 root sys' -i '"scsiclass,03" "scsiclass,XX"' sgen`.

---

**Note**

XX would be the device type from a previously run command `grep sgen /etc/driver_aliases`. It is critical that you use single and double quotes exactly as shown.

---

An example of this command is as follows: `add_drv -m '* 0600 root sys' -i '"scsiclass,03" "scsiclass,06"' sgen`.

This command should return to the prompt with no errors or warnings. Check connectivity to the Data Domain system. There should be at least one file in `/dev/scsi/processor`.

- c. To confirm at least one entry for three configuration files, run the command `grep sgen /etc/minor_perm /etc/name_to_major` and `/etc/driver_aliases`.
- d. Open the `/kernel/drv/sgen.conf` file. If the `device-type-config-list` is uncommented, add “processor” to the list to ensure that the driver is recognized. For example, if the `device-type-config-list` is uncommented as in this example:

```
device-type-config-list="direct", "sequential", "worm",
"rodirect", "optical", "changer";
```

Change the entry to:

```
device-type-config-list="direct", "sequential", "worm",
"rodirect", "optical", "changer", "processor";
```

# CHAPTER 4

## Configuring the Oracle RMAN Agent

This chapter contains the following topics:

- [Installing and configuring the Oracle RMAN agent on AIX, HP-UX, Linux, and Solaris systems](#)..... 50
- [Installing and configuring the Oracle RMAN agent on Windows systems](#)..... 57
- [Registering each Data Domain system](#)..... 61
- [Performing backups of a clustered environment to Data Domain](#)..... 65
- [Data Domain replication](#)..... 65
- [Specifying subdirectories for Oracle RMAN agent backups](#)..... 70
- [Performing backups and restores of Oracle CDBs and PDBs](#)..... 73
- [Performing DD Boost compressed restores](#)..... 74
- [Display and deletion of backups](#)..... 74
- [Configuring the use of Data Domain Cloud Tier for data movement to the cloud](#)..... 81
- [Usage limits on Data Domain resources](#)..... 83

# Installing and configuring the Oracle RMAN agent on AIX, HP-UX, Linux, and Solaris systems

## Installing the Oracle RMAN agent on database servers

Install the Oracle RMAN agent on all the database servers that must access the Data Domain system.

The pre-4.0 versions of Oracle RMAN agent are installed under the `ORACLE_HOME` directory. Starting with Oracle RMAN agent version 4.0, the software must be installed in a directory other than the `ORACLE_HOME` directory.

Run the `install.sh` script to install the Oracle RMAN agent 4.5 or to upgrade from an earlier version of the Oracle RMAN agent. The script installs the Oracle RMAN agent in a user-specified directory or in the default installation directory, `$HOME/opt/dpsapps/rmanagent`.

Run the `install.sh -h` or `install.sh --help` command to obtain more information about the script operation.

---

### Note

The user-specified installation directory must be a directory that is created specifically for the Oracle RMAN agent 4.5 installation, and must not be the `ORACLE_HOME` directory. The complete directory pathname must be specified, without a slash (/) at the end.

---

To upgrade from a pre-4.0 version, use either of the following methods:

- Run the uninstallation script, `uninstall.sh`, to uninstall the pre-4.0 software and then install the Oracle RMAN agent 4.5. [Uninstalling the Oracle RMAN agent on database servers](#) on page 54 describes how to use the uninstallation script.
- Run the installation script to install the version 4.5 software and at the same time remove or update the pre-4.0 Oracle RMAN agent settings.

For an upgrade, if you want to continue using a pre-4.0 RMAN script without any changes, it is recommended that you run the installer script for each `ORACLE_HOME`. If you do not want to reuse the RMAN script, run the installer script once and then modify the RMAN script for Oracle RMAN agent 4.5.

You can run the `install.sh` script manually or automatically. To enable the automatic operation, you must set the environment variables as listed in [Table 3](#) on page 51:

- When the variables are not set, the script runs manually and prompts you for the required values.
- When the variables are set, the script runs automatically and performs the installation or upgrade according to the environment variable settings.

---

### Note

`ORACLE_HOME` is an Oracle environment variable. The other variables are the Oracle RMAN agent environment variables.

---

**Table 3** Environment variables for installation or upgrade of Oracle RMAN agent

Environment variable	Description	Default and valid values
<i>ORACLE_HOME</i>	Specifies the installation location of the pre-4.0 Oracle RMAN agent.	<ul style="list-style-type: none"> <li>Undefined (default).</li> <li>Valid value of the <code>ORACLE_HOME</code> directory under which the pre-4.0 Oracle RMAN agent is installed.</li> </ul>
<i>RMAN_AGENT_HOME</i>	Specifies the installation directory for the Oracle RMAN agent 4.5.	<ul style="list-style-type: none"> <li><code>\$HOME/opt/dpsapps/rmanagent</code> (default).</li> <li>Valid complete pathname of the directory for the Oracle RMAN agent 4.5 installation.</li> </ul> <hr/> <p><b>Note</b></p> <p>The directory pathname must not end with a slash (/).</p>
<i>RMAN_AGENT_INSTALL_OPTIONS</i>	Required only when Oracle RMAN agent already exists in the installation directory. Specifies whether to bypass the version 4.5 installation or continue and overwrite the existing installation.	<ul style="list-style-type: none"> <li>Undefined (default).</li> <li><code>BYPASS</code> or <code>bypass</code>—Specifies to bypass the version 4.5 installation.</li> <li><code>OVERWRITE</code> or <code>overwrite</code>—Specifies to continue the version 4.5 installation and overwrite the existing installation in the destination directory.</li> </ul>
<i>RMAN_AGENT_UPGRADE_OPTIONS</i>	Required only if a pre-4.0 Oracle RMAN agent is installed. Specifies whether to keep or remove the pre-4.0 installation, and whether to update the Oracle settings to enable the reuse of pre-4.0 RMAN scripts without any script changes.	<ul style="list-style-type: none"> <li>Undefined (default).</li> <li><code>KEEP</code> or <code>keep</code>—Specifies to keep the pre-4.0 installation.</li> <li><code>UNINSTALL</code> or <code>uninstall</code>—Specifies to uninstall the pre-4.0 installation.</li> <li><code>UPGRADE</code> or <code>upgrade</code>—Specifies to uninstall the pre-4.0 installation and also update the Oracle settings so that the pre-4.0 RMAN scripts can be reused without any script changes.</li> </ul>

Perform the download and installation steps for the environment as described in the following topic.

## Installing the Oracle RMAN agent on Linux and UNIX

### Procedure

1. Ensure that no backups are running. Stop the RMAN processes before you install the Oracle RMAN agent.

**Note**

If you use AIX and you experience copy failures, you might need to run `/usr/bin/slibclean` as the root user after you stop the RMAN processes.

2. Download the Oracle RMAN agent as the software package for the operating system from the Support website at <https://support.emc.com>.

The following table shows the software package names for the supported operating systems.

**Table 4** Software package names on UNIX and Linux

Operating system	Software package name
AIX (64-bit)	ddrman45_aixpower
HP Itanium (64-bit)	ddrman45_hp11ia64
Linux x64 (64-bit)	ddrman45_linux_x86_64
Solaris AMD (64-bit)	ddrman45_solaris_x86_64
Solaris SPARC (64-bit)	ddrman45_solaris64

3. Use the `gunzip` and `tar` commands to uncompress the downloaded file:

```
# gunzip <file_name>.gz
# tar -vxf <file_name>
```

4. If you want the installation script to run automatically, ensure that you set the multiple environment variables that are described in [Installing the Oracle RMAN agent on database servers](#) on page 50.

For example, you want to run an automatic installation that overwrites an existing Oracle RMAN agent 4.5 in the default directory and uninstalls a pre-4.0 installation under `ORACLE_HOME` without any prompts. In this case, run the following commands to set the required environment variables before you run the installation script:

```
# export ORACLE_HOME=/home/oracle/app/oracle/product/11.2.0/dbhome_1
# export RMAN_AGENT_INSTALL_OPTIONS=overwrite
# export RMAN_AGENT_HOME=/opt/dpsapps/rmanagent
# export RMAN_AGENT_UPGRADE_OPTIONS=uninstall
```

To verify the value of an environment variable, run the `echo` command. For example:

```
# echo $RMAN_AGENT_HOME
```

```
/home/oracle/opt/dpsapps/rmanagent
```

5. As an Oracle user, run the `install.sh` script:

```
# ./install.sh
```

If `RMAN_AGENT_HOME` is not set, the script displays the following prompt:

```
Do you want to install under the default installation directory
$HOME/opt/dpsapps/rmanagent? (y or n)
```

Type the appropriate value:

- To install to the default directory, type **y**.
- To install in a non-default directory, type **n**.

The script prompts you to enter the installation directory pathname:

```
A directory outside of the ORACLE_HOME directory must have
been created specifically for this installation. Enter the
full pathname of the installation destination:
```

Type the pathname of an already created installation directory.

6. If a previous version of the Oracle RMAN agent is installed:

a. The installer prompts whether to uninstall pre-4.0 software. Type the appropriate value:

- To uninstall the pre-4.0 software, type **y**. The lockbox is kept under the `$ORACLE_HOME/config` directory.
- To keep and use the pre-4.0 software, type **n**. You can later run the uninstallation script if required to uninstall the pre-4.0 software, as described in [Uninstalling the Oracle RMAN agent on Linux and UNIX](#) on page 56.

b. If you selected to uninstall pre-4.0 software in the preceding step, the installer prompts whether to update the Oracle settings under `$ORACLE_HOME`. The updated settings will enable the reuse of RMAN scripts from the previous version without any script changes. Type the appropriate value:

- To enable the reuse of the pre-4.0 RMAN scripts without script changes, type **y**.
- To prevent updates of the Oracle settings, type **n**. Then, to make any pre-4.0 RMAN script usable with the Oracle RMAN agent 4.5, you must change the following items in the script:
  - Set `SBT_LIBRARY` to `$RMAN_AGENT_HOME/lib/libddobk.so`.
  - Replace the `ORACLE_HOME` setting with the `RMAN_AGENT_HOME` setting.

The Oracle RMAN agent 4.5 files are installed in the following directories on a Linux or UNIX system.

**Table 5** Product software components on Linux or UNIX

Software components	Installation directory
DD Boost library files:	<code>\$RMAN_AGENT_HOME/lib</code>

**Table 5** Product software components on Linux or UNIX (continued)

Software components	Installation directory
libddobk.so libDDBoost.so	
<b>RSA library files:</b> libccme_base.so libcryptocme.so libcryptocme.sig libCSP-lb.so libCSP-Security.so	\$RMAN_AGENT_HOME/lib
<b>Lockbox files:</b> ddboost.config ddboost.config.bak ddboost.config.FCD ddboost.config.bak.FCD	\$RMAN_AGENT_HOME/config  <b>Note</b> If lockbox files exist that were created by a pre-4.0 installation and no new lockbox files are created by the version 4.5 software, the pre-4.0 lockbox files are copied to this location during the installation.
ddutil binary	\$RMAN_AGENT_HOME/bin
ddutil log files	\$RMAN_AGENT_HOME/logs
rman_agent_home.cfg  <b>Note</b> This configuration file is created during an update from a pre-4.0 installation.	\$ORACLE_HOME/config

[Installing and Configuring the Oracle RMAN Agent](#) on page 121 provides examples.

## Uninstalling the Oracle RMAN agent on database servers

Run the `uninstall.sh` script to uninstall the Oracle RMAN agent 4.5. You can also run the script to uninstall a previous version of the Oracle RMAN agent.

Run the `uninstall.sh -h` or `uninstall.sh --help` command to obtain more information about the script operation.

You can run the `uninstall.sh` script manually or automatically. To enable the automatic operation, you must set the appropriate environment variables as listed in [Table 6](#) on page 55:

- When the variables are not set, the script runs manually and prompts for the required values.
- When the variables are set, the script runs automatically and performs the uninstallation according to the environment variable settings.

**Note**

*ORACLE\_HOME* is an Oracle environment variable. The other variables are the Oracle RMAN agent environment variables.

**Table 6** Environment variables for uninstallation of Oracle RMAN agent

Environment variable	Description	Default and valid values
<i>ORACLE_HOME</i>	Specifies the installation location of a pre-4.0 Oracle RMAN agent.	<ul style="list-style-type: none"> <li>Undefined (default).</li> <li>Valid value of the <i>ORACLE_HOME</i> directory under which the pre-4.0 Oracle RMAN agent is installed.</li> </ul>
<i>RMAN_AGENT_HOME</i>	Specifies the installation directory for the Oracle RMAN agent 4.5.	<ul style="list-style-type: none"> <li><code>/home/oracle1/opt/dpsapps/rmanagent</code> (default).</li> <li>Valid complete pathname of the directory for the Oracle RMAN agent 4.5 installation.</li> </ul> <p><b>Note</b></p> <p>The directory pathname must not end with a slash (/).</p>
<i>RMAN_AGENT_UNINSTALL_OPTIONS</i>	Required only for an uninstallation of Oracle RMAN agent 4.5. Specifies the components to uninstall for the version 4.5 software.	<ul style="list-style-type: none"> <li>Undefined (default).</li> <li>NONE or none—Specifies to keep the version 4.5 software, and not perform the uninstallation.</li> <li>BINARY or binary—Specifies to uninstall the version 4.5 software, but not the lockbox.</li> <li>FULL or full—Specifies to uninstall both the version 4.5 software and the lockbox.</li> </ul>
<i>RMAN_AGENT_UNINSTALL_PRE40_OPTIONS</i>	Required only for an uninstallation of pre-4.0 Oracle RMAN agent software. Specifies the components to uninstall for the pre-4.0 software.	<ul style="list-style-type: none"> <li>Undefined (default).</li> <li>NONE or none—Specifies to keep the pre-4.0 software, and not perform the uninstallation.</li> <li>BINARY or binary—Specifies to uninstall the pre-4.0 software, but not the lockbox.</li> <li>FULL or full—Specifies to uninstall both the pre-4.0 software and the lockbox.</li> </ul>

Perform the uninstallation steps as described in the following topic.

## Uninstalling the Oracle RMAN agent on Linux and UNIX

---

### Note

It is not necessary to uninstall the Oracle RMAN agent for an upgrade on Linux or UNIX. An existing Oracle RMAN agent is overwritten during an upgrade.

---

### Procedure

1. Stop all the RMAN and Oracle services before you uninstall the Oracle RMAN agent.
2. If you want the uninstallation script to run automatically, ensure that the required environment variables are set as described in [Uninstalling the Oracle RMAN agent on database servers](#) on page 54. To verify the value of an environment variable, run the `echo` command. For example:

```
# echo $RMAN_AGENT_HOME
```

```
/home/oracle/opt/dpsapps/rmanagent
```

To set the `RMAN_AGENT_HOME` environment variable, run the `export` command. For example:

```
# export RMAN_AGENT_HOME=/opt/dpsapps/rmanagent
```

3. As an Oracle user, run the `uninstall.sh` script:

```
# ./uninstall.sh
```

4. If the script does not run automatically, type the appropriate values at the prompts.

When prompted for a directory pathname for the uninstallation, you can specify the pathname of the Oracle RMAN agent that is installed. Type the complete pathname without a slash (/) at the end.

If the script detects an additional installation of Oracle RMAN agent, the script prompts whether to uninstall that version. You can specify to keep or uninstall the software.

## Configuring the software on Linux and UNIX systems

From within RMAN, configure a channel.

---

### Note

The storage unit specified in the following command must already exist on the Data Domain system.

---

The following example shows the command syntax for configuring a channel in a Linux or UNIX environment.

```

RMAN> CONFIGURE CHANNEL DEVICE TYPE 'SBT_TAPE' TRACE <trace-level>
PARMS 'BLKSIZE=1048576,
SBT_LIBRARY=<path-to-$RMAN_AGENT_HOME>/lib/libddobk.so,
ENV=(STORAGE_UNIT=<storage-unit-on-Data-Domain-system>,
BACKUP_HOST=<hostname-of-Data-Domain-system>,
RMAN_AGENT_HOME=<path-to-$RMAN_AGENT_HOME>';

```

The configuration that you specify is also displayed in the Media Management Settings field of Oracle Enterprise Manager (OEM).

## Installing and configuring the Oracle RMAN agent on Windows systems

### Installing the Oracle RMAN agent on database servers

Install the Oracle RMAN agent on all the database servers that must access the Data Domain system. The installation can be performed by an administrator or an Oracle user.

The pre-4.0 versions of Oracle RMAN agent are installed under the `ORACLE_HOME` directory. Starting with Oracle RMAN agent version 4.0, the software must be installed in a directory other than the `ORACLE_HOME` directory.

Before you install the Oracle RMAN agent, ensure that all the Oracle and RMAN processes are stopped and no RMAN backup or restore processes are running.

During an upgrade of a preexisting Windows Oracle RMAN agent, the installation might fail unless all the Oracle instances are stopped. Stop all the Oracle instances before you perform a Windows installation.

The user that performs the Oracle RMAN agent installation must have permissions to install the Microsoft Installer (MSI) technology binaries. If the local users do not have such permissions, the administrator must perform the installation on behalf of the local users.

For each user that performs the installation, a set of Oracle RMAN agent files is installed. By default, the software files are installed under the user profile directory. If you want a single copy of the software files to be installed per computer, specify a non-default shared directory for the installation and manually set the appropriate read and write permissions on the directory.

An installation of the Oracle RMAN agent requires at least 50 MB of available disk space in the installation user profile directory.

Perform the download and installation steps for the environment as described in the following topic.

### Installing the Oracle RMAN agent on Windows systems

---

#### Note

During an upgrade of Oracle RMAN agent on Windows, you do not need to uninstall the previous-version software. An existing Oracle RMAN agent is overwritten during an upgrade.

---

## Procedure

1. Stop all the RMAN and Oracle processes before you install the Oracle RMAN agent.
2. Download the Oracle RMAN agent as the Windows software package from the Support website at <https://support.emc.com/>.

The Windows software package name is `ddrman45_ntx64`.

3. Install the Oracle RMAN agent by using either of the following methods:

- Interactive method with the installation wizard:
  - a. Double-click `Oracle RMAN agent.exe` to launch the installation wizard.

The **Welcome to the Setup Wizard** page appears.

---

### Note

Run only one instance of the installation wizard at a time. If you try to run more than one instance, the wizard operation terminates with an error.

---

- b. Read the terms of the Software License and Support Services Agreement, select **I Agree**, and then click **Next** to continue.
- c. The **Change Install Location** page displays the default destination directory for the installation as `%LOCALAPPDATA%\RMANAgent`, for example, `C:\Users\Administrator\AppData\Local\RMANAgent`. The Oracle RMAN agent components will be installed in subdirectories in this destination directory.

Accept or change the installation directory:

- To accept the displayed directory for the installation, click **Install**.
  - To change the displayed directory, click **Change** and browse to select a different directory. When the correct directory is displayed, click **Install**.
- d. If any previous versions of Oracle RMAN agent are installed, an **Upgrade Options** page displays a list of the previous installations. Select to keep or remove the previous installations as required:
    - To uninstall all the previous installations, select **Remove all the previous Oracle RMAN agent installations**.
    - To copy the pre-4.0 lockbox files from one of the displayed installations to the `%RMAN_AGENT_HOME%\config` directory:
      - Select **Copy a legacy lockbox into the new installation directory (select only one item)**.
      - Select only one of the displayed pre-4.0 installations in the list.

### Note

The uninstallation of previous versions is irreversible. Pre-4.0 lockbox files are not removed by the version 4.5 installation.

---

When you have completed the selections on the **Upgrade Options** page, click **Install** and respond to any displayed prompts.

- e. The **Installation Progress** window displays a progress bar. To cancel the installation at any time, click **Cancel**.

f. When the installation completes, the **Complete the Setup** window displays information about the log, any errors, and the software installation status. Click **Finish** to exit the wizard.

- Silent installation method:

Run the `Oracle RMAN agent.exe` command with the appropriate command options:

```
Oracle RMAN agent.exe <installation_options>
```

where *<installation\_options>* includes the following options:

- `/s /qn` are the only mandatory options, to suppress the GUI output. Without these options, the Windows installer GUI is launched to perform the installation.
- `ProductInstallPath <pathname>` specifies the installation directory for the Oracle RMAN agent files. Enclose *<pathname>* in double quotes if it includes a space. The default directory is `%LOCALAPPDATA%\RMANAgent`, for example, `C:\Users\Administrator\AppData\Local\RMANAgent`. Use this option to specify a non-default installation directory.
- `RemoveOldVersions=1` specifies to uninstall all pre-4.0 versions of Oracle RMAN agent that are installed on the Windows system. The uninstallation is irreversible, and does not display any prompts for confirmation.
- `CopyLockboxBinaries <pathname>` specifies to copy the pre-4.0 lockbox files from the *<pathname>* directory to the `%RMAN_AGENT_HOME%\config` directory. The Oracle RMAN 4.5 installation directory is `%RMAN_AGENT_HOME%`.

---

#### Note

You can use the `CopyLockboxBinaries` option only when `RemoveOldVersions=1` appears in the same command. The command can copy the lockbox files from only one pre-4.0 directory. Pre-4.0 lockbox files are not removed by a version 4.5 installation.

---

For example, the following command silently installs the Oracle RMAN agent 4.5 software in the directory `C:\Program Files\RMANAgent`, and uninstalls all pre-4.0 versions of the software. The command also copies the pre-4.0 lockbox files from `C:\app\oracle\product\12.1.0\dbhome_2\config` to `%RMAN_AGENT_HOME%\config`:

```
"Oracle RMAN agent.exe" /s /qn ProductInstallPath="C:\Program Files\RMANAgent" RemoveOldVersions=1 CopyLockboxBinaries="C:\app\oracle\product\12.1.0\dbhome_2"
```

The Oracle RMAN agent 4.5 files are installed in the following directories on a Windows system.

**Table 7** Product software components on Windows

Software components	Installation directory
<b>DD Boost library files:</b> libddobk.dll libDDBoost.dll libDataDomainoncrpc.dll libDataDomainpthread.dll	%RMAN_AGENT_HOME%\BIN
<b>RSA library files:</b> ccme_base.dll ccme_base_non_fips.dll cryptocme.dll cryptocme.sig CSP-lb.dll CSP-Security.dll	%RMAN_AGENT_HOME%\cst
<b>Other library files:</b> msvcr80.dll msvcr100.dll	%RMAN_AGENT_HOME%\BIN
<b>Lockbox files:</b> ddbboost.config ddbboost.config.bak ddbboost.config.FCD ddbboost.config.bak.FCD	%RMAN_AGENT_HOME%\config
ddutil.exe <b>binary</b>	%RMAN_AGENT_HOME%\BIN
ddutil.exe <b>log files</b>	%RMAN_AGENT_HOME%\logs
rman_agent_home.cfg	%ORACLE_HOME%\config
<b>Note</b> This configuration file is created during an update from a pre-4.0 installation.	

## Uninstalling the Oracle RMAN agent on Windows systems

**Note**

During an upgrade of Oracle RMAN agent on Windows, you do not need to uninstall the previous-version software. An existing Oracle RMAN agent is overwritten during an upgrade.

## Procedure

1. Stop all the RMAN and Oracle services.
2. Uninstall the Oracle RMAN agent by using the uninstall/change program feature in the Windows Control Panel, as used to uninstall a typical Windows program. Be careful to select and uninstall the correct item for the Oracle RMAN agent.

Open the Windows Control Panel, and then select **Program and Features**.

3. In the **Uninstall or change a program** window, select **Oracle RMAN agent** and then click **Uninstall**.
4. In the **Install Modification** window, select the **Remove** radio button and then click **Remove**.

---

### Note

The **Change** option is not supported for the Oracle RMAN agent. The **Repair** option can be selected to restore the installation to its original state without affecting any user-generated files or configuration files.

The **Uninstallation Progress** window displays a progress bar. The software files are uninstalled without removing any user-generated files or configuration files.

5. When the uninstallation completes, the **Complete the Setup** window displays information about the log, any errors, and the software uninstallation status. Click **Finish** to exit the wizard.

## Configuring the software on Windows systems

From within RMAN, configure a channel.

---

### Note

The storage unit specified in the following command must already exist on the Data Domain system.

The following example shows the Oracle RMAN command syntax for configuring a channel in a Windows environment:

```
RMAN> CONFIGURE CHANNEL DEVICE TYPE 'SBT_TAPE' TRACE <trace-level>
PARMS 'BLKSIZE=1048576,
SBT_LIBRARY=<path-to-%RMAN_AGENT_HOME%\bin\libDDobk.dll,
ENV=(STORAGE_UNIT=<storage-unit-on-Data-Domain-system>,
BACKUP_HOST=<hostname-of-Data-Domain-system>,
RMAN_AGENT_HOME=<path-to-%RMAN_AGENT_HOME%>');
```

The configuration that you specify is also displayed in the Media Management Settings field of Oracle Enterprise Manager (OEM).

## Registering each Data Domain system

To connect to a Data Domain system, you must first register the Data Domain system with the Oracle RMAN agent. This registration provides the Oracle RMAN agent with the necessary username and password for connecting to the Data Domain system. The Oracle RMAN agent stores the username and password securely in an encrypted lockbox file and retrieves the values when connecting to a Data Domain system.

When a Data Domain system is registered for the first time, the encrypted lockbox file is created in the `/config` subdirectory of the Oracle home directory. For security purposes, the lockbox file is typically created so that it can only be accessed by the host system that creates it. If another host system tries to access the lockbox, the access is not permitted. This prevents other systems from connecting to the Data Domain system. Typically, these other systems do not have their own Oracle instance and Oracle home directory with RMAN installed and therefore will have their own lockbox file.

In the typical cases when only a single Oracle host system is in use, or when multiple host systems exist and each has its own Oracle instance and its own Oracle home directory, the Data Domain systems should be registered as described in [Registering a standard Data Domain system](#) on page 62. Each client system has its own lockbox file, accessible only to that system.

When a Data Domain system has DD Boost over Fibre Channel, the Data Domain system should be registered as described in [Registering a Data Domain system with DD Boost over Fibre Channel](#) on page 63.

In some cases, such as in an Oracle RAC environment where multiple system nodes share an Oracle home directory and thus share a lockbox file, it is necessary to create a lockbox file that can be accessed by multiple systems. In these cases, follow the instructions in [Registering a Data Domain system for shared lockbox access](#) on page 63.

### Registering a standard Data Domain system

To register a Data Domain system so that RMAN can connect to it for backup operations, use an RMAN command as shown in the following example. Edit the command for your environment. This registration task is required only once for each Data Domain system. Do not include the registration command in each backup script. Perform this registration before you perform any backups.

```

RUN {
  ALLOCATE CHANNEL C1 TYPE 'SBT_TAPE' TRACE 5 PARMS 'BLKSIZE=1048576,
  SBT_LIBRARY=/home/oracle1/rman_agent_home/lib/libddobk.so,
  ENV=(STORAGE_UNIT=rman_cloud, BACKUP_HOST=bu-ddbea-
  lin-16.lss.emc.com,
  RMAN_AGENT_HOME=/home/oracle1/rman_agent_home)';
  send 'set username q_ost password lX3quest servername bu-ddbea-
  lin-16.lss.emc.com';
  RELEASE CHANNEL C1;
}

```

Ensure that you use the correct values in this RMAN command:

- The username is the DD Boost username that you set with the `user add` command and `ddboost set username` command in step 4 in [Enabling DD Boost on a Data Domain system](#) on page 34. The username is case-sensitive, and must match exactly the value set on the Data Domain system.
- The password is the password that you set up with the DD OS CLI `user` command, as referenced in the Notes at the end of the topic [Enabling DD Boost on a Data Domain system](#) on page 34. The password is case-sensitive, and must match exactly the value set on the Data Domain system.
- The servername value is the name of the Data Domain system exactly as specified by the `BACKUP_HOST` value in the `CONFIGURE CHANNEL` command. The servername value is case-sensitive, and the same form of the name must be used in this command as in the `CONFIGURE CHANNEL` command.

For example, if `server1.domain.com` was used when configuring the channel, you must also use `server1.domain.com` here. Do not use `SERVER1.domain.com` or `server1`. An IP address can be used, if that same IP address was used in the `CONFIGURE CHANNEL` command. You can register multiple forms of the server name. For example, you can register both `server1` and `server1.domain.com` and use either of those names in `CONFIGURE CHANNEL` commands.

Now you are ready to perform backups and restores with your Data Domain system by using `'device type sbt_tape'`. If you want all the backups to go to the Data Domain system, configure `SBT_TAPE` as the default device by using the following command:

```
RMAN> CONFIGURE DEFAULT DEVICE TYPE TO 'SBT_TAPE';
```

### Registering a Data Domain system with DD Boost over Fibre Channel

Registering a Data Domain system with DD Boost over Fibre Channel is almost the same as registering a standard Data Domain system, but with one important difference:

```
RUN {
  ALLOCATE CHANNEL C1 TYPE 'SBT_TAPE' TRACE 5 PARMS 'BLKSIZE=1048576,
  SBT_LIBRARY=/home/oracle1/rman_agent_home/lib/libddobk.so,
  ENV=(STORAGE_UNIT=rman_cloud, BACKUP_HOST=bu-dfc-lin-16.lss.emc.com,
  RMAN_AGENT_HOME=/home/oracle1/rman_agent_home)';
  send 'set username q_ost password lX3quest servername bu-dfc-
  lin-16.lss.emc.com';
  RELEASE CHANNEL C1;
}
```

---

#### Note

For the `servername` value, enter the configured Fibre Channel server name.

Specify the configured Fibre Channel server name for the `BACKUP_HOST` parameter.

```
RUN {
  CONFIGURE CHANNEL DEVICE TYPE 'SBT_TAPE' PARMS 'BLKSIZE=1048576,
  SBT_LIBRARY=/home/oracle1/ddboost_home/lib/libddobk.so,
  ENV=(STORAGE_UNIT=<storage-unit>, BACKUP_HOST=<DFC-server>,
  RMAN_AGENT_HOME=/home/oracle1/ddboost_home) ' FORMAT '%d_%U_%t';
}
```

### Registering a Data Domain system for shared lockbox access

If multiple database host systems must use the same lockbox file, as might be the case in an RAC environment, the procedure for creating the lockbox file includes an additional step. This step is required so that host database systems other than the one that created the lockbox can also access the lockbox. Perform this step before any Data Domain systems are registered. After this step is performed, each Data Domain system can be registered.

The `hostnames` command specifies all the hosts that must access the lockbox file. When you run the `hostnames` command to create the lockbox, all the hosts that you specify in the command are given access to the lockbox.

You run the `hostnames` command only once, to create the lockbox. All the hosts that share access to the lockbox must be specified in a single `hostnames` command. It is

not possible to grant additional hosts access to the lockbox later. In that case, you must delete the existing lockbox file and then create a new lockbox file with all the hostnames.

The `hostnames` command format is as follows:

```

RUN {
  ALLOCATE CHANNEL C1 TYPE SBT TAPE PARMS
  'SBT_LIBRARY=<path-to-$RMAN_AGENT_HOME>/lib/libddobk.so,
  ENV=(RMAN_AGENT_HOME=<path-to-$RMAN_AGENT_HOME>)' ;
  send 'hostnames <hostname-1> <hostname-2> ... <hostname-n>';
  RELEASE CHANNEL C1;
}

```

Run this command as the first RMAN command, before the registration command. All the hosts that will access the shared lockbox file must be listed in the *hostname-1... hostname-n* list. Specify each *hostname* value as a fully qualified hostname, for example, `test1.datadomain.com`.

For example, the following RMAN command creates a lockbox that can be shared by three database client host systems:

```

RUN {
  ALLOCATE CHANNEL C1 TYPE SBT TAPE PARMS
  'SBT_LIBRARY=<path-to-$RMAN_AGENT_HOME>/lib/libddobk.so,
  ENV=(RMAN_AGENT_HOME=<path-to-$RMAN_AGENT_HOME>)' ;
  send 'hostnames dbclient1.datadomain.com
  dbclient2.datadomain.com dbclient3.datadomain.com';
  RELEASE CHANNEL C1;
}

```

When the Oracle RMAN agent receives the `hostnames` command, the Oracle RMAN agent checks if a lockbox file already exists. If so, an error is reported, and a message (logged in the RMAN trace file) reports this and instructs the user to delete the lockbox file if preferred and repeat the command. The Oracle RMAN agent does not delete an existing lockbox file.

If no lockbox file exists, the Oracle RMAN agent creates a lockbox file, and all the specified hosts are granted access to the new lockbox. The lockbox can then be used and shared by all these hosts. Next, register each Data Domain system that these hosts will access. You can register additional Data Domain systems at any time after the lockbox is created. Register each Data Domain system as described in [Registering a standard Data Domain system](#) on page 62. Because all the specified hosts have access to the lockbox, you can register the Data Domain systems on any of the listed host systems. Each Data Domain system must be registered only once.

It is not possible to add additional hostnames after the `hostnames` command has been run. The only way to add additional hostnames is to delete the current lockbox file and then re-create the file by running the `hostnames` command again, this time with all the required hostnames. However, the lockbox contents are lost, and all the Data Domain systems must be re-registered.

Users of RMAN versions prior to RMAN 1.2 who want to share lockbox files among multiple client systems must delete the existing lockbox files and create new lockbox files by using RMAN 1.2 and the instructions in this topic. It is not possible to upgrade a lockbox that was created with earlier versions of RMAN so that it can be accessed by multiple client database systems.

---

**Note**

A lockbox file created with RMAN 1.0 or 1.1 continues to work with RMAN 1.2 when accessed from the client database system that created it. It is not necessary to modify the lockbox file.

---

## Performing backups of a clustered environment to Data Domain

The following examples provide the procedures in different lockbox scenarios for Oracle RMAN agent backups and restores with Oracle RAC nodes in a clustered environment.

### Scenario 1 - Lockbox in a shared location

1. Mount an NFS/CIFS path on both of the Oracle RAC nodes to be used for backups and restores.
2. Install the Oracle RMAN agent on the NSF-mounted path.
3. Run the following command to create a shared lockbox that includes the RAC node hostnames in the cluster:

```
send 'hostnames <hostname-1> <hostname-2> ... <hostname-n>';
```

4. Back up the database on node 1.
5. Restore the database backup on node 2.

### Scenario 2 - Using a lockbox copy

1. Install the Oracle RMAN agent on both of the Oracle RAC nodes separately.
2. Create a lockbox, and add both of the node hostnames to the lockbox on node 1.
3. Copy the lockbox files from node 1 to node 2.
4. Back up the database on node 1.
5. Restore the database backup on node 2.

---

**Note**

As an alternative, you can back up the database on node 2 and then restore the backup on node 1.

---

### Scenario 3 - Without a shared lockbox

1. Install the Oracle RMAN agent on both of the Oracle RAC nodes separately.
2. Register the lockbox separately on each node, node 1 and node 2.
3. Back up the database on node 1.
4. Restore the database backup on node 2.

## Data Domain replication

The Data Domain Replicator provides automated encrypted replication for disaster recovery and multisite backup and archive consolidation. The Data Domain Replicator software asynchronously replicates only compressed, deduplicated data over a wide area network (WAN).

You can use one of two methods to replicate the Oracle backups. Oracle RMAN agent does not support the use of both methods together:

- **Managed file replication (MFR)**

MFR can be used to create multiple copies of a backup while the backup is running. As each RMAN file is created on the primary Data Domain system during a backup, network-efficient replication automatically begins to copy the unique and compressed file segments to the remote Data Domain system. 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. [Performing managed file replication \(MFR\)](#) on page 66 provides details about setting up MFR in the Oracle RMAN agent and restoring from a backup copy.

- **MTree replication**

MTree replication can be used to create a copy or copies of a backup independently from the RMAN backup process. You use the Oracle RMAN agent to create the backup on a primary Data Domain system, and the Data Domain administrator performs the MTree replication from the primary system to the secondary system. The Oracle RMAN agent does not initiate or monitor the MTree replication. However, the product can restore from the replicated copy on a secondary Data Domain system. [Restoring from MTree replicated backups](#) on page 69 provides details about the restore process.

The *Data Domain Operating System Administration Guide* provides more information about the license requirement for Data Domain Replicator and the compatibility between the source DD OS and destination DD OS. The guide includes detailed steps to perform and manage the MTree replication.

## Performing managed file replication (MFR)

File replication enables you to make multiple copies of an RMAN backup. With the Oracle RMAN agent, you can make two, three, or four copies of a backup by specifying `COPIES <n>` in the backup command, where `<n>` is 2, 3, or 4.

To perform file replication, use a command similar to the one shown as follows. Edit the script command for your environment, specifying the Data Domain system that is the destination system for the file replication. Specifying `COPIES <n>` requires an installed Replicator license on both the source and destination Data Domain systems. Performing a backup by specifying `COPIES <n>` requires Oracle Enterprise Edition.

For example, specifying `COPIES 2` instructs the Oracle RMAN agent to perform managed file replication (MFR) to the second Data Domain system. The second Data Domain system name is specified as part of the file name of the second copy, as `<destination-system>` in the following command.

The two file names specified in this command must be identical except for the inclusion of `<destination-system>` in the second name. The Oracle RMAN agent requires that the file names of the two copies are identical. The underscore character (`_`) is also required. Omitting the underscore in the file names results in an error, and the backup fails. Any changes made to the file name formats shown must meet both of the following requirements:

- The file name formats include the underscore character.
- The file names are identical, except for `<destination-system>` in the name for the second copy.

The use of Oracle formats such as `%U` and `%c` is not supported, as these Oracle formats include the copy number in the generated file name. As a result, the file name

generated for the second copy is different from the file name generated for the first copy.

```
backup copies 2 database filesperset=1 format '%u_%p',
 '<destination-system>/%u_%p';
```

#### Note

Use the lowercase `u` in `'%u_%p'`. An uppercase `U`, while a valid Oracle format, is invalid for the Oracle RMAN agent and causes the replication to fail.

To use `COPIES 2`, the value of `BACKUP_TAPE_IO_SLAVES` must be `TRUE`, which you can configure as follows:

```
SQL> show parameter backup_tape
```

NAME	TYPE	VALUE
backup_tape_io_slaves	boolean	FALSE

```
SQL> alter system set backup_tape_io_slaves = true scope=spfile;
```

System altered.

bounce the oracle instance( shutdown and startup) the database

```
SQL> show parameter backup_tape
```

NAME	TYPE	VALUE
backup_tape_io_slaves	boolean	TRUE

The only difference from performing a second backup to a different `BACKUP_HOST` is that the data flow for managed file replication is between the two Data Domain systems. A typical backup sends data between the backup server and a Data Domain system.

To make three or four copies, use a backup command with `COPIES 3` or `COPIES 4` instead of `COPIES 2`. The format of the backup command is as follows:

```
backup copies <n> database filesperset=1 format '%u_%p',
 '<destination-system2>/%u_%p', '<destination-system3>/%u_%p',
 '<destination-system4>/%u_%p'
```

Specifying two, three, or four copies creates the additional copies on the second, third, and fourth destination systems, respectively. All the formats must be identical except for the initial Data Domain system name so that all the file copies are created with the same file names. Each Data Domain system specified must have a storage unit with the name specified by the `STORAGE_UNIT` value.

File replication must be performed to a destination storage unit with the same name as the source storage unit. The target Data Domain systems must have a storage unit that has the same name as the storage unit on the source Data Domain system

containing the file being replicated. The file is replicated to the identically named storage unit on the destination Data Domain systems.

File replication using RMAN requires that both the source Data Domain system and the destination Data Domain systems are connected to, and accessible from, the RMAN system that performs the backups. The network configuration must enable backups to the source Data Domain system and to all the destination Data Domain systems. The destination systems must all be accessible from the source system. When more than two copies are made, the destination systems do not need to be accessible to each other, only to the source system. Replication does not succeed in configurations where the destination Data Domain system is only connected to the source Data Domain system through a private network or link.

The following example shows the output of the RMAN `list` command, which displays the destination and the storage unit associated with `COPIES 2`.

```

RMAN> list backupset 52;

List of Backup Sets
=====

BS Key   Type LV Size
-----
52       Full  1.62G
List of Datafiles in backup set 52
File LV Type Ckp SCN      Ckp Time  Name
-----
1       Full  4614599  09-MAY-12 /u02/app/oracle/oradata/Demo/system01.dbf
2       Full  4614599  09-MAY-12 /u02/app/oracle/oradata/Demo/sysaux01.dbf
3       Full  4614599  09-MAY-12 /u02/app/oracle/oradata/Demo/undotbs01.dbf
4       Full  4614599  09-MAY-12 /u02/app/oracle/oradata/Demo/users01.dbf
5       Full  4614599  09-MAY-12 /u02/app/oracle/oradata/Demo/example01.dbf

Backup Set Copy #2 of backup set 52
Device Type Elapsed Time Completion Time Compressed Tag
-----
SBT_TAPE    00:00:34    09-MAY-12      NO          TAG20120509T101338

List of Backup Pieces for backup set 52 Copy #2
BP Key  Pc# Status      Media              Piece Name
-----
78      1  AVAILABLE  Demo              dd670-ostqa06.datadomain.com/20nahpei_1

Backup Set Copy #1 of backup set 52
Device Type Elapsed Time Completion Time Compressed Tag
-----
SBT_TAPE    00:00:34    09-MAY-12      NO          TAG20120509T101338

List of Backup Pieces for backup set 52 Copy #1
BP Key  Pc# Status      Media              Piece Name
-----
77      1  AVAILABLE  Demo              20nahpei_1

RMAN>

```

### Restoring from multiple backup copies

If there are multiple copies of a backup that are generated through MFR, you can perform a restore from any Data Domain system that contains a copy of the backup. Typically, you perform a restore from the source (local) Data Domain system. In this case, specify this Data Domain system as the backup host in the restore command.

If this source Data Domain system is unavailable for any reason, you can restore the backup from any of the other available Data Domain systems. To do this, determine

which Data Domain systems are accessible over the network. If `COPIES 2` was used, there is only one other system that can perform the restore, so it must be available. If `COPIES 3` or `COPIES 4` was used, any currently available system can perform the restore. If more than one system with the required backup is available, select a system based on the following factors:

- Available bandwidth
- Speed of the connection
- Other performance or logistical considerations

Specify the preferred system as the backup host in the restore command.

In some cases, it is possible that a backup file is unavailable on the specified backup host. In such a case, if another available Data Domain system has a copy of the backup, RMAN tries to retrieve the file from the other system. This retrieval attempt occurs although this system was not specified as the backup host.

For example, three copies of a backup were made to `ddr-host-1`, `ddr-host-2`, and `ddr-host-3`. When a restore is required, `ddr-host-1` is down, but both `ddr-host-2` and `ddr-host-3` are available. If you specify `ddr-host-2` as the backup host in the restore command, RMAN retrieves the backup files from `ddr-host-2`. If a file cannot be found on `ddr-host-2`, RMAN tries to retrieve the file from `ddr-host-3`. If that retrieval does not succeed, the restore fails.

If you select `ddr-host-3` as the backup host, RMAN tries to restore the backup files from `ddr-host-3`. If a file cannot be found on `ddr-host-3`, RMAN tries to retrieve the file from `ddr-host-2`.

---

#### Note

If `ddr-host-2` is down but `ddr-host-1` and `ddr-host-3` are available, specifying `ddr-host-3` as the backup host causes the restore to be performed from `ddr-host-3`. However, if a file is not found on `ddr-host-3`, RMAN does not try to find the missing file on `ddr-host-1`.

---

The system specified as the backup host must always be available. If the system specified as the backup host is unavailable, the restore always fails, even if the backup is available from other copies.

## Restoring from MTREE replicated backups

To enable the restore of an MTREE replicated backup, you must register the secondary Data Domain system in the lockbox. Use an RMAN command as shown in the following example, and edit the command for the environment.

---

#### Note

The registration of the secondary Data Domain system is required only once. Do not include the registration command in each restore script. Perform this registration before you perform any restores from the secondary Data Domain system.

---

```
RUN {
  ALLOCATE CHANNEL C1 TYPE 'SBT_TAPE' TRACE 5 PARMS 'BLKSIZE=1048576,
  SBT_LIBRARY=/home/oracle1/rman_agent_home/lib/libddobk.so,
  ENV=(STORAGE_UNIT=rman_su, BACKUP_HOST=bu-ddbealin-17.lss.emc.com,
  RMAN_AGENT_HOME=/home/oracle1/rman_agent_home)';
  send 'set username q_ost password lX3quest servername bu-
  ddbealin-17.lss.emc.com';
  RELEASE CHANNEL C1;
}
```

Ensure that you specify the correct values in the RMAN `send` command for the registration:

- The `username` value is the DD Boost username on the secondary Data Domain system. The username is case-sensitive, and must match exactly the value that is set on the secondary Data Domain system.
- The `password` value is the password that you set for the username. The password is case-sensitive, and must match exactly the value that is set on the secondary Data Domain system.
- The `servername` value is the case-sensitive name of the secondary Data Domain system. Specify the same form of the name in this command as in the RMAN restore script, shown as follows. For example, if you use `server1.domain.com` in the RMAN restore script, you must also use `server1.domain.com` in the `send` command. Do not use `SERVER1.domain.com` or `server1`.

You can specify an IP address if the same IP address is used in the RMAN restore script. You can register multiple forms of the server name. For example, you can register both `server1` and `server1.domain.com`, and use either of these names in the RMAN restore script.

After the registration is complete, you can perform restores from the secondary Data Domain system. Use an RMAN restore script as shown in the following example, and edit the command for the environment:

```
RUN {
ALLOCATE CHANNEL C1 TYPE 'SBT_TAPE' TRACE 5 PARMS 'BLKSIZE=1048576,
SBT_LIBRARY=/home/oracle1/rman_agent_home/lib/libddobk.so,
ENV=(STORAGE_UNIT=rman_su, BACKUP_HOST=bu-ddbealin-17.lss.emc.com,
RMAN_AGENT_HOME=/home/oracle1/rman_agent_home)';
restore database;
RELEASE CHANNEL C1;
}
```

Ensure that you specify the correct values in the RMAN restore script:

- The `STORAGE_UNIT` value is the destination storage unit of the MTree replication on the secondary Data Domain system. The value is case-sensitive, and must match exactly the value that is set on the secondary Data Domain system.
- The `BACKUP_HOST` value is the case-sensitive name of the secondary Data Domain system. Specify the same form of the name in this script as in the RMAN registration command. For example, if you use `server1.domain.com` in the RMAN registration command, you must also use `server1.domain.com` in the restore script. Do not use `SERVER1.domain.com` or `server1`.

You can specify an IP address if the same IP address is used in the RMAN registration command.

---

#### Note

For MTree replication, the destination storage unit name on the secondary Data Domain system can be the same as or different than the source storage unit name on the source Data Domain system.

---

## Specifying subdirectories for Oracle RMAN agent backups

You can specify one or more subdirectories in the pathname that Oracle RMAN agent uses as the location to store the backup pieces created during a backup. You must specify the subdirectories in the `FORMAT` specification in the RMAN backup script. If the specified subdirectories do not yet exist when the backup starts, the Oracle

RMAN agent automatically creates the subdirectories during the backup to the Data Domain system.

---

#### Note

When you perform a restore, crosscheck, or delete operation, you do not need to specify any subdirectories in the RMAN script. The pathname is recorded as part of the backup piece name in the Oracle catalog. When you perform the operation, Oracle provides the required pathname to the Oracle RMAN agent.

You must use the correct subdirectory naming in the FORMAT specification, depending on whether the backup is performed to a local or remote Data Domain system:

- For a backup to a local system, use the following types of FORMAT specification:
  - When the pathname includes a single subdirectory:

```
FORMAT './<directory-name>/%u_%p';
```

- When the pathname includes multiple subdirectories:

```
FORMAT './<directory-name-1>/<directory-name-2>/.../
<directory-name-n>/%u_%p';
```

- For a backup to a remote system, use the following types of FORMAT specification:

- When the pathname includes a single subdirectory:

```
FORMAT './<directory-name>/%u_%p', '<destination-system>/./
<directory-name>/%u_%p';
```

- When the pathname includes multiple subdirectories:

```
FORMAT './<directory-name-1>/<directory-name-2>/.../
<directory-name-n>/%u_%p', '<destination-system>/./<directory-
name-1>/<directory-name-2>/.../<directory-name-n>/%u_%p';
```

---

#### Note

When you specify any subdirectories in a single `BACKUP` command that is used for a database plus archive log backup, you must use the FORMAT specification twice: once for the database backup pieces and once for the archive log backup pieces. The following RMAN script examples include a database plus archive log backup.

The following examples provide sample RMAN scripts with the correct subdirectory naming in the FORMAT specification for local and remote backups.

**Example 1** Specifying a single subdirectory in the FORMAT specification for a local backup

The following sample RMAN script for a local backup includes a FORMAT specification with a single subdirectory in the backup pathname:

**Example 1** Specifying a single subdirectory in the FORMAT specification for a local backup (continued)

```

RUN {
  ALLOCATE CHANNEL CH1 TYPE 'SBT_TAPE' TRACE 5 PARMS
    'BLKSIZE=1048576,
    SBT_LIBRARY=/home/oracle1/rman_agent_home/lib/libddobk.so,
    ENV=(STORAGE_UNIT=rman_cloud,
    BACKUP_HOST=bu-ddbea-lin-16.lss.emc.com,
    RMAN_AGENT_HOME=/home/oracle1/rman_agent_home)';
  BACKUP DATABASE FORMAT './testdir/%u_%p';
  RELEASE CHANNEL CH1;
}

```

**Example 2** Specifying a subdirectory in the FORMAT specification for a local backup of database plus archive log

The following sample RMAN script for a local backup includes a FORMAT specification with a single subdirectory in the backup pathname. Because the BACKUP command includes both the `database` and `archive log` options, you must specify the FORMAT twice in this case:

```

RUN {
  ALLOCATE CHANNEL CH1 TYPE 'SBT_TAPE' TRACE 5 PARMS
    'BLKSIZE=1048576,
    SBT_LIBRARY=/home/oracle1/rman_agent_home/lib/libddobk.so,
    ENV=(STORAGE_UNIT=rman_cloud,
    BACKUP_HOST=bu-ddbea-lin-16.lss.emc.com,
    RMAN_AGENT_HOME=/home/oracle1/rman_agent_home)';
  BACKUP FORMAT './testdir/%u_%p' DATABASE PLUS ARCHIVELOG DELETE
  INPUT FORMAT './testdir/%u_%p';
  RELEASE CHANNEL CH1;
}

```

**Example 3** Specifying multiple subdirectories in the FORMAT specification for a local backup

The following sample RMAN script for a local backup includes a FORMAT specification with multiple subdirectories in the backup pathname:

```

RUN {
  ALLOCATE CHANNEL CH1 TYPE 'SBT_TAPE' TRACE 5 PARMS
    'BLKSIZE=1048576,
    SBT_LIBRARY=/home/oracle1/rman_agent_home/lib/libddobk.so,
    ENV=(STORAGE_UNIT=rman_cloud,
    BACKUP_HOST=bu-ddbea-lin-16.lss.emc.com,
    RMAN_AGENT_HOME=/home/oracle1/rman_agent_home)';
  BACKUP TABLESPACE PSAPCERUSR FORMAT './testdir1/testdir2/.../
  testdir<n>/%u_%p';
  RELEASE CHANNEL CH1;
}

```

**Example 4** Specifying a single subdirectory in the FORMAT specification for a remote backup

**Example 4** Specifying a single subdirectory in the FORMAT specification for a remote backup (continued)

The following sample RMAN script for a backup to a remote Data Domain system includes a FORMAT specification with a single subdirectory in the backup pathname. The two file names specified in the `BACKUP COPIES 2` command must be identical except for the inclusion of `<destination-system>` in the second name. The underscore character (`_`) is also required:

```

RUN {
  ALLOCATE CHANNEL CH1 TYPE 'SBT_TAPE' TRACE 5 PARMS
    'BLKSIZE=1048576,
    SBT_LIBRARY=/home/oracle1/rman_agent_home/lib/libddobk.so,
    ENV=(STORAGE_UNIT=rman_cloud,
    BACKUP_HOST=bu-ddbea-lin-16.lss.emc.com,
    RMAN_AGENT_HOME=/home/oracle1/rman_agent_home)';
  BACKUP COPIES 2 DATABASE FORMAT './testdir1/%u_%p', '<destination-
system>./testdir1/%u_%p';
  RELEASE CHANNEL CH1;
}

```

**Example 5** Specifying multiple subdirectories in the FORMAT specification for a remote backup

The following sample RMAN script for a backup to a remote Data Domain system includes a FORMAT specification with multiple subdirectories in the backup pathname. The two file names specified in the `BACKUP COPIES 2` command must be identical except for the inclusion of `<destination-system>` in the second name. The underscore character (`_`) is also required:

```

RUN {
  ALLOCATE CHANNEL CH1 TYPE 'SBT_TAPE' TRACE 5 PARMS
    'BLKSIZE=1048576,
    SBT_LIBRARY=/home/oracle1/rman_agent_home/lib/libddobk.so,
    ENV=(STORAGE_UNIT=rman_cloud,
    BACKUP_HOST=bu-ddbea-lin-16.lss.emc.com,
    RMAN_AGENT_HOME=/home/oracle1/rman_agent_home)';
  BACKUP COPIES 2 DATABASE FORMAT './testdir1/testdir2/.../
testdir<n>/%u_%p', '<destination-system>./testdir1/testdir2/.../
testdir<n>/%u_%p';
  RELEASE CHANNEL CH1;
}

```

## Performing backups and restores of Oracle CDBs and PDBs

Oracle 12c introduced support for a multitenant database that allows a single container database (CDB) to include multiple user-created pluggable databases (PDBs). You can perform backups and restores of a whole CDB or individual PDBs. No additional configuration steps are needed in the Oracle RMAN agent to support the environment.

It is recommended that you perform regular backups of the whole CDB to ensure that all PDBs and archived logs are backed up. You can recover the whole CDB or individual

PDBs to any point-in-time from a CDB backup. The Oracle documentation provides details.

## Performing DD Boost compressed restores

Oracle RMAN agent 4.5 introduced support of DD Boost compressed restores. A compressed restore uses less bandwidth by restoring the backed-up data in a compressed form from the Data Domain system to the application host.

A compressed restore can be beneficial in a constrained bandwidth environment or a cloud environment, but can impact the restore performance due to the usage of compression resources on the Data Domain system and application host. By default, the compression of restores is disabled.

Configure a compressed restore of an Oracle RMAN agent backup by setting the `DDBOOST_COMPRESSED_RESTORE` parameter to the value `TRUE` in the `SBT_PARMS` or `ENV` section of the RMAN restore script.

For example, the following RMAN restore script includes the parameter setting:

```
RUN {
  ALLOCATE CHANNEL CH1 TYPE 'SBT_TAPE' TRACE 5 PARMS
  'BLKSIZE=1048576,
  SBT_LIBRARY=C:\Users\Administrator\AppData\Local\RMANAgent\bin
  \libDDobk.dll,
  ENV=(STORAGE_UNIT=ddbea,
  BACKUP_HOST=bu-dbe4500.lss.emc.com,
  RMAN_AGENT_HOME=C:\Users\Administrator\AppData\Local\RMANAgent,
  DDBOOST_COMPRESSED_RESTORE=TRUE)';
  RESTORE DATABASE;
  RECOVER DATABASE;
  RELEASE CHANNEL CH1; }
```

---

### Note

When the `DDBOOST_COMPRESSED_RESTORE` parameter is not set to `TRUE`, the Oracle RMAN agent performs an uncompressed restore of the backed-up data from the Data Domain system.

---

## Display and deletion of backups

When the Oracle catalog records are removed for some backups that were performed with the Oracle RMAN agent but the backups still remain on the Data Domain system, those backups become orphaned backups. You can delete such orphaned backups with the `ddutil` program that is installed with the Oracle RMAN agent software.

---

### Note

The `ddutil` program deletes backup files on the Data Domain system but does not delete any information on the database server. Backups are normally deleted by using the `RMAN DELETE` command, which deletes both the backup on the Data Domain system and the corresponding RMAN catalog record.

The `ddutil` program user must have permissions to the lockbox as described in [Registering each Data Domain system](#) on page 61. The program uses the DD Boost library for operations, and does not require you to mount the storage unit.

---

The `ddutil` program is installed in the `bin` directory, created by the Oracle RMAN agent installer. The program performs deletions of individual backup files or groups of backup files, as specified.

You can perform any of the following operations, based on the `ddutil` command and options that you specify on the command line:

- Display information about backup files at the individual file level, directory level, or storage unit level.
- Delete the backup files at the individual file level or directory level.
- Prune (delete) backup files that were created or modified during a specified time interval.

You must run the `ddutil` command on the command line with the required options. The `-z <hostname>:<storage_unit>` option is the only mandatory option, and specifies the credentials that the Oracle RMAN agent requires for lockbox access during the display, delete, and prune operations.

---

#### Note

You must run the `ddutil` command from the original installation location or set the `RMAN_AGENT_HOME` environment variable to enable the dependent libraries and lockbox to be found.

---

The following topics provide details about how to use the `ddutil` command and options to perform the display, delete, and prune operations.

## Using the `ddutil` command to display and delete backups

An Oracle DBA user can run the `ddutil` command with the appropriate options to perform the following operations:

- Display information about the specified backup files and directories:

```
ddutil [-f] -z <hostname>:<storage_unit> [-b <start_time>] [-e
<end_time>] [-R] [-t] [-D <debug_level>] [<pathname1>
<pathname2>]... [<pathname><n>]]
```

---

#### Note

The default operation of the `ddutil` command is the display of backup information. If you run the `ddutil` command with no options other than the mandatory `-z` option, the command displays all the backup files that are located directly under the storage unit.

---

[Using the `ddutil` command to display backup information](#) on page 78 provides details.

- Delete the specified backup files and directories:

```
ddutil -d -z <hostname>:<storage_unit> [-c] [-x] [-D
<debug_level>] [<pathname1> [<pathname2>]... [<pathname><n>]]
```

---

[Using the `ddutil` command to delete backups](#) on page 79 provides details.

- Prune (delete) the backup files and directories within a save time range:

```
ddutil -k -z <hostname>:<storage_unit> [-b <start_time>] [-e <end_time>] [-c] [-R] [-x] [-D <debug_level>] [<pathname1> [<pathname2>] ... [<pathname><n>]]
```

Using the [ddutil command to prune backups based on age](#) on page 80 provides details.

Options enclosed in brackets ([ ]) are optional. The following table describes the `ddutil` command options.

**Table 8** The `ddutil` command options for backup display and deletion

Option	Description
-b <start_time>	<p>Optional. Specifies the start of a time range, in a date and time format. <a href="#">Date and time format used with the ddutil command options</a> on page 77 provides details about the supported date and time formats for &lt;start_time&gt;.</p> <p>Without this option, the earliest backup time is used by default for the start of the time range.</p> <hr/> <p><b>Note</b></p> <p>The operation terminates with an error if the &lt;start_time&gt; value is later than the &lt;end_time&gt; value that is specified with the -e option.</p>
-c	<p>Optional. Specifies to suppress prompting during a delete or prune operation. With this option, the program proceeds to delete or prune the specified backup files without issuing any user prompts.</p>
-d	<p>Required for a delete operation. Specifies to delete the specified backup files and directories from the Data Domain system.</p>
-D <debug_level>	<p>Optional. Generates debugging information during the operation, where &lt;debug_level&gt; is an integer from 0 to 9. This option is used to troubleshoot operational issues.</p> <p>This option causes the <code>ddutil</code> command to print the debugging information to <code>stderr</code> and to the operational log file. The higher the integer value, the more log information is produced. The operational log file also contains the normal operational messages that are printed to <code>stdout</code>.</p> <p>The operational log file, <code>ddutil_&lt;timestamp&gt;.log</code>, resides in the <code>log</code> subdirectory of the installation location. For example, the operational log file name has the format <code>ddutil_Wed May 17 21:49:08 2017.log</code>.</p>
-e <end_time>	<p>Optional. Specifies the end of a time range, in a date and time format. <a href="#">Date and time format used with the ddutil command options</a> on page 77 provides details about the supported date and time formats for &lt;end_time&gt;.</p> <p>The option setting <code>-e now</code> specifies the current time, which is the default time value.</p> <hr/> <p><b>Note</b></p> <p>The operation terminates with an error if the &lt;end_time&gt; value is earlier than the &lt;start_time&gt; value that is specified with the -b option.</p>

**Table 8** The ddutil command options for backup display and deletion (continued)

Option	Description
-f	Optional. Specifies to display information about the specified backup files and directories that were created or modified within the specified time range on the Data Domain system. This display operation is the default operation of the <code>ddutil</code> command.
-k	Required for a prune operation. Specifies to prune (delete) the specified backup files that were created or modified within the specified time range on the Data Domain system.
-R	Optional. Specifies to perform the operation on the backup files in every subdirectory.
-t	Optional. Specifies to include the cloud tier information in the backup file and directory information that a display operation produces.
-x	Optional. Specifies to delete all the empty directories during a delete or prune operation.
-z <i>&lt;hostname&gt;</i> : <i>&lt;storage_unit&gt;</i>	<p>Mandatory. Specifies the hostname of the Data Domain system, followed by a colon, a forward slash (/), and the storage unit name. For example:</p> <pre>-z dd_host1:/dd_storage_unit1</pre> <hr/> <p><b>Note</b></p> <p>The hostname can be a DFC server name if you must run the <code>ddutil</code> command over Fibre Channel.</p> <hr/>
<i>&lt;pathname1&gt;</i> <i>&lt;pathname2&gt;</i> ... <i>&lt;pathname&gt;</i> < <i>n</i> >	<p>Optional. Specifies the pathnames of the files or directories for display or deletion. For each specified directory, the operation is performed on all the files in the directory. When <code>-R</code> is included, the operation processes all the subdirectories of each directory. Without <code>-R</code>, the operation does not process any subdirectories.</p> <p>The following guidelines apply to the pathnames in a <code>ddutil</code> command:</p> <ul style="list-style-type: none"> <li>• The pathnames must be the last options on the command line. No other types of options, such as <code>-R</code> or <code>-c</code>, can appear after the pathnames.</li> <li>• When no pathnames are specified, the operation occurs only at the level of the storage unit root directory.</li> <li>• When a pathname starts with a forward slash (/), the storage unit name must appear after the slash.</li> <li>• When a pathname does not start with a slash (/), the pathname is assumed to be located directly under the storage unit.</li> </ul> <p>For example, when the storage unit is named <code>su1</code>, the pathnames <code>/su1/dir1</code> and <code>dir1</code> both specify the directory <code>dir1</code> under the storage unit.</p> <p>The wildcard character <code>*</code> is supported in pathnames, and represents any number of unspecified characters.</p>

## Date and time format used with the ddutil command options

You can use specific date and time formats with the options `-b <start_time>` and `-e <end_time>` in the `ddutil` command. The following date and time formats are valid:

- Time of day—A time of day is in the form `hh[:mm[:ss]]` (or `hhmm`) [*meridian*] [*zone*]. If you do not specify a meridian (am or pm), a 24-hour clock is used. You can specify a time of day as just `hh` followed by a *meridian*. If you do not specify a

timezone (for example, GMT), then the current timezone is used, as determined by the second parameter, *now*.

- **Date**—A date is a specific month and day, and possibly a year. The acceptable formats are *mm/dd[yy]* and *month\_name dd[, yy]*. If omitted, the year defaults to the current year. If you specify a year as a number in the range 70 and 99, 1900 is added. If a year is in the range 00 and 30, 2000 is added. The treatment of other years less than 100 is undefined. If a number is not followed by a day or relative time unit, the number is interpreted as a year if a *time\_of\_day*, *month\_name*, and *dd* have already been specified; otherwise, it will be treated as a *time\_of\_day*.
- **Day**—A day of the week can be specified. The current day is used if appropriate. A day can be preceded by a number, indicating which instance of that day is preferred; the default is 1. Negative numbers indicate times past. Some symbolic numbers are accepted: last, next, and the ordinals first through twelfth (second is ambiguous, and is not accepted as an ordinal number). The symbolic number next is equivalent to 2; thus, next monday does not refer to the coming Monday, but refers to the one a week later.
- **relative time**—Specifications relative to the current time can be used. The format is [*number*] *unit*. Acceptable units are decade, year, quarter, month, fortnight, week, day, hour, minute, and second.

Most common abbreviations for days, months, and so on are acceptable, including an uppercase or lowercase first letter and three-letter abbreviations, with or without a trailing period. Units, such as weeks, can be specified as singular or plural. Timezone and meridian values can be uppercase or lowercase, with or without periods.

The actual date is formed as follows. First, any absolute date and/or time is processed and converted. Using that time as the base, day-of-week specifications are added. Last, relative specifications are used. If a date or day is specified, and no absolute or relative time is given, midnight is used. Finally, a correction is applied so that the correct hour of the day is produced after allowing for daylight savings time differences.

## Using the `ddutil` command to display backup information

You can run the following `ddutil` command to display information about the backup files and directories:

```
ddutil [-f] -z <hostname>:<storage_unit> [-b <start_time>] [-e
<end_time>] [-R] [-t] [-D <debug_level>] [<pathname1>
[<pathname2>]... [<pathname><n>]]
```

[Table 8](#) on page 76 provides details about the command options.

Typically, you run this operation to prepare for a delete or prune operation with the `ddutil` command. You can compare the displayed backup information with details in the database backup catalog and determine which backup files to delete or prune.

The `ddutil -f <options>` command output includes one row for each backup file and directory that was found, where the row contains the file or directory name and the backup size and time.

The default operation of the `ddutil` command is the display of backup information. If you run the command with no pathnames and no options other than the `-z` option, the

command displays all the backup files that are located directly under the storage unit. The following commands display the same output:

```
ddutil -z <hostname>:<storage_unit>
ddutil -f -z <hostname>:<storage_unit>
```

The `ddutil` command ignores all subdirectories unless you specify the `-R` option. For example, the following command displays all the backup files and directories on the storage unit `su1`, including the files in all the subdirectories:

```
ddutil -f -z host1:/su1 -R
```

Type	Perm	Size	Time	Path
dir	777	489	Wed May 3 12:15:12 2017 (1493828112)	/su1
file	644	682098688	Wed May 3 12:12:47 2017 (1493827967)	/su1/6rs3aer4_1
file	644	3181641728	Wed May 3 12:14:35 2017 (1493828075)	/su1/6ts3aet5_1_1
file	644	775946240	Wed May 3 12:14:53 2017 (1493828093)	/su1/6us3aev7_1_1
file	644	690225152	Wed May 3 12:15:07 2017 (1493828107)	/su1/6vs3aevm_1_1
file	644	18350080	Wed May 3 12:12:56 2017 (1493827976)	/su1/c-4088691966-20170503-00
file	644	18350080	Wed May 3 12:15:12 2017 (1493828112)	/su1/c-4088691966-20170503-01

To display the backup files in specific directories, include the directory pathnames on the command line. For example, the following command displays the backup files in the directories `/su1/dir1` and `/su1/dir2`:

```
ddutil -f -z host1:/su1 dir1 dir2
```

To display the backup files that were created after March 1, use the `-b Mar 1` option. For example:

```
ddutil -f -z host1:/su1 -b Mar 1 -R
```

## Using the `ddutil` command to delete backups

You can run the following `ddutil` command to delete the specified backup files and directories:

```
ddutil -d -z <hostname>:<storage_unit> [-c] [-x] [-D <debug_level>]
[<pathname1> [<pathname2>] ... [<pathname><n>]]
```

[Table 8](#) on page 76 provides details about the command options.

The `-d` option specifies to delete from the Data Domain system any backup files and directories that are specified on the command line. When you specify a directory, the command also deletes any subdirectories in that directory.

By default, the `ddutil -d <options>` command prompts you before deleting any files or directories. The `-c` option specifies to perform the deletions without prompting.

When the `-x` option is specified, the `ddutil` command also deletes any empty directories that it finds. Without the `-x` option, the directories are not deleted.

**NOTICE**

Use the `ddutil -d <options>` command with caution. The deletion of backup files and directories cannot be undone.

---

To delete specific backup files, run the following command:

```
ddutil -d -z <hostame>:/<storage_unit> <file_pathname1>
<file_pathname2> ...
```

To delete specific backup directories, run the following command:

```
ddutil -d -x -z <hostame>:/<storage_unit> <directory_pathname1>
<directory_pathname2> ...
```

To delete both backup files and directories, run the following command:

```
ddutil -d -x -z <hostame>:/<storage_unit> <file_pathname1>
<file_pathname2> ... <directory_pathname1> <directory_pathname2> ...
```

## Using the ddutil command to prune backups based on age

Run the following `ddutil` command to prune (delete) the backup files in the specified pathnames that were created (or later modified) during the time range specified by the `-b` and `-e` options:

```
ddutil -k -z <hostname>:<storage_unit> [-b <start_time>] [-e
<end_time>] [-c] [-R] [-x] [-D <debug_level>] [<pathname1>
[<pathname2>]... [<pathname><n>]]
```

[Table 8](#) on page 76 provides details about the command options.

The `-k` option specifies to delete from the Data Domain system the backup files based on their age:

- If a specified pathname is a file, the file is deleted if it meets the time criteria.
- If a specified pathname is a directory, all the files in the directory are deleted if they meet the time criteria.
- If the `-b` and `-e` options are omitted, the files in the specified pathnames are deleted, without considering a time range.

**NOTICE**

Use the `ddutil -k <options>` command with caution. After you delete certain backup files, you might be unable to restore data from a backup. Therefore, improper file deletion with the `ddutil` command can lead to failed restores and data loss.

The deletion of backup files and directories cannot be undone.

---

When the `-x` option is specified, the `ddutil` command also deletes any empty directories that it finds. Without the `-x` option, the directories are not deleted.

By default, the `ddutil` command prompts you before deleting any files. The command prompts for every file to be deleted, but does not prompt for the removal of

empty directories. The `-c` option specifies to perform the deletions without prompting.

The `ddutil` command ignores all subdirectories unless you specify the `-R` option. Without the `-R` option, empty subdirectories are not deleted; only the current directory is deleted if it is empty of all files and subdirectories after the prune operation.

Both of the following example commands prune all the backup files and empty directories that are older than 30 days under `dir1/subdir1` on the storage unit `su1`. The commands also prune the directory `dir1/subdir1` itself:

```
ddutil -k -z host1:/su1 -e "-30 days" -R -x dir1/subdir1
ddutil -k -z host1:/su1 -e "now-30 days" -R -x dir1/subdir1
```

## Configuring the use of Data Domain Cloud Tier for data movement to the cloud

You can configure the Oracle RMAN agent to use the Data Domain Cloud Tier for the movement of backup data to the cloud and the subsequent recall of the backup data from the cloud.

Data Domain (DD) Cloud Tier is a native feature of DD OS 6.0 and later for data movement from the active tier to low-cost, high-capacity object storage in the public, private, or hybrid cloud for long-term retention. The Oracle RMAN agent 4.0 introduced support of the DD Cloud Tier for movement of DD Boost backup data to the cloud, which frees up space on the Data Domain system (active tier).

You must set up a DD Cloud Tier policy, also known as a data movement policy, for each MTree or storage unit that the Oracle RMAN agent uses for data movement to the cloud.

After you have set up the data movement policies, you can configure and perform the following operations:

- Movement of backup data from the Data Domain system to the cloud.
- Recall of backup data from the cloud to the Data Domain system.

The following topics describe how to set up the required DD Cloud Tier policies to enable the data movement to the cloud and how to perform the data recall from the cloud.

### Setting up the data movement to the cloud

The Oracle RMAN agent moves the backup data from the active tier to the cloud according to the DD Cloud Tier policy. To enable the data movement to the cloud, you must set up the required policy for each MTree or storage unit.

DD Cloud Tier provides two types of data movement policy, the application-based policy and the age-based policy. The Oracle RMAN agent supports only the age-based policy. A file is moved from the active tier to the cloud tier based on the date that it was last modified. For data integrity, the entire file is moved at this time. The age-based policy can specify the file age threshold, the age range for the files to be moved, and the destination. For example, files older than 1 month and younger than 3 months can be moved to the cloud.

Set the age-based data movement policy on a per-MTree basis on the Data Domain system by using the Data Domain Data Management console, as described in the *Data*

*Domain Operating System Administration Guide.* For example, the following procedure adds a data movement policy to an MTree.

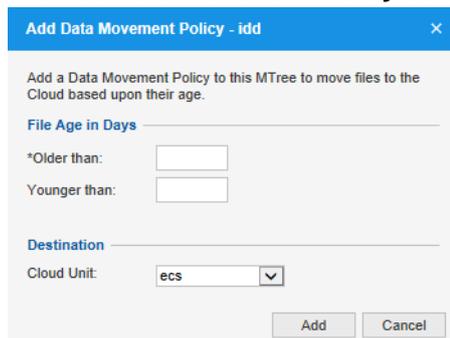
**Note**

Data movement can be initiated manually or automatically by using the policy schedule. Once the policy is set up, the backup files on the Data Domain system are automatically moved to the cloud according to the schedule, and no manual intervention is required. You can also initiate the data movement manually from the Data Domain Data Management console.

**Procedure**

1. Select **Data Management > MTree**.
2. In the top panel, select the MTree to which you want to add a data movement policy.
3. Click the **Summary** tab.
4. Under **Data Movement Policy**, click **Add**.

The **Add Data Movement Policy** dialog box appears:



5. For **File Age in Days**, set the file age threshold (**Older than**) and optionally the age range (**Younger than**).

**Note**

The minimum number of days for **Older than** is 14. Files moved to the cloud tier cannot be accessed directly and must be recalled to the active tier before you can access them. Select the age threshold value as appropriate to minimize or avoid the need to access a file that is moved to the cloud tier.

6. For **Destination**, specify the destination cloud unit.
7. Click **Add**.

## Performing the data recall from the cloud

After the backup data has been moved to the cloud through the DD Cloud Tier policies, the data can be recalled from the cloud automatically or manually. By default, the Oracle RMAN agent automatically recalls data from the cloud to the Data Domain system as required to complete a restore operation. You can also recall the backup files manually, prior to a restore operation.

During a restore operation, if a backup file is requested and not in the active tier, the Oracle RMAN agent automatically recalls the file from the cloud. The Oracle RMAN agent waits for the file to become available before continuing the restore. The time

taken by the recall process depends on the file size and the speed of the network between the cloud and the Data Domain system.

As an alternative, to avoid waiting a long time for the recall of files during a restore, you can recall files manually prior to a restore by running the following RMAN command:

```
RESTORE ... PREVIEW RECALL
```

This command produces the following results:

1. RMAN generates a list of the files that are required for the restore.
2. RMAN sends the list of files to the Oracle RMAN agent.
3. The Oracle RMAN agent recalls all the backup pieces from the list that are in the cloud tier.

If a file from the list is already in the active tier, the Oracle RMAN agent performs no action related to the file.

## RMAN crosscheck and deletion of backup data in the cloud

An RMAN crosscheck or deletion operates on a backup piece file that is created by the Oracle RMAN agent, whether the file exists on the Data Domain system (active tier) or in the cloud tier.

The RMAN `crosscheck` command checks whether a backup piece file exists on the backup media and updates the file information accordingly in the RMAN catalog. For a backup piece file created by the Oracle RMAN agent, the `crosscheck` command reports the file as AVAILABLE when the file exists on the Data Domain system or in the cloud. The `crosscheck` command does not recall a file that exists in the cloud; the file remains in the cloud after the `crosscheck` operation.

An RMAN deletion removes a backup piece file, regardless of the file location. The RMAN deletion removes the file, whether it exists on the Data Domain system or in the cloud. If the backup piece file exists in the cloud, it is removed from the cloud and can no longer be recalled from the cloud.

## Usage limits on Data Domain resources

The Data Domain administrator can use the Data Domain OS commands or the Data Domain Administration GUI to set limits on capacity or streams usage:

- Capacity refers to the amount of hard drive capacity that the Oracle RMAN agent uses on the Data Domain host.

---

### Note

Capacity limits are based on the used logical space. Logical space depends on how much data is written to a storage unit before deduplication. Logical capacity is the size of the uncompressed data. For example, when a 1 GB file is written twice to the same storage unit that was initially empty, then the storage unit has a logical size of 2 GB but a physical size of 1 GB.

- Streams refers to the number of DD Boost streams that the Oracle RMAN agent uses to read data from a storage unit or write data to a storage unit on the Data Domain host.

The Data Domain OS supports soft limits and hard limits on capacity or streams usage:

- When a soft limit is exceeded, the Data Domain host generates an alert, and if a tenant-unit notification list is configured, sends an email to the addresses on the list. An application can continue to use more of the limited resource after a soft limit is exceeded.
- When a hard limit is exceeded, an application cannot use any more of the limited resource.

The Data Domain documentation provides details on the Data Domain versions that support the soft and hard limits for capacity or streams.

The Data Domain administrator can create a separate storage unit for each Oracle RMAN agent host or for a set of hosts that will be limited.

For example, to limit the storage unit capacity used by each Oracle RMAN agent host, where there are 10 Oracle RMAN agent hosts, the Data Domain administrator must create at least 10 storage units. If the Data Domain administrator creates fewer storage units, you must group the Oracle RMAN agent hosts and assign each group of hosts to a storage unit. In this case, you cannot limit the amount of storage that will be consumed by each host. One of the Oracle RMAN agent hosts could consume 100% of a storage unit resource. The resources are consumed on a first come, first served basis.

## Oracle RMAN agent operations with Data Domain usage limits on capacity

At the start of a backup, the Oracle RMAN agent cannot determine how much capacity will be required for the backup. The Oracle RMAN agent always tries to perform a requested backup when the destination device has any available space or storage capacity.

A capacity limit can be set on a storage unit. When the capacity soft limit is exceeded during a backup, alerts appear in the **Current Alerts** pane in the Data Domain Administration GUI.

If the storage unit is part of a tenant-unit with a notification list, the Data Domain host sends an email to the addresses on the list. The Data Domain administrator and the Oracle RMAN agent user should be included in the list.

The backup or restore operation continues without any adverse impact when the capacity soft limit is exceeded. The Oracle RMAN agent does not generate any warning or error message in its log file or operational output.

When the capacity hard limit of a storage unit is exceeded during a backup, the Oracle RMAN agent terminates the backup and generates a message to inform the user about the lack of space on the storage unit.

## Oracle RMAN agent operations with Data Domain usage limits on streams

When the stream soft limit is exceeded during a backup and the storage unit is part of a tenant-unit with a notification list, the Data Domain host sends an email to the addresses on the list. The Data Domain administrator and the Oracle RMAN agent user should be included in the list.

Alerts appear in the **Current Alerts** pane in the Data Domain Administration GUI when the soft limit is exceeded, whether or not the storage unit is part of a tenant-unit.

The backup or restore operation continues without any adverse impact when the stream soft limit is exceeded. The Oracle RMAN agent does not generate any warning or error message in its log file or operational output.

When the stream hard limit is exceeded during an operation, the Oracle RMAN agent terminates the operation with a message to inform the user that a higher streams limit is required. The method for displaying and logging messages is Oracle-specific.

## Configuring usage limits on Data Domain resources

You must complete the required procedures on the Data Domain host to configure the capacity or streams usage limits for the Oracle RMAN agent. The following topics provide the configuration details. Additional following topics provide guidelines and best practices related to the capacity or streams usage limits on Oracle systems.

### Configuring usage quota on Data Domain capacity

To configure a capacity usage quota for the Oracle RMAN agent, the Data Domain administrator must set the hard capacity limit for the storage unit that the Oracle RMAN agent uses for backups.

#### Procedure

1. Determine which Oracle RMAN agent hosts will use the storage unit.
2. Determine how much capacity to allow for the storage unit.
3. Create the storage unit, and then set the capacity quota in the GUI or the CLI command. The Data Domain documentation provides more details.
4. Provide the Oracle RMAN agent users with the DD hostname, storage unit name, username, and password of the storage unit to be used for backups.

The Data Domain administrator can also set the soft capacity quota for the storage unit, which triggers alerts and notifications but does not limit the capacity usage.

The Data Domain administrator can use the Data Domain OS commands or the Data Domain Administration GUI to add or modify the capacity quota of storage units. The Data Domain documentation provides more details.

#### NOTICE

Use caution when decreasing a capacity quota. When a storage unit is almost full and the capacity quota is decreased, the next backup might fail. Notify the Oracle RMAN agent users when a capacity quota is decreased so that the users can evaluate the potential impact on backups.

### Configuring usage limits on Data Domain streams

A storage unit can have soft and hard limits for streams. Soft limits can be set both individually for read, write, and replication streams, and collectively for the total number of all types of streams. A hard limit can be set only for the total number of all types of streams.

To configure a streams usage limit for a storage unit, the Data Domain administrator must set the hard limit for the storage unit that the Oracle RMAN agent uses for backups.

#### Procedure

1. Determine which Oracle RMAN agent hosts will use the storage unit.
2. Determine how many backup and restore streams to allow for the storage unit.
3. Create the storage unit. You can set the streams limit as part of the `ddboost storage-unit create` command or (after the storage unit is created) with

the `ddboost storage-unit modify` command. The Data Domain documentation provides more details.

---

**Note**

A streams limit cannot be set in the Data Domain Administration GUI.

---

4. Provide the Oracle RMAN agent users with the DD hostname, storage unit name, username, and password of the storage unit to be used for backups.

The Data Domain administrator can also set soft limits for the storage unit, which trigger alerts and notifications but do not limit the number of streams used.

The Data Domain administrator can use the `ddboost storage-unit modify` command to modify the streams limits of storage units. The Data Domain documentation provides more details.

**NOTICE**

Use caution when setting a streams hard limit. Setting the streams limit to a low value can impact the backup and restore performance. Decreasing a streams limit can cause a restore to fail. Notify the Oracle RMAN agent users when a streams limit is decreased so that the users can evaluate the potential impact on backups.

---

## Estimating the Data Domain resource usage on Oracle systems

The following topics provide additional guidelines and best practices related to the Data Domain resource usage for Oracle systems.

### Capacity usage on Oracle systems

If the storage capacity of the Data Domain system is exceeded, the backup operation fails. The Oracle RMAN agent generates the following type of error message in the operational log:

```
SBT-10156 2017-04-24 12:12:18 PM sbtwrite2: ddp_write returned error 5194. Unable to write to a file due to reaching the hard quota limit.
```

### Streams usage on Oracle systems

---

**Note**

The streams usage varies, depending on the number and type of parallel operations that are performed at a given time. This topic provides typical numbers for the streams usage of a single operation. To determine more exact numbers, you must monitor the number of streams that the storage units use over a period of time.

---

The number of streams that RMAN typically uses for backups, restores, and maintenance operations (such as crosscheck) corresponds to the number of used channels plus 1.

**If the Data Domain system runs out of streams during a backup, the RMAN channel fails with the following type of error message in the operational log:**

```
SBT-10156 2017-04-24 12:12:18 PM sbtwrite2: ddp_write returned error
5519. Unable to write to a file because the streams limit was
exceeded.
```

**The backup continues and succeeds because RMAN reassigns the backup job to other channels:**

```
RMAN-03009: failure of backup command on CH5 channel at 04/24/2017
16:38:20 ORA-19502: write error on file "CER4_lsrlpqjo_1_1", block
number 33 (block size=8192)
ORA-27030: skgfwr: sbtwrite2 returned error ORA-19511: Error
received from media manager layer, error text: asdf_output_section1()
failed xdr=0x0xb8183f8: bp=0x0xc19c538: send_len=262144: type=12800:
fhand=0x0xc17e688: wrapper=0x(nil): directp=0x0x7f268ad9e000 (1:4:22)
channel CH5 disabled, job failed on it will be run on another channel
```

**If the Data Domain system runs out of streams during a restore, the RMAN channel fails with the following type of error message in the operational log:**

```
SBT-10156 2017-04-24 12:12:18 PM sbtread2: ddp_read returned error
5519. Unable to read from a file because the streams limit was
exceeded.
```



# CHAPTER 5

## Data Domain System Administration

This chapter contains the following topics:

- [Restricting DD Boost access to specific database servers](#)..... 90
- [File replication version compatibility](#)..... 91
- [Modifying an interface group](#)..... 91

## Restricting DD Boost access to specific database servers

By default, when the DD Boost service is enabled on a Data Domain system, the service is accessible to all database servers. However, you can use the `ddboost access` command to override this default, and restrict access to a limited set of database servers.

The following example commands remove the default access permission for all servers and then add new access permissions for two specific database servers, `databaseserver1.datadomain.com` and `databaseserver2.datadomain.com`:

```
ddboost disable
ddboost access del clients *
ddboost access add clients databaseserver1.datadomain.com
databaseserver2.datadomain.com
ddboost enable
```

These commands establish a set of access controls that enables DD Boost access only to the two specified database servers, `databaseserver1.datadomain.com` and `databaseserver2.datadomain.com`.

If no specific access controls are established, running the `ddboost enable` command configures the default access control, which allows all the hosts to access the DD Boost service. However, if there are any access control entries already established, running the `ddboost enable` command does not modify the access control list, thus enabling the access controls that are already established.

---

### Note

- Ensure that no backup jobs are running to the Data Domain system when you change the access control. Run the `ddboost disable` command to prevent operations while access is changed. Disabling DD Boost disables data access to all the database servers.
  - Specify only a fully qualified domain name, IP address, or resolvable DNS name for the client when you modify the client access control list.
  - If the username is changed or deleted, the change in access rights does not affect any current operations. For example, deleting the current clients from the DD Boost access list by running the `ddboost access del clients` command does not stop a file backup that is already in progress. All operations in progress will continue. The change in access rights does not cause existing operations to fail.
  - Run the `ddboost enable` command to re-enable DD Boost and enable RMAN operations after changing the access, which are now subject to the changed access rights.
  - When access control is changed on a Data Domain system for a connected RMAN application, RMAN is still able to perform according to the access rights established at the first connection. The first connection is when the first RMAN command was invoked by using the Oracle RMAN agent. The next time RMAN is stopped and restarted, the new access rights are used.
-

## File replication version compatibility

The Data Domain policy of upgrade compatibility for file replication is as follows:

- All the maintenance and patch versions within a *family* are backward compatible. A family is identified by the first two digits of the release number, such as 5.1. For example, 5.1.0.0, 5.1.0.2, 5.1.1.0, and 5.1.2.0 are backward compatible.
- File replication is backward compatible across two consecutive release families, such as 5.1 and 5.2.
- File replication over WAN should use the Data Domain `ddboost file-replication` option to adjust for low bandwidth and long latency.

## Modifying an interface group

After the interface group is set up, you can add or delete interfaces from the group. The following example shows how to remove an interface from the configured interface group on the Data Domain system.

### Procedure

1. Ensure that no jobs are active from RMAN to the Data Domain system on the interface that you wish to remove from the group. You can do this from the Data Domain system by checking the status of existing connections in the interface group, by running the following command:

```
# ddboost show connections
```

2. Remove the interface from the group on Data Domain system:

```
# ddboost ifgroup del default 192.168.1.3
```

After this, the interface 192.168.1.3 is released from the group and is no longer used by the DD Boost storage server for any jobs from the database servers.

---

### Note

Removing the interface registered with RMAN makes the Data Domain system inaccessible to the database servers. The configuration of the ifgroup on the Data Domain system is not deleted.

To make any changes to any interface that is added to the interface group on the Data Domain system at the network layer, remove the interface from the group and add it back.

If you make changes by running the `net` command to modify the interfaces, such as enabling an interface that is configured for ifgroup, then run the `ddboost show connections` command to update the load balancing view. This updating enables the ifgroup to use the interface. [show connections](#) on page 108 provides more information.

## Removing the advanced load balancing and link failover configuration

The following example shows how to remove a configured interface group on the Data Domain system.

### Procedure

1. Ensure that no jobs are active from RMAN to the Data Domain system. You can do this from the Data Domain system by checking the status of existing connections in the interface group, by running the following command:

```
# ddbost show connections
```

[show connections](#) on page 108 provides information about sample output.

2. Ensure that there are no pending jobs from any of the connected database servers to this system.
3. Disable the feature on the Data Domain system:

```
# ddbost ifgroup disable default
```

4. Reset the interface group:

```
# ddbost ifgroup default reset
```

### Results

All the interfaces are released from the group.

# CHAPTER 6

## Troubleshooting

This chapter contains the following topics:

- [Troubleshooting overview](#) ..... 94
- [Investigating problems](#) ..... 94
- [Oracle limitations](#) ..... 94
- [Reporting problems](#) ..... 94
- [Data Domain system settings for file replication](#) ..... 95
- [Lockbox troubleshooting information](#) ..... 95
- [Failure to load the media management library](#) ..... 95

## Troubleshooting overview

This chapter provides basic troubleshooting tips that can enable you to resolve issues with the product. For issues that cannot be resolved, contact the contracted support providers.

For more information, see the Data Domain Knowledge Base, which is available at <https://support.emc.com>.

## Investigating problems

When you investigate a problem, be aware that the DD Boost software has components on both a Data Domain system and an RMAN system. The two environments must be compatible. The following troubleshooting considerations apply to both systems:

- Supported configurations
 

Ensure that you have a supported configuration as specified in the online software compatibility guide, which is available at <http://compatibilityguide.emc.com:8080/CompGuideApp/getDataDomainBoostCompGuidePage.do>. A supported configuration can become unsupported if any component changes.
- Authorization failures
 

If you encounter authorization failures, ensure that all the systems have correct access credentials for the other systems.

## Oracle limitations

Many RMAN functions require specific editions of Oracle. For example, using `COPIES 2` or parallelism requires the Enterprise Edition. For a complete list, refer to the Oracle documentation at <http://oracle.su/docs/11g/license.112/e10594/editions.htm.com>.

RMAN has numerous restrictions and limitations on how backups can be performed and what can be backed up. The Oracle RMAN agent can only do backup and restore operations that RMAN supports. For example, control file autobackups are never multiplexed. This means that if `COPIES 2` is used to create a replica of a backup, the control files do not get replicated to the second (destination) Data Domain system. Using the `INCLUDE CURRENT CONTROLFILE` option in the backup statement is one possible solution; this includes a copy of the current control file in a backup piece.

## Reporting problems

When reporting a problem to Technical Support, always include the following information:

- The `sbtio.log` file, which is critical for analysis.
- Output of the RMAN command `show all`.
- All the RMAN commands that you used.

If possible, delete the `sbtio.log` file, re-create the problem in as few steps as possible, then capture the `sbtio.log` file to send with your problem report. Set the TRACE level to 5 with the `CONFIGURE CHANNEL` command when you try to re-create the problem.

RMAN determines where the `sbtio.log` file is located. The `sbtio.log` file is usually located in the directory defined by the `user_dump_dest` initialization variable or in the `$ORACLE_HOME/rdbms/log` directory. If all else fails, go to the `$ORACLE_HOME` directory and run a `find` command to locate the file.

## Data Domain system settings for file replication

For all DD OS versions, the `replication throttle` command controls replication. Setting the throttle too low can lead to file replication problems. The *Data Domain Operating System Administration Guide* provides information about using the `replication throttle` command to display or control the throttle settings.

## Lockbox troubleshooting information

The `ddboost.config` file is the configuration file for Oracle RMAN agent. The file contains the saved username/password pairs for the Data Domain systems that you use. If you connect to system `ddr-system` as user `ddr-user` with password `user-password`, there is a corresponding entry in the file. The entry specifies that connections to `ddr-system` are made by using username `ddr-user` with password `user-password`. The configuration file is encrypted and maintained with the RSA lockbox libraries, which are installed as part of the Oracle RMAN agent, and hence some of the error strings returned in `sbtio.log` refer to `lockbox`.

If you change the hostname of the Data Domain system without changing the hostname in all the scripts and so on, then attempts to connect to `ddr-system` fail because there is no host with the specified hostname. If you use the new hostname in all scripts, you can add a new `username-password` entry for that hostname to the `ddboost.config` file. [Registering each Data Domain system](#) on page 61 provides instructions.

The `ddboost.config` file is located in the `$RMAN_AGENT_HOME/config` directory (Linux or UNIX) or in the `%RMAN_AGENT_HOME%\config` directory (Windows).

You cannot reference a lockbox except from the database server where the lockbox was created, unless at the time you created the lockbox, you set it up to be accessible from multiple servers. [Registering a Data Domain system for shared lockbox access](#) on page 63 provides more information.

If you uninstall the boost library, check to see that the old `ddboost.config` files are deleted. There are four files:

- `ddboost.config`
- `ddboost.config.bak`
- `ddboost.config.FCD`
- `ddboost.config.bak.FCD`

Deleting the `ddboost.config` files never causes a problem because the entries will be re-created on demand if you re-run the RMAN commands that register the Data Domain systems.

## Failure to load the media management library

After installation and configuration of the Oracle RMAN agent is completed, attempts to use the Oracle RMAN agent sometimes result in RMAN failures with the RMAN

**error message** Failed to load Media Management Library as shown in the following RMAN error report:

```

RMAN-00571: =====
RMAN-00569: ===ERROR MESSAGE STACK FOLLOWS=====
RMAN-00571: =====
=====
RMAN-03009: failure of allocate command on dd0 channel at 01/07/2013
13:40:36
ORA-19554: error allocating device, device type: SBT_TAPE, device
name:
ORA-27211: Failed to load Media Management Library
    
```

This error usually occurs in an AIX, HP-UX, or Linux environment. The error message indicates that RMAN is unable to locate one of the required libraries, either `libddobk.so` or `libDDBoost.so` or another library that is used by these libraries.

The `ddutil` program can also fail due to an inability to load the required dependent libraries. If this occurs, ensure that the `ddutil` program is running from the installation directory of the Oracle RMAN agent. Alternatively, set the `RMAN_AGENT_HOME` environment variable to enable the dependent libraries to be found.

If you encounter the preceding RMAN error message in the environment, you can use the following steps to help find the source of the problem and correct it.

**Procedure**

1. Check that the `libddobk.so` library pathname was specified correctly in the RMAN configuration command.

[Configuring the software on Linux and UNIX systems](#) on page 56 or [Configuring the software on Windows systems](#) on page 61 provides more information. If the pathname is not correct, the `libddobk.so` library cannot be found. Check that `libddobk.so` exists in the specified directory and that the directory and library file have appropriate access for the user.

2. Check that the home directory of the Oracle RMAN agent was specified correctly in the RMAN configuration command.

[Configuring the software on Linux and UNIX systems](#) on page 56 or [Configuring the software on Windows systems](#) on page 61 provides more information. If the directory is not correct, the `libDDBoost.so` library cannot be found. Check that the `lib` subdirectory exists in the home directory of the Oracle RMAN agent, the `libDDboost.so` library exists in this directory, and the directory and library file have appropriate access for the user.

3. Run the `ldd` command on the `libddobk.so` library to ensure that the libraries that it uses can be found, as shown in the following example:

# CHAPTER 7

## Using DD Boost Commands

This chapter contains the following topics:

- [DD Boost command summary](#) ..... 98
- [The ddbboost command options](#) ..... 98

## DD Boost command summary

The topics in this chapter describe how to use the `ddbboost` command with the following options.

**Table 9** Summary of DD Boost commands

<code>access</code>	Manage the DD Boost access for clients.	<a href="#">access</a> on page 98
<code>clients</code>	Enable the encrypted backups and restores over the LAN, and specify the encryption strength.	<a href="#">clients</a> on page 99
<code>destroy</code>	Delete all the storage units and their contents from the Data Domain system.	<a href="#">destroy</a> on page 99
<code>disable</code>	Disable DD Boost.	<a href="#">disable</a> on page 99
<code>enable</code>	Enable DD Boost.	<a href="#">enable</a> on page 100
<code>fc</code>	Add or modify the DD Boost over Fibre Channel devices.	<a href="#">fc</a> on page 100
<code>file-replication</code>	Set, reset, or display the status of the low-bandwidth optimization or encryption option for file replication. Show or reset statistics, history, and performance.	<a href="#">file-replication</a> on page 100
<code>ifgroup</code>	For advanced load balancing and link failover, manage the interface group (IP addresses) on a Data Domain private network that processes the data transfer.	<a href="#">ifgroup</a> on page 104
<code>option</code>	Set, reset, or display the status of the distributed segment processing feature.	<a href="#">option</a> on page 106
<code>reset</code>	Delete the DD Boost user, and reset the statistics. Clear all the job connection counters shown in the output of the <code>ddbboost show connections</code> command in the event that the network connection is lost.	<a href="#">reset</a> on page 107
<code>set user-name</code>	Set the DD Boost username when DD Boost is enabled.	<a href="#">set user-name</a> on page 108
<code>show</code>	Display the DD Boost connections, DD Boost histogram, or statistics.	<a href="#">show</a> on page 108
<code>status</code>	Show whether DD Boost is enabled or disabled.	<a href="#">status</a> on page 109
<code>storage-unit</code>	Create and delete a storage unit, or list the storage units and files in a storage unit.	<a href="#">storage-unit</a> on page 110

## The `ddbboost` command options

The following topics provide details about the purpose of each `ddbboost` command option.

### **access**

Use this command option to manage the DD Boost access for clients.

**ddboost access add clients**

```
ddboost access add clients client-list
```

Use this command to add clients to the DD Boost access list.

**ddboost access delete clients**

```
ddboost access del clients client-list
```

Use this command to delete clients from the DD Boost access list.

**ddboost access reset**

```
ddboost access reset
```

Use this command to reset the DD Boost client access list to the factory default.

**ddboost access show**

```
ddboost access show
```

Use this command to display the DD Boost client access list.

**clients**

```
ddboost clients add <client-list> [encryption-strength {medium | high}]
```

Administrators use this command to enable encrypted backup and restore over LAN, and specify the encryption strength on a specified client or group of clients.

**destroy**

```
ddboost destroy
```

Administrators use this command to delete all the storage units and their contents from the Data Domain system. The command permanently removes all the data (files) contained in the storage units.

You must also manually remove (expire) the corresponding catalog entries.

**disable**

```
ddboost disable
```

Administrators use this command to disable DD Boost for the Data Domain system.

## enable

```
ddboost enable
```

Administrators use this command to enable DD Boost for the Data Domain system. Whenever the user, user ID (UID), or group ID (GID) changes, the Data Domain system updates all the files and storage units the next time that the `ddboost enable` command is issued.

## fc

The *Data Domain Operating System Command Reference Guide* provides a full list of `ddboost fc` commands.

### fc group add

```
ddboost fc group add groupname initiator "initiator_names"
```

Administrators use this command to grant access to these DD Boost over Fibre Channel devices through an initiator, by adding the initiator to a `ddboost-type scsitarget` access group.

### fc group modify

```
ddboost fc group modify groupname device-set count count
```

Administrators use this command to control the number of DD Boost over Fibre Channel devices advertised to the initiator by configuring the device-set of the `scsitarget` access group. The maximum number of supported DD Boost over Fibre Channel devices per Data Domain system is 64.

## file-replication

### file-replication option reset

```
ddboost file-replication option reset {low-bw-optim | encryption}
```

Administrators use this command to reset the low-bandwidth optimization or encryption to its default value, which is disabled.

### file-replication option set

```
ddboost file-replication option set encryption {enabled | disabled}
ddboost file-replication option set low-bw-optim {enabled |
disabled}
```

This command must be run on both Data Domain systems, the source system and the destination target system. Only an administrator can specify these options.

Encryption must be enabled on both the source and destination systems.

Low-bandwidth optimization, which is disabled by default, is designed for use on networks with less than 6 Mbps aggregate bandwidth. Do not use this option if the maximum file system write performance is required.

After you enable low-bandwidth optimization on both systems, both systems must undergo a full cleaning cycle to prepare the existing data. Run the following command on the source and destination systems:

```
filesystem clean start
```

The amount of time used by the cleaning cycle depends on the amount of data currently on the Data Domain system.

#### file-replication option show

```
ddboost file-replication option show [low-bw-optim | encryption]
```

Use this command to show whether low-bandwidth optimization and/or encryption is enabled or disabled for file replication.

#### file-replication reset stats

```
ddboost file-replication reset stats
```

Administrators use this command to reset the file-replication statistics when DD Boost is enabled.

#### file-replication show active

```
ddboost file-replication show active
```

Use this command to display the status of a DD Boost file-replication transfer to a Data Domain system. Prior to DD OS 5.1, this command only showed the status of DD Boost file-replication transfers outbound from a Data Domain system. With DD OS 5.1, this command also shows the status of DD Boost file-replication transfers inbound to a Data Domain system.

The low-bandwidth optimization status output indicates either that it is enabled and running, or that it is enabled but there is a configuration mismatch.

#### file-replication show detailed-file-history

```
ddboost file-replication show detailed-file-history [duration
duration{day | hr}]
```

Use this command to display the file-replication history over time. The output is organized by file; it is not cumulative data. You must specify a duration.

The data for each file name is organized by date, time, and direction (outbound or inbound). Values are given in KB for pre-compressed, post-filtered, post-low-bw-option, post-compressed, and network data. The last value is an error code; 0 indicates that no error occurred. The following example provides a description of each column:

- Date: 2010/06/16

- Time: 07:48:22
- Direction: Outbound
- File name: /data/coll/x-ddp1/PEOST\_00000001:0000000000:BEOST:4:0::
- Pre-compressed: 3,765,837
- Post-filtered: 30,802
- Post-low-bw-optim: 30,802
- Post-compressed: 17,423
- Network: 65,550
- Error code: 0

**file-replication show detailed-history**

```
ddboost file-replication show detailed-history [duration duration{day | hr}] [interval interval{hr}]
```

Use this command to display the file-replication history over time. A cumulative history is displayed, as opposed to a view by file name. You must specify a duration.

Data is organized by date, time, and direction (outbound or inbound). Cumulative values are given in KB for pre-compressed, post-filtered, post-low-bw-option, post-compressed, and network data. The last value is an error code; 0 indicates no error occurred.

**file-replication show file-history**

```
ddboost file-replication show file-history [duration duration{day | hr}]
```

Use this command to display the data transfer history for inbound and outbound traffic on each file in Data Domain system's `backup` directory for a given duration at a specified interval. The following example shows the inbound and outbound traffic by date and time:

- Amount of pre-compressed data in KB
- Amount of post-compressed data in KB
- Network transfer data in KB
- Low-bandwidth optimization factor
- Number of errors

**Table 10** Inbound and outbound file transfer

Date	Time	Direction	Filename	Pre-Comp (KB)	Post-Comp (KB)	Network (KB)	Low-bw-optim	Errors
2009/08/31	15:13:09	Inbound	/backup/source.tar	47,738	26,634	27,527	1.32	0
2009/08/31	15:13:09	Outbound	/backup/destin.tar	14,327	8,638	9,034	1.21	0

**file-replication show history**

```
ddboost file-replication show history [duration duration{day | hr}]
[interval interval{hr}]
```

Use this command to display the data transfer history between the source and destination (target) Data Domain systems. The following information is shown for inbound and outbound traffic for a given duration at a specified interval:

- Amount of pre-compressed data in KB
- Amount of post-compressed data in KB
- Network transfer data in KB
- Low-bandwidth optimization factor
- Number of errors

**Table 11** Inbound and outbound data transfer

Date	Time	Direction	Pre-Comp (KB)	Post-Comp (KB)	Network (KB)	Low-bw-optim	Errors
2010/08/31	11:14	Inbound	47,738	26,634	27,527	1.32	0
2010/08/31	11:14:09	Outbound	14,327	8,638	9,034	1.21	0

**file-replication show performance**

```
ddboost file-replication show performance [interval sec] [count count]
```

Use this command to display in real time, for a given duration at a specified interval, the amount of pre-compressed outbound and inbound data versus the actual network throughput, or post-compressed data.

```
04/30 11:14:16
      Outbound
Pre-comp   Network   Pre-comp   Network
(KB/s)     (KB/s)     (KB/s)     (KB/s)
-----
          0          0          1644        325
          0          0          2901        470
          0          0          2430        458
          0          0         14390        300
          0          0           0          70
          0          0          2308        446
          0          0          2814        552
          0          0          3010        529
```

**file-replication show stats**

```
ddboost file-replication show stats
```

Use this command to monitor outbound and inbound traffic on a Data Domain system during replication. The compression ratio increases when low-bandwidth optimization is enabled.

```

Direction:                               Outbound
Network bytes sent:                       88,383,976
Pre-compressed bytes sent:                436,874,240
Bytes after filtering:                    153,080,417
Bytes after low-bw-optim:                 145,547,868
Bytes after local compression:            83,628,271
Compression ratio:                        4.9

Direction:                               Inbound
Network bytes received:                   88,383,976
Pre-compressed bytes received:            436,874,240
Bytes after filtering:                    153,080,417
Bytes after low-bw-optim:                 145,547,868
Bytes after local compression:            83,628,271
Compression ratio:                        4.9

```

## ifgroup

The `ifgroup` option specifies an interface group.

### ifgroup add interface

```
ddboost ifgroup add interface group-name {interface ipaddr | client host}
```

Use this command to add an interface, client, or both to the *group-name* or default group. Prior to adding an interface, you must create the *group-name* unless the group name is the default group.

#### Note

The *group-name* “default” is created during an upgrade of a fresh install and is always used if *group-name* is not specified.

Additionally, the IP address must be configured on the Data Domain system and its interface must be enabled. You can add public or private IP addresses for data transfer connections. After adding an IP address as an interface, you must enable advanced load balancing and link failover.

The *Data Domain Operating System Administration Guide* provides more information about interface groups.

### ifgroup create group-name

```
ddboost ifgroup create group-name
```

Use this command to create a group with the name *group-name* for the interface. You can use the command option on a single node and on a master node in a Global Deduplication Array (GDA).

The *group-name* can contain the characters ^, [0-9, a-z, A-Z], \* \$, underscore (\_), and hyphen (-). Hostnames and fully qualified hostnames can be used. You can specify wildcard hostnames, as indicated by an asterisk (\*). Reserved group names that cannot be used are: default, all, none.

**ifgroup del interface**

```
ddboost ifgroup del group-name {interface ipaddr | client host}
```

Use this command to remove an interface, client, or both from the *group-name* or default group. Deleting the last IP address interface disables the ifgroup. If this is the case, you have the option of terminating this command option.

**ifgroup destroy group-name**

```
ddboost ifgroup destroy group-name
```

Use this command to destroy the group name. Only empty groups can be destroyed. Interfaces or clients cannot be destroyed, but can be removed sequentially or by using the command option `ddboost ifgroup reset group-name`.

**Note**

The group-name default cannot be destroyed.

**ifgroup disable group-name**

```
ddboost ifgroup disable group-name
```

Use this command to disable a specific group by entering the *group-name*. If *group-name* is not specified, the command applies to the default group.

**ifgroup enable group-name**

```
ddboost ifgroup enable group-name
```

Use this command to enable the group specified by *group-name*. If *group-name* is not specified, the command enables the default group.

**ddboost ifgroup rename**

```
ddboost ifgroup rename source-group-name destination-group-name
```

Use this command to rename the ifgroup *source-group-name* to *destination-group-name*. This command option does not require disabling of the group.

**ifgroup reset group-name**

```
ddboost ifgroup reset group-name
```

Use this command to reset the group specified by *group-name*. If *group-name* is not specified, the command disables the default group.

**ifgroup show config**

```
ddboost ifgroup show config {interfaces | clients | groups | all}
group-name
```

Use this command to display the selected configuration options. If no selection is made, all the information about the specified *group-name* is displayed.

If *group-name* is not specified, all the information about the default group is displayed. Specify the `all` option to view the configuration options of all the groups. All users can specify this command option.

**ifgroup status group-name**

```
ddboost ifgroup status group-name
```

Use this command to display the status of link aggregation, as enabled or disabled. The status is displayed for all the groups unless *group-name* is specified. All users can specify this command option.

**option****option reset distributed-segment processing**

```
ddboost option reset distributed-segment processing
```

Administrators use this command to reset the distributed segment processing to the default option, which is enabled.

**option set distributed-segment processing**

```
ddboost option set distributed-segment-processing {enabled |
disabled}
```

Administrators use this command to enable or disable the distributed segment processing feature on the DD OS. By default, it is enabled for DD OS 5.1. If this feature presents any problem for a database server, use this command to disable the feature on the Data Domain system.

Distributed segment processing is supported only if the feature is enabled on the Data Domain system.

**Note**

- Distributed segment processing is enabled by default in DD OS 5.1.
- If a system is upgraded from DD OS 4.9 to DD OS 5.0 and the `boost` option was disabled before the upgrade, then the `distributed-segment-processing` option is disabled after the upgrade.

**option show distributed-segment processing**

```
ddboost option show distributed-segment-processing
```

Use this command to show whether the distributed segment processing option is enabled or disabled.

**option reset virtual-synthetics**

```
ddboost option reset virtual-synthetics
```

Use this command to reset the virtual synthetics to the default option, which is disabled. Virtual synthetics is supported on single-node configurations and DD Extended Retention systems only.

**option set virtual-synthetics {enabled | disabled}**

```
ddboost option set virtual-synthetics {enabled | disabled}
```

Use this command to enable or disable the virtual synthetics feature on the DD OS. Virtual synthetics is supported on single-node configurations and DD Extended Retention systems only.

**option show virtual-synthetics**

```
ddboost option show [virtual-synthetics]
```

Use this command to display the status of the virtual synthetics, as enabled or disabled. All users can run this command.

**reset****reset stats**

```
ddboost reset stats
```

Administrators use this command for either of the following purposes:

- To reset all the statistics when DD Boost is enabled.
- As a network recovery procedure to clear job connections after the network connection has been lost.

**reset user-name**

```
ddboost reset user-name username
```

Use this command to delete and reset the DD Boost username.

**Note**

This command is deprecated. It is recommended that you use the `ddboost user unassign` command instead.

## set user-name

```
ddboost set user-name user-name
```

Use this command to set the DD Boost username when DD Boost is enabled.

## show

### show connections

```
ddboost show connections
```

Use this command to display the following information:

- Number of active clients, and for each active client:
  - Client's name.
  - Whether or not the client is idle (yes/no).
  - Number of installed CPUs.
  - Amount of memory in MiB.
  - Version number of Installed Oracle RMAN agent.
  - Name and version of the operating system.
  - Oracle version number.
  - RMAN version number.
- Number of connections that a system uses for DD Boost, and whether these connections are spread across interfaces by using advanced load balancing and link failover. The following sample shows the output for client connections.
- Number of connections used for a given group, even when the interface has been removed from the ifgroup.
- Overview of interfaces available for DD Boost.

### show histogram

```
ddboost show histogram
```

Use this command to display a DD Boost histogram for the Data Domain system.

**Table 12** Output definitions for DD Boost histogram

OPER	Name of the operation request sent from a DD Boost client to the Data Domain system.
mean	Mathematical mean time for completion of the operations, in milliseconds (ms).
std-dev	Standard deviation for the time to complete operations, derived from the mean time, in ms.
<1ms	Number of operations that took less than 1 ms.

**Table 12** Output definitions for DD Boost histogram (continued)

<5ms	Number of operations that took between 1 ms and 5 ms.
<10ms	Number of operations that took between 5 ms and 10 ms.
<100ms	Number of operations that took between 10 ms and 100 ms.
<1s	Number of operations that took between 100 ms and 1 second.
<10s	Number of operations that took between 1 second and 10 seconds.
>10s	Number of operations that took more than 10 seconds.
total	Total time taken for a single operation, in ms.
max	Maximum time taken for a single operation, in ms.
min	Minimum time taken for a single operation, in ms.

**show stats**

```
ddboost show stats [interval seconds] [count count]
```

Use this command to display the output of the previous `show stats` command, the number of bytes written to and read from files contained in storage units, and the number of files created and deleted from storage units. The number of errors encountered for each operation is also displayed. Optionally, you can display the statistics for a specific number of seconds. The count displays the number of lines equal to the count value.

If distributed segment processing is enabled, the number of bytes transferred through distributed segment processing and the amount of compression achieved is displayed.

**show user-name**

```
ddboost show user-name
```

Use this command to display the default DD Boost user if one is configured, or otherwise display that there is no default user.

**Note**

This command is deprecated. It is recommended that you use the `ddboost user show` command instead.

**status**

```
ddboost status
```

Use this command to determine whether DD Boost is enabled or disabled.

## storage-unit

### storage-unit create

```
ddboost storage-unit create storage_unit_name [quota-soft-limit n
{MiB|GiB|TiB|PiB}] [quota-hard-limit n {MiB|GiB|TiB|PiB}]
```

Use this command to create a storage unit and set limits. The `mtree` command description in the *Data Domain Operating System Command Reference Guide* provides details on quota limits. If the quota feature is not enabled, the quota is created but a message appears stating that the feature is disabled and limits are not enforced, as shown in Example 2. The `quota` command description in the *Data Domain Operating System Command Reference Guide* provides details.

Quotas can cause OpenStorage backup applications to report non-intuitive sizes and capacities. The Knowledge Base article 85210, available on the Support portal, provides details.

#### Example 1

```
ddboost storage-unit create SU_1 quota-soft-limit 793 GiB quota-
hard-limit 1078 GiB
```

```
Created storage-unit SU_1
quota-soft-limit: 793 GiB, quota-hard-limit: 1078 GiB
```

#### Example 2

```
ddboost storage-unit create SU_2 quota-soft-limit 1586 GiB
```

```
** Quota is disabled. Quota limits are not enforced.
Created storage-unit SU_2
quota-soft-limit: 1586 GiB, quota-hard-limit: N/A
```

### storage-unit delete

```
ddboost storage-unit delete storage-unit-name
```

Administrators use this command to delete a specified storage unit and all of its contents. You must also manually remove (expire) the corresponding catalog entries.

### storage-unit show

```
ddboost storage-unit show [compression] [storage-unit-name]
```

Use this command to display the names of all the storage units or, optionally, the names of all the files in a specified storage unit.

To display the compression for all the storage units (the original byte size, global compression, and local compression for all storage units), use the `compression` option.

---

**Note**

To interrupt the output of this command, press **Ctrl+C**.

---



# APPENDIX A

## Installation and Configuration Checklist

This appendix contains the following topics:

- [Configure the Data Domain system](#)..... 114
- [Download and install the Oracle RMAN agent](#)..... 114
- [Configure RMAN for the environment](#)..... 114
- [Test the installation by using sbtttest \(optional\)](#) .....115
- [Backup examples](#).....115

## Configure the Data Domain system

[Configuring the Data Domain System](#) on page 33 provides details about the following steps.

### Procedure

1. Add the DD Boost license.
2. Establish the server name, username, and password for the Data Domain system.

---

### Note

The server name, username, and password must exactly match what is used in the `send` command of Step 2 in [Configure RMAN for the environment](#) on page 114.

---

3. Enable DD Boost.
4. Create a logical storage unit, as described in [Creating storage units](#) on page 36.
5. Enable the distributed segment processing.
6. Check the connectivity. (Can I ping the DDR server name from the backup server?) Ensure that the firewall is not blocking ports.

## Download and install the Oracle RMAN agent

[Configuring the Oracle RMAN Agent](#) on page 49 provides details about the following steps.

### Procedure

1. Download the Oracle RMAN agent software.
2. Complete the software installation and setup procedures.

## Configure RMAN for the environment

[Configuring the Oracle RMAN Agent](#) on page 49 provides details about the following steps.

### Procedure

1. Run the command `CONFIGURE CHANNEL DEVICE TYPE 'SBT_TAPE' .`

[Configuring the software on Linux and UNIX systems](#) on page 56 or [Configuring the software on Windows systems](#) on page 61 provide more information. Ensure that the `PARMS` option is set correctly.

2. Run the command `ALLOCATE CHANNEL .`

[Registering each Data Domain system](#) on page 61 provides more information. Set the library, and run the `send` command to set the server name, username, and password. The server name, username, and password must match exactly what is used in Step 2 in [Configure the Data Domain system](#) on page 114.

3. Run the command `CONFIGURE DEFAULT DEVICE TYPE 'SBT_TAPE'`.

[Registering each Data Domain system](#) on page 61 provides more information.

## Test the installation by using sbttest (optional)

Oracle provides the `sbttest` command as a diagnostic tool. Before you run `sbttest` with the Oracle RMAN agent, set the following system environment variables by using the same values as required for the `PARMS` option in step 1 of [Configure RMAN for the environment](#) on page 114:

- `ORACLE_HOME`
- `RMAN_AGENT_HOME`
- `BACKUP_HOST`
- `STORAGE_UNIT`

### Linux example

```
export ORACLE_HOME=/oracle_home
export BACKUP_HOST=MyDDR
export STORAGE_UNIT=MyStorageUnitName
sbttest MyFileName -trace MyTraceFileName -libname
<ddbost_directory>/lib/libddobk.so
```

## Backup examples

```
RMAN> backup copies 2 database format '%u_%p',
'destinationsystem/%u_%p' filesperset=1;

RMAN> backup database format '%u_%p';
```



# APPENDIX B

## Performing a Backup with Oracle Enterprise Manager

This appendix contains the following topics:

- [Example screens](#).....118

# Example screens

Figure 9 Login window

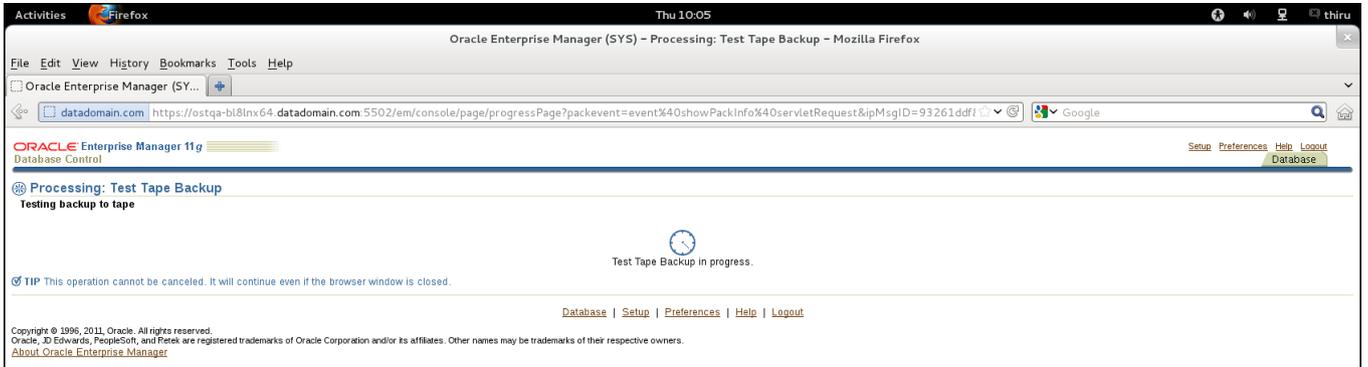


Figure 10 Database instance page



Figure 11 Processing page

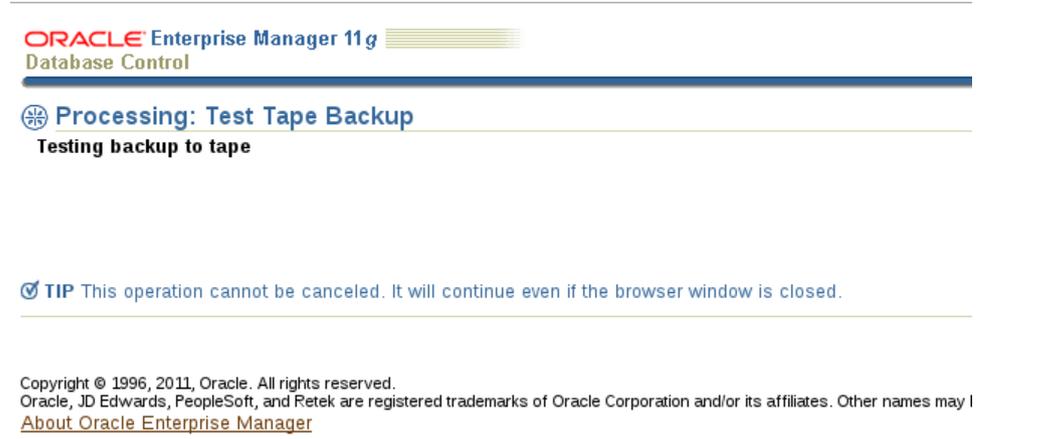


Figure 12 Schedule Customized Backup: Options page

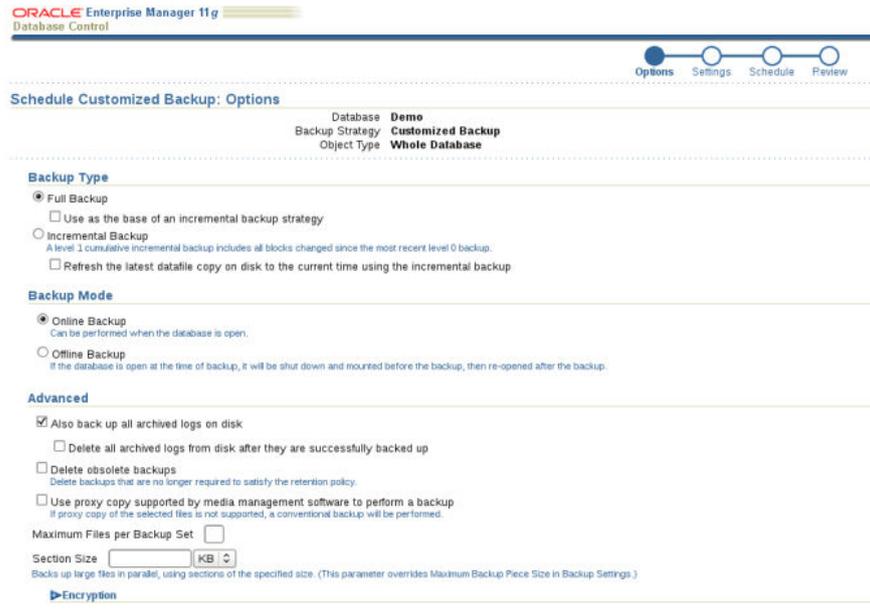


Figure 13 Schedule Customized Backup: Schedule page

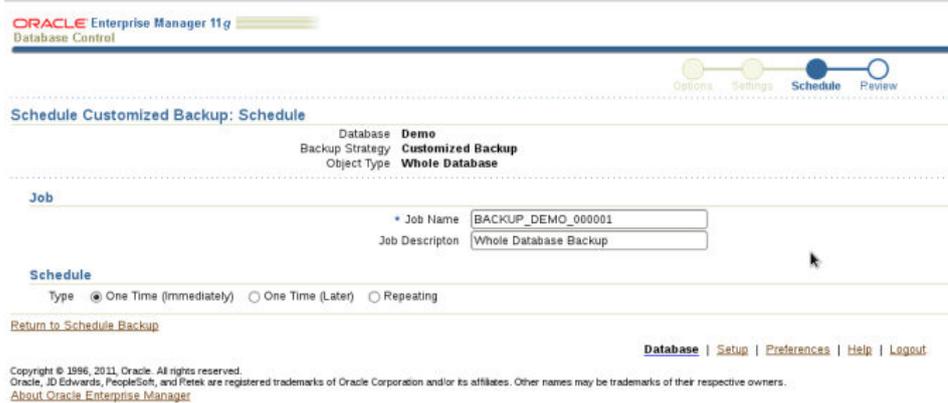
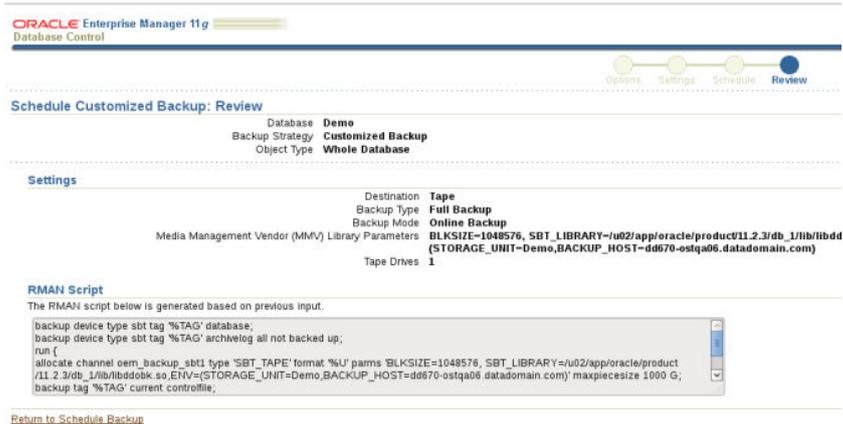


Figure 14 Schedule Customized Backup: Review page



**Figure 15** Execution page

ORACLE Enterprise Manager 11g  
Database Control

Execution: Demo Page Refreshed Apr 26, 2012 4:20:25 PM EDT

---

**Summary**

Status	<b>Succeeded</b>	Type	<b>Database Backup</b>
Scheduled	Apr 26, 2012 4:19:06 PM (UTC-04:00)	Owner	SYS
Started	Apr 26, 2012 4:19:06 PM (UTC-04:00)	Description	Archived Logs Backup
Ended	Apr 26, 2012 4:19:55 PM (UTC-04:00)	Oracle Home	/u02/app/oracle/product/11.2.3/...
Elapsed Time	48 seconds	Oracle SID	Demo
Notification	No	Host Username	oracle
		Database Username	SYS
		Database Role	sysdba
		Backup Strategy	advanced
		Version 10g or higher	YES
		Database Connect String	DESCRIPTION=ADDRESS_LIST=ADDR...
		Database Name	DEMO
		Blackout	NO
		Encryption Mode	None
		Online Backup	NO
		Backup Script	Show

Targets

Status

Expand All   Collapse All	Name	Targets	Status	Started	Ended
▼	Execution: Demo	Demo	Succeeded	Apr 26, 2012 4:19:06 PM (UTC-04:00)	Apr 26, 2012 4:19:55 PM (UTC-04:00)
	Step_PreBackup	Demo	Succeeded	Apr 26, 2012 4:19:24 PM (UTC-04:00)	Apr 26, 2012 4:19:25 PM (UTC-04:00)
	Step_Backup	Demo	Succeeded	Apr 26, 2012 4:19:34 PM (UTC-04:00)	Apr 26, 2012 4:19:50 PM (UTC-04:00)
	Step_PostBackup	Demo	Succeeded	Apr 26, 2012 4:19:55 PM (UTC-04:00)	Apr 26, 2012 4:19:55 PM (UTC-04:00)

# APPENDIX C

## Installing and Configuring the Oracle RMAN Agent

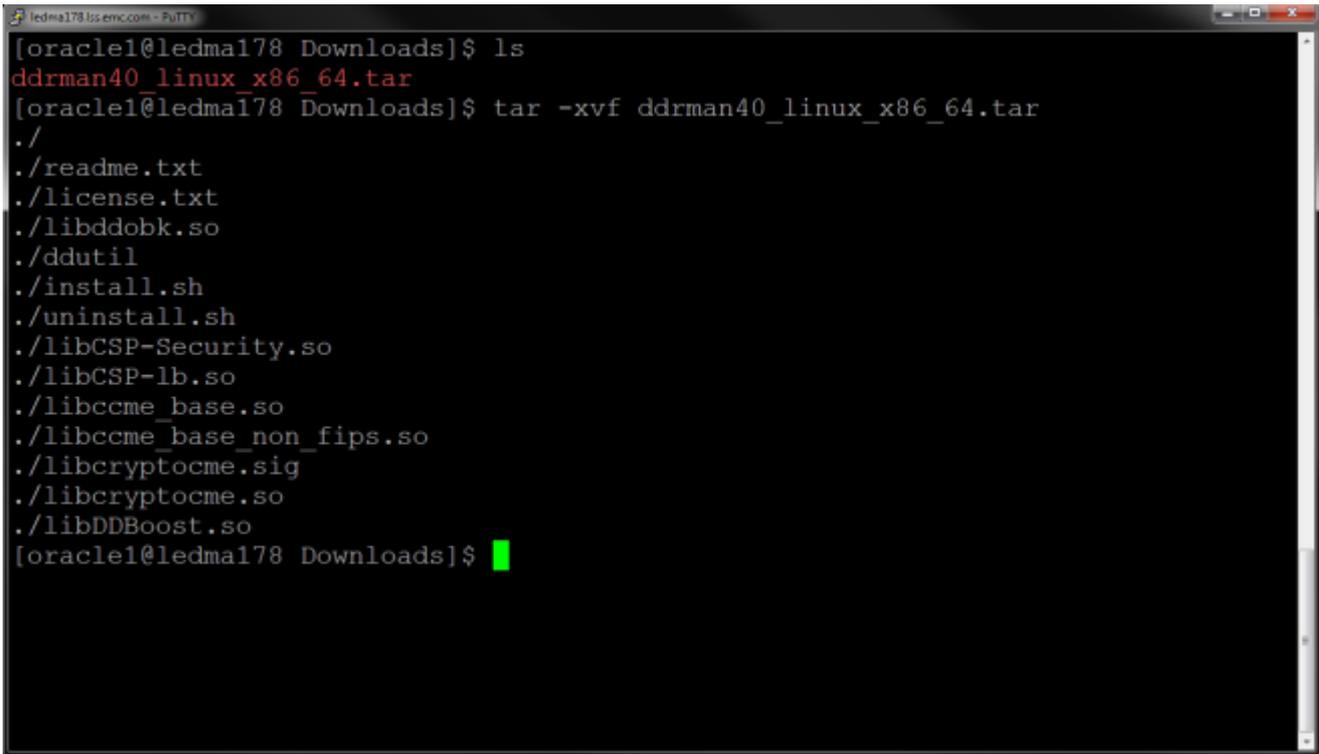
This appendix contains the following topics:

- [Installation and configuration examples](#) ..... 122

## Installation and configuration examples

The following figure shows an example of the software extraction for the Oracle RMAN agent on Linux. You must run the `tar -xvf <file_name>` command to extract the Oracle RMAN agent files from the tar file.

**Figure 16** Oracle RMAN agent package and its contents



```

[oracle1@ledma178 Downloads]$ ls
ddrman40_linux_x86_64.tar
[oracle1@ledma178 Downloads]$ tar -xvf ddrman40_linux_x86_64.tar
./
./readme.txt
./license.txt
./libddobk.so
./ddutil
./install.sh
./uninstall.sh
./libCSP-Security.so
./libCSP-lb.so
./libccme_base.so
./libccme_base_non_fips.so
./libcryptocme.sig
./libcryptocme.so
./libDDBoost.so
[oracle1@ledma178 Downloads]$

```

The following figure shows an example of the installation, in which the `install.sh` script guides the user through the process. The directory `RMAN_AGENT_HOME:/home/oracle1/rman_agent_home` has been selected. The Oracle RMAN agent is being installed for the first time on the system, which has no previous installations of pre-4.5 Oracle RMAN agent software.

Figure 17 Oracle RMAN agent installation

```

[oracle1@ledma178 Downloads]$ ./install.sh
Do you want to install under the default installation directory /home/oracle1/opt/dpsapps/rmanagent? (y or n)
n
A directory outside of the ORACLE_HOME directory must have been created specifically for this installation.
Enter the full pathname of the installation destination:
/home/oracle1/rman_agent_home
The lib directory /home/oracle1/rman_agent_home/lib is created.
The config directory /home/oracle1/rman_agent_home/config is created.
The bin directory /home/oracle1/rman_agent_home/bin is created.
Installing the Oracle RMAN agent.
Copying the lockbox libraries to /home/oracle1/rman_agent_home/lib/.
Copying libddobk.so to /home/oracle1/rman_agent_home/lib/.
Copying libDDBBoost.so to /home/oracle1/rman_agent_home/lib/.
Copying ddutil to /home/oracle1/rman_agent_home/bin/.
Successfully installed the Oracle RMAN agent.
The Oracle RMAN agent library, libddobk.so, does not exist in /home/oracle1/app/oracle1/product/11.2.0/dbhome_1/lib.
Do you want to update settings in /home/oracle1/app/oracle1/product/11.2.0/dbhome_1 directory so that RMAN scripts from previous installation can be reused? (y or n)
n
Installation is completed.
[oracle1@ledma178 Downloads]$ █

```

The following figure shows an example of the Oracle RMAN agent files that were installed:

- The shared library `libddobk.so` and its dependent libraries were installed under the `$RMAN_AGENT_HOME/lib` directory.
- The new `ddutil` binary was installed under the `$RMAN_AGENT_HOME/bin` directory.
- A `config` directory was created for the lockbox.

---

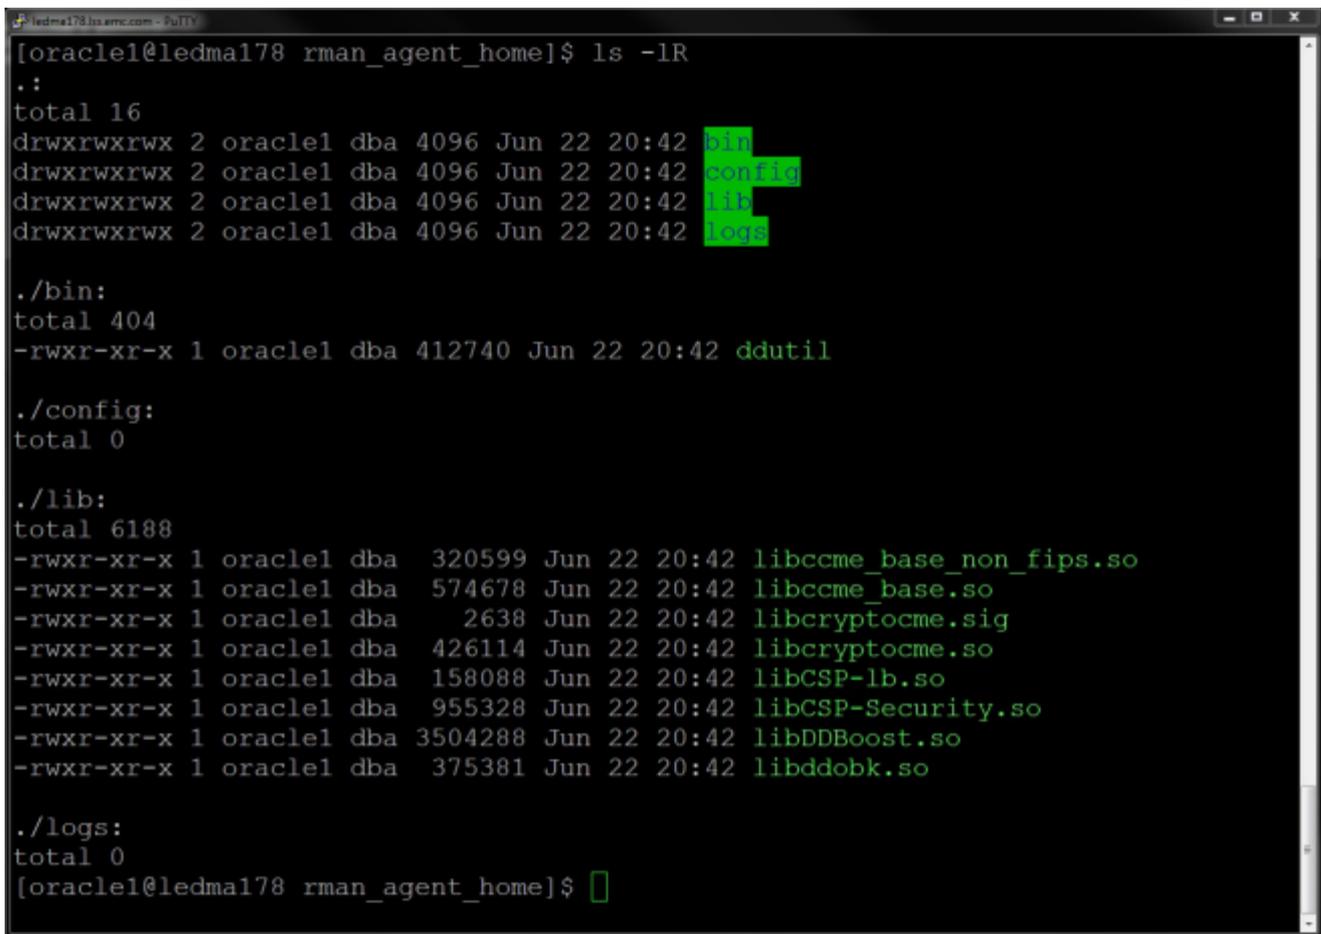
**Note**

Currently, there is no lockbox because this is a new installation.

---

- A `log` directory was created for the `ddutil` program logs.

Figure 18 Oracle RMAN agent installed files

A terminal window showing the output of the command 'ls -lR' in the directory 'rman\_agent\_home'. The output lists the directory structure and file permissions. The root directory contains subdirectories: bin, config, lib, and logs. The bin directory contains the file 'ddutil'. The lib directory contains several shared object files (so) including libccme\_base\_non\_fips.so, libccme\_base.so, libcryptocme.sig, libcryptocme.so, libCSP-lb.so, libCSP-Security.so, libDDBoost.so, and libddobk.so.

```
[oracle1@ledma178 rman_agent_home]$ ls -lR
.:
total 16
drwxrwxrwx 2 oracle1 dba 4096 Jun 22 20:42 bin
drwxrwxrwx 2 oracle1 dba 4096 Jun 22 20:42 config
drwxrwxrwx 2 oracle1 dba 4096 Jun 22 20:42 lib
drwxrwxrwx 2 oracle1 dba 4096 Jun 22 20:42 logs

./bin:
total 404
-rwxr-xr-x 1 oracle1 dba 412740 Jun 22 20:42 ddutil

./config:
total 0

./lib:
total 6188
-rwxr-xr-x 1 oracle1 dba 320599 Jun 22 20:42 libccme_base_non_fips.so
-rwxr-xr-x 1 oracle1 dba 574678 Jun 22 20:42 libccme_base.so
-rwxr-xr-x 1 oracle1 dba 2638 Jun 22 20:42 libcryptocme.sig
-rwxr-xr-x 1 oracle1 dba 426114 Jun 22 20:42 libcryptocme.so
-rwxr-xr-x 1 oracle1 dba 158088 Jun 22 20:42 libCSP-lb.so
-rwxr-xr-x 1 oracle1 dba 955328 Jun 22 20:42 libCSP-Security.so
-rwxr-xr-x 1 oracle1 dba 3504288 Jun 22 20:42 libDDBoost.so
-rwxr-xr-x 1 oracle1 dba 375381 Jun 22 20:42 libddobk.so

./logs:
total 0
[oracle1@ledma178 rman_agent_home]$
```

The following figure shows an example registration of a Data Domain user with the Data Domain server. This example shows the registration of the Data Domain user `q_ost` with password `9_o3s21t!` to enable access to the Data Domain server `bu-ddbea-lin-16.lss.emc.com`.

Figure 19 Registration of the Data Domain user

```

Recovery Manager: Release 11.2.0.1.0 - Production on Thu Jun 22 20:52:26 2017
Copyright (c) 1982, 2009, Oracle and/or its affiliates. All rights reserved.

RMAN> connect target *
2>
3> run {
4> allocate CHANNEL ch1 TYPE 'SBT_TAPE' TRACE 5 PARMS
5> 'BLKSIZE=1048576,
6> SBT_LIBRARY=/home/oracle1/ddboost_home/lib/libddobk.so,
7> ENV=(STORAGE_UNIT=rman_cloud, BACKUP_HOST=bu-ddbea-lin-16.lss.emc.com,
8> RMAN_AGENT_HOME=/home/oracle1/rman_agent_home)';
9> SEND "set username q_ost password 9_03s2!t! servername bu-ddbea-lin-16.lss.emc
.com";
10> }
11>
12>
connected to target database: CER6 (DBID=2273787666)

using target database control file instead of recovery catalog
allocated channel: ch1
channel ch1: SID=150 device type=SBT_TAPE
channel ch1: Data Domain Boost API

sent command to channel: ch1
released channel: ch1

Recovery Manager complete.
[oracle1@ledma178 script]$

```

The following figure shows an example of the lockbox files that were created after the registration.

Figure 20 Oracle RMAN agent lockbox files

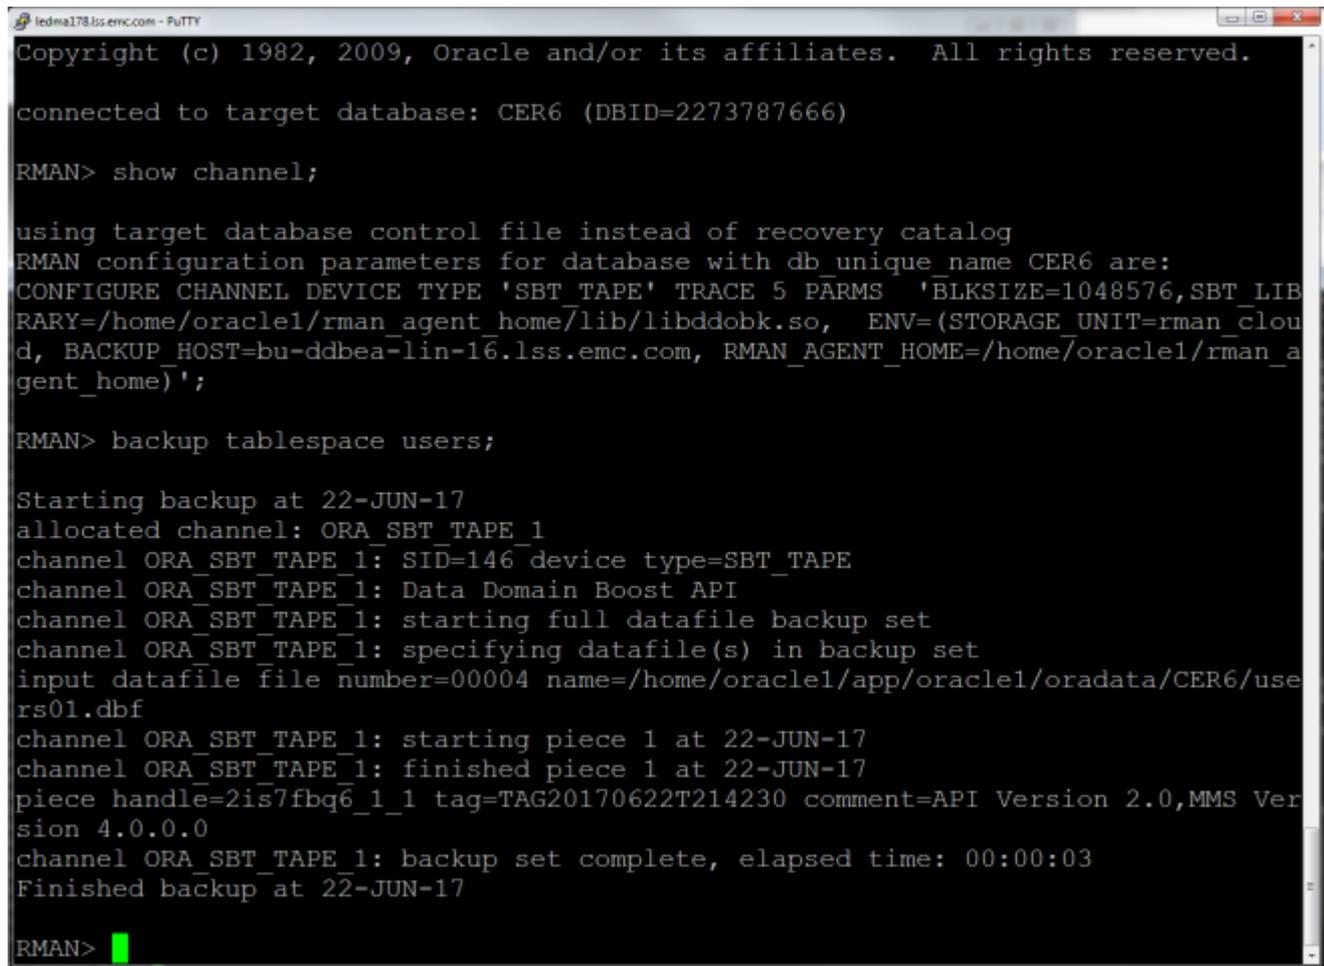
```

[oracle1@ledma178 rman_agent_home]$ ls -l config
total 16
-rw-r--r-- 1 oracle1 dba 3398 Jun 22 20:52 ddbboost.config
-rw-r--r-- 1 oracle1 dba 3114 Jun 22 20:52 ddbboost.config.bak
-rw-r--r-- 1 oracle1 dba 2 Jun 22 20:52 ddbboost.config.bak.FCD
-rw-r--r-- 1 oracle1 dba 4 Jun 22 20:52 ddbboost.config.FCD
[oracle1@ledma178 rman_agent_home]$

```

The following figure shows an example of the channel configuration that is required to run backups and a subsequent backup operation.

Figure 21 Oracle RMAN agent channel configuration and backup



```
Copyright (c) 1982, 2009, Oracle and/or its affiliates. All rights reserved.
connected to target database: CER6 (DBID=2273787666)

RMAN> show channel;

using target database control file instead of recovery catalog
RMAN configuration parameters for database with db_unique_name CER6 are:
CONFIGURE CHANNEL DEVICE TYPE 'SBT_TAPE' TRACE 5 PARMS 'BLKSIZE=1048576,SBT_LIBRARY=/home/oracle1/rman_agent_home/lib/libddobk.so, ENV=(STORAGE_UNIT=rman_cloud, BACKUP_HOST=bu-ddbea-lin-16.lss.emc.com, RMAN_AGENT_HOME=/home/oracle1/rman_agent_home)';

RMAN> backup tablespace users;

Starting backup at 22-JUN-17
allocated channel: ORA_SBT_TAPE_1
channel ORA_SBT_TAPE_1: SID=146 device type=SBT_TAPE
channel ORA_SBT_TAPE_1: Data Domain Boost API
channel ORA_SBT_TAPE_1: starting full datafile backup set
channel ORA_SBT_TAPE_1: specifying datafile(s) in backup set
input datafile file number=00004 name=/home/oracle1/app/oracle1/oradata/CER6/users01.dbf
channel ORA_SBT_TAPE_1: starting piece 1 at 22-JUN-17
channel ORA_SBT_TAPE_1: finished piece 1 at 22-JUN-17
piece handle=2is7fbq6_1_1 tag=TAG20170622T214230 comment=API Version 2.0,MMS Version 4.0.0.0
channel ORA_SBT_TAPE_1: backup set complete, elapsed time: 00:00:03
Finished backup at 22-JUN-17

RMAN>
```