

EMC[®] Data Domain Boost for Oracle Recovery Manager

Version 1.3

Administration Guide

302-001-184

REV. 02

Copyright © 2012-2014 EMC Corporation. All rights reserved. Published in USA.

Published November, 2014

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

The information in this publication is provided as is. EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose. Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

EMC², EMC, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners.

For the most up-to-date regulatory document for your product line, go to EMC Online Support (<https://support.emc.com>).

EMC Corporation
Hopkinton, Massachusetts 01748-9103
1-508-435-1000 In North America 1-866-464-7381
www.EMC.com

CONTENTS

	Preface	7
Chapter 1	Introducing Data Domain Boost for RMAN	9
	Revision History.....	10
	Overview of DD Boost for RMAN.....	10
	Supported Configurations.....	11
Chapter 2	DD Boost Features	13
	Distributed Segment Processing.....	14
	Advanced Load Balancing and Link Failover.....	16
	Configuration Restrictions.....	18
	DD Boost-over-Fibre Channel Transport.....	18
	DD Boost-over-Fibre Channel Path Management.....	20
	Initial Path Selection.....	21
	Client Path Failover.....	22
	Dynamic Re-Balancing.....	22
	Queue-Depth Constraints.....	22
	Encrypted Managed File Replication.....	23
	Low-Bandwidth Optimization.....	24
	Client Access Validation.....	24
	In-flight Encryption.....	24
	Marker Support.....	24
Chapter 3	Configuring the Data Domain System	25
	RMAN Plugin and DD OS Upgrades.....	26
	Firewalls and Ports.....	26
	Enabling DD Boost on a Data Domain System.....	26
	Configuring the DD Boost Server.....	27
	Creating Storage Units.....	27
	Configuring Distributed Segment Processing.....	28
	Configuring Advanced Load Balancing and Link Failover.....	28
	Configuring DD Boost-over-FC Service.....	29
	DD Boost-over-Fibre Channel Path Management.....	30
	Enabling Low-Bandwidth Optimization.....	31
	Enabling Encrypted File Replication.....	32
	Configuring Client Access Validation.....	32
Chapter 4	Configuring DD Boost for RMAN	35
	Install and Configure the RMAN Plug-in AIX, HPUX, Linux or Solaris Environments.....	36
	Installing the RMAN Plugin on Database Servers.....	36
	Installing the RMAN Plugin (Linux or UNIX).....	36
	Uninstalling the UNIX Plugin.....	37
	AIX, HPUX, Linux, or UNIX Environment.....	37
	Install and Configure the RMAN Plug-in in Windows Environments.....	38
	Installing the RMAN Plugin on Database Servers.....	38

	Installing the RMAN Plugin (Windows).....	38
	Uninstalling the Windows Plugin.....	41
	Windows Environment.....	43
	Registering Each Data Domain System.....	44
	Performing File Replication.....	47
	Restoring From Multiple Copies.....	49
Chapter 5	Data Domain System Administration	51
	Restricting DD Boost Access to Specific Database Servers.....	52
	File Replication Version Compatibility.....	52
	Modifying an Interface Group.....	53
	Removing Advanced Load Balancing and Link Failover Configuration	53
	Sizing DD Boost-over FC Device-Set.....	54
Chapter 6	Troubleshooting	57
	About Troubleshooting.....	58
	Investigating Problems.....	58
	Oracle Limitations.....	58
	Reporting Problems.....	58
	Data Domain System Settings for File-Replication.....	59
	Lockbox Troubleshooting Information.....	59
	Failure to Load Media Management Library.....	59
Chapter 7	Using DD Boost Commands	63
	Command Summary.....	64
	ddboost Command Options.....	64
	access.....	64
	clients.....	65
	destroy.....	65
	disable.....	65
	enable.....	65
	fc.....	65
	file-replication.....	66
	ifgroup.....	69
	option.....	70
	reset.....	71
	set user-name.....	71
	show.....	72
	status.....	73
	storage-unit.....	73
Appendix A	Checklists and Notes	75
	Installation and Configuration Checklist.....	76
	Configuring the Data Domain system.....	76
	Downloading and Installing the RMAN Plugin.....	76
	Configuring RMAN for Your Environment.....	76
	Testing the Installation by Using sbttest (Optional)	76
	Backup Examples.....	77
Appendix B	Performing a Backup with Oracle Enterprise Manager	79

	Example Screens.....	80
Appendix C	Installing and Configuring the RMAN Plugin for Multiple SIDs Under One ORACLE_HOME	83
	Installation and Configuration Examples.....	84

CONTENTS

Preface

As part of an effort to improve its product lines, EMC periodically releases revisions of its software and hardware. Therefore, some functions described in this document might not be supported by all versions of the software or hardware currently in use. The product release notes provide the most up-to-date information on product features.

Contact your EMC technical support professional if a product does not function properly or does not function as described in this document.

Note

This document was accurate at publication time. Go to EMC Online Support (<https://support.emc.com>) to ensure that you are using the latest version of this document.

Purpose

This document describes the EMC Data Domain Boost (DD Boost) for Recovery Manager (RMAN) software.

Audience

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

Related EMC documentation

The following EMC publications provide additional information:

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

RMAN Documentation

11g Release 2: From the Oracle Database Documentation Library at 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, 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 used in this document

EMC uses the following conventions for special notices:

NOTICE

Addresses practices not related to personal injury.

Note

Presents information that is important, but not hazard-related.

Typographical conventions

EMC uses the following type style conventions in this document:

Table 1 Typography in EMC Publications

Bold	Use for names of interface elements, such as names of windows, dialog boxes, buttons, fields, tab names, key names, and menu paths (what the user specifically selects or clicks)
<i>Italic</i>	Use for full titles of publications referenced in text
Monospace	Use for: <ul style="list-style-type: none"> • System code • System output, such as an error message or script • Pathnames, filenames, prompts, and syntax • Commands and options
<i>Monospace italic</i>	Use for variables
Monospace bold	Use for user input
[]	Square brackets enclose optional values
	Vertical bar indicates alternate selections - the bar means “or”
{ }	Braces enclose content that the user must specify, such as x or y or z
...	Ellipses indicate nonessential information omitted from the example

Where to get help

EMC support, product, and licensing information can be obtained as follows:

EMC product information

For documentation, release notes, software updates, or information about EMC products, go to EMC Online Support at <https://support.emc.com>.

Technical support

Go to EMC Online Support and click Service Center. You will see several options for contacting EMC Technical Support. Note that to open a service request, you must have a valid support agreement. Contact your EMC sales representative for details about obtaining a valid support agreement or with questions about your account.

Your comments

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications. Send your feedback about this document to DPAD.Doc.Feedback@emc.com.

CHAPTER 1

Introducing Data Domain Boost for RMAN

This chapter contains the following topics:

- [Revision History](#)..... 10
- [Overview of DD Boost for RMAN](#)..... 10
- [Supported Configurations](#)..... 11

Revision History

The following table presents the revision history of this document.

Table 2 Revision History of DD Boost for RMAN Release 1.3

Revision	Date	Description
02 (1.3.0)	January 2015	Added information about DD Boost client access validation and in-flight encryption.
01 (1.3.0)	October 2014	Initial publication.

Overview of DD Boost for RMAN

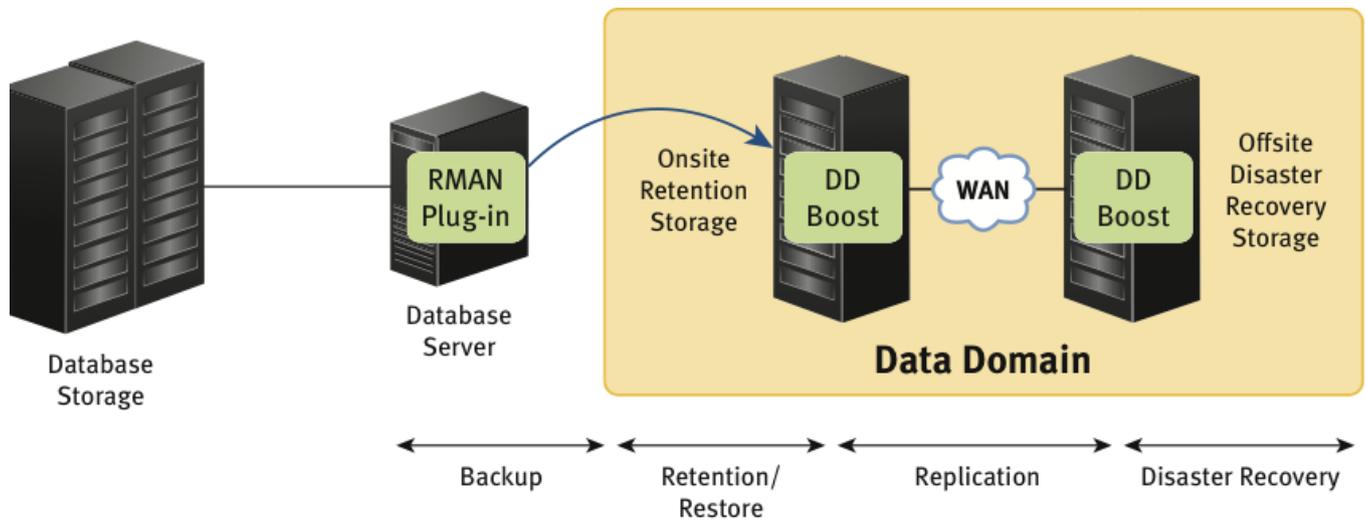
EMC Data Domain Boost (DD Boost) for Oracle Recovery Manager (RMAN) enables database servers to communicate with Data Domain systems in an optimized way, without the need to use a backup application. Use of EMC DD Boost for RMAN improves performance while reducing data transferred over the LAN. In the context of Oracle RMAN, the software has two components:

- An RMAN plugin that you install on each database server. This plugin includes the DD Boost libraries for communicating with the DD Boost server running on the Data Domain system.
- The DD Boost server that runs on 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 of the features of DD Boost, including WAN-efficient replicator software. RMAN manages all 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 Data Domain RMAN plugin, 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 next figure shows an example configuration of Data Domain Boost for RMAN.

Figure 1 DD Boost for RMAN 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 RMAN plugin version must be compatible with your Data Domain system and RMAN configurations. Data Domain does not support combinations other than those detailed in the *EMC Data Domain Boost Compatibility Guide*, which is available from the Data Domain Support portal (<https://support.emc.com>). Navigate to the **Documentation > Compatibility Matrices** page and select **Data Domain Boost Compatibility Guide**.

CHAPTER 2

DD Boost Features

This chapter contains the following topics:

- [Distributed Segment Processing](#) 14
- [Advanced Load Balancing and Link Failover](#) 16
- [DD Boost-over-Fibre Channel Transport](#) 18
- [DD Boost-over-Fibre Channel Path Management](#) 20
- [Encrypted Managed File Replication](#) 23
- [Low-Bandwidth Optimization](#) 24
- [Client Access Validation](#) 24
- [In-flight Encryption](#) 24
- [Marker Support](#) 24

Distributed Segment Processing

Distributed segment processing allows parts of the deduplication process to be performed by the RMAN plugin, which avoids sending 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 RMAN plugin negotiates with the Data Domain system for the current setting of the option and behaves appropriately. Therefore, the RMAN plugin 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:

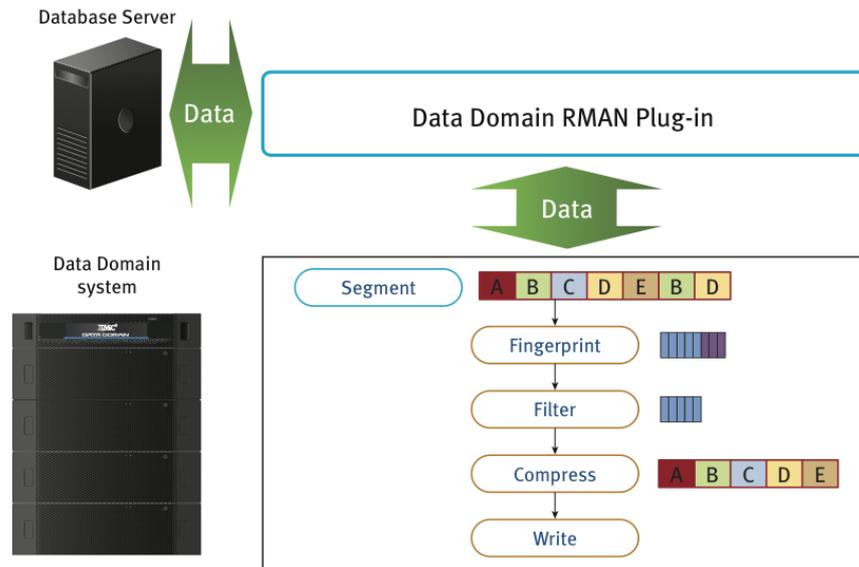
- Potentially higher throughput—because the RMAN plugin sends only unique data to a Data Domain system versus sending all of the data over the LAN. Throughput improvements depend on the redundant nature of the data being backed up, the overall workload on the database server, and the capability of the 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 since 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.

See the *EMC Data Domain Operating System Administration Guide* for more information about local compression and its configuration.

You manage distributed segment processing via the `ddboost option` commands. Data Domain recommends that you use distributed segment processing if your network connection is 1 Gb Ethernet. See [Configuring Distributed Segment Processing on page 28](#).

Distributed Segment Processing Disabled

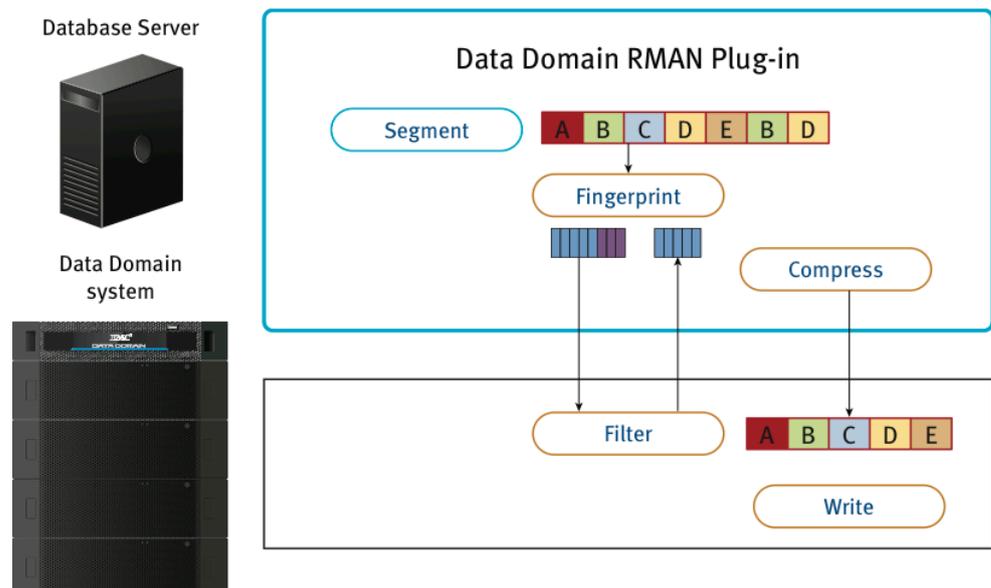
In this mode, the RMAN plugin sends the data directly to the Data Domain system over the LAN. The Data Domain system then segments, deduplicates, and compresses the data before it is written to the disk. A high-level data flow scenario for a single-system Data Domain system with distributed segment processing disabled is shown in the next figure. (Distributed segment processing cannot be disabled on an Extended Retention Data Domain system.) The RMAN plugin runs on the database server.

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 RMAN plugin and the Data Domain system. The data flow for this mode is shown in the next figure.

Parts of the deduplication process are run on the RMAN plugin so that only unique data is sent to a Data Domain system over the LAN. The RMAN plugin 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. A high-level data flow for this scenario on a Data Domain system is shown in the next figure. The RMAN plugin runs on the database server.

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

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 allows for combining multiple Ethernet links into a group. Only one of the interfaces on the Data Domain system is registered with RMAN. The RMAN plugin 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.

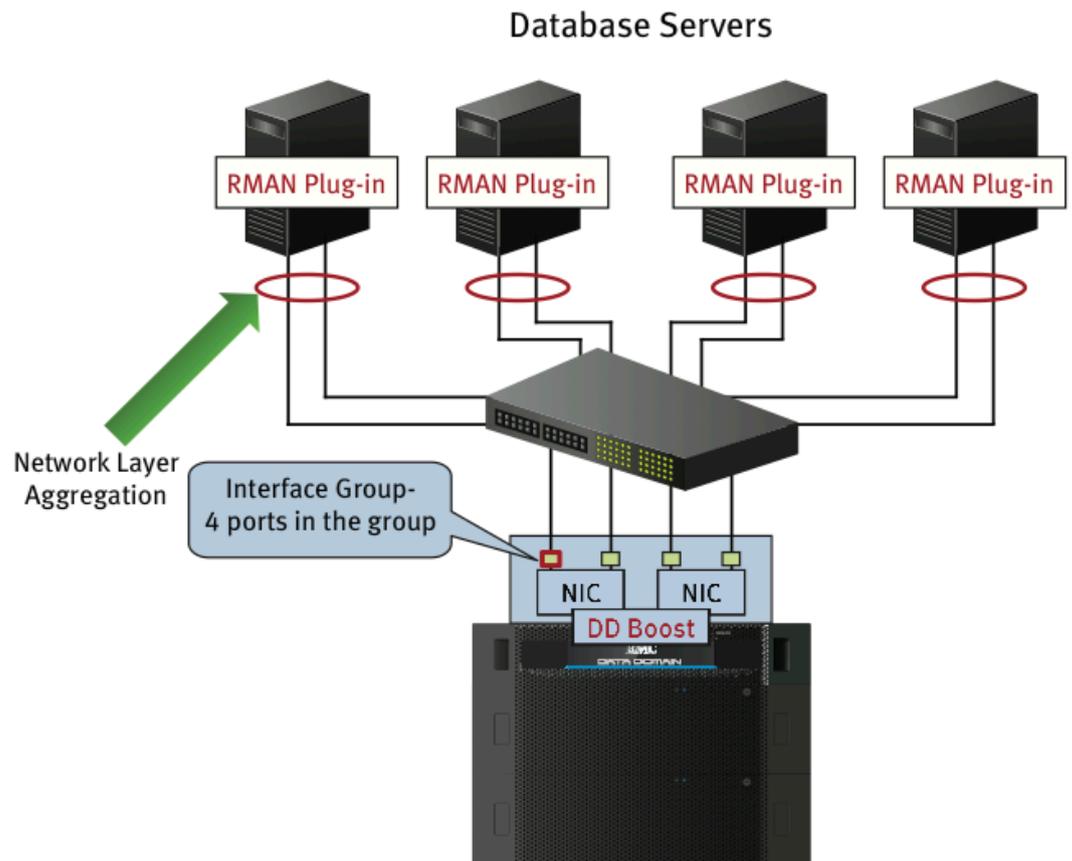
The Data Domain system load balances the connections coming in from multiple database servers on all the interfaces in the group. The load balancing is transparent to RMAN. Because Advanced Load Balancing and Link Failover works at the RMAN plugin software layer, it is seamless to the underlying network 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 RMAN plugin must be able to connect to every interface that is a member of the interface group on the Data Domain system.

The next 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 in conjunction with other network layer aggregation and failover technologies. The sample configuration in the previous 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 in conjunction with other network layer functionality on the Data Domain systems, including VLAN tagging and IP aliasing. This functionality allows additional flexibility in segregating traffic into multiple virtual networks, all of which run over the same physical links on the Data Domain system.

Note

See the *EMC Data Domain Operating System Administration Guide* for more information about how to configure VLAN tagging and IP aliasing on a Data Domain system.

Benefits of Advanced Load Balancing and Link Failover are as follows:

- 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 transparently are failed over to healthy operational links. From the point of view of RMAN, the jobs continue uninterrupted.

For more information, see [Configuring Advanced Load Balancing and Link Failover on page 28](#).

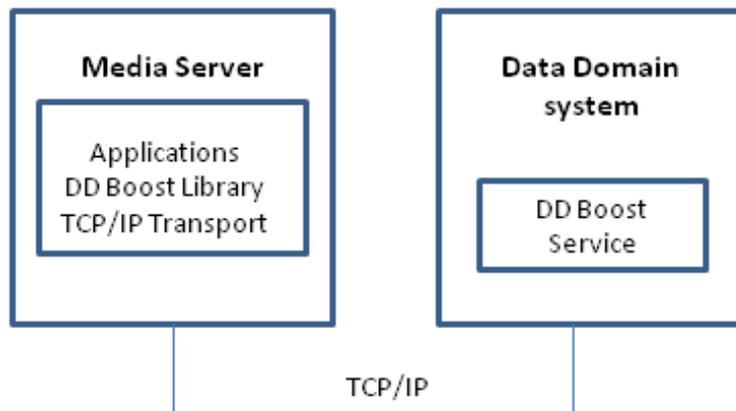
Configuration Restrictions

- Interfaces can be added only to the group using an IP address.
- Even though the Advanced Load Balancing and Link Failover feature works with mixed 1GbE 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 using IP networking. The application specifies the Data Domain system using its hostname or IP address, and the DD Boost Library uses TCP/IP connections to transmit requests to, and receive responses from, the Data Domain system, as shown below.

Figure 5 DD Boost-over-IP Transport



DD OS release 5.3 introduces an alternative transport mechanism for communication between the DD Boost Library and the Data Domain system—Fibre Channel.

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) allows 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 (e.g. Windows, 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.

Under 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).

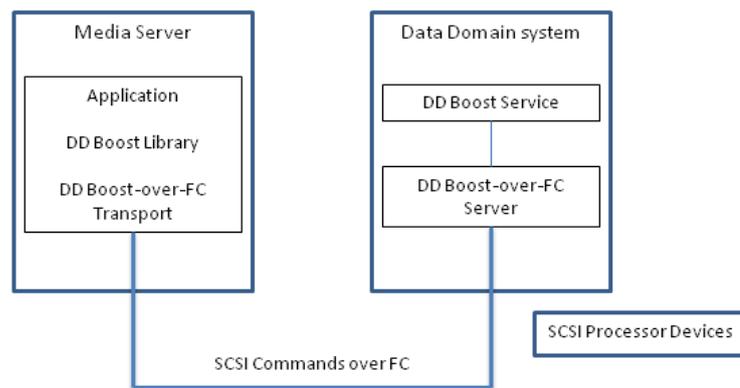
To request access to a Data Domain system using the DD Boost-over-FC transport, the application specifies the Data Domain system using the special string

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.

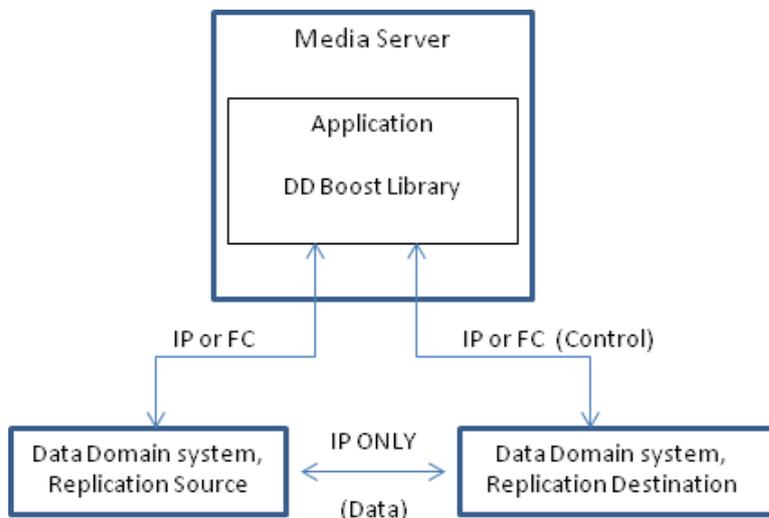
Referencing this catalog, the DD Boost-over-FC Transport logic issues SCSI commands to the identified generic SCSI devices, to transfer DD Boost protocol requests and responses between the library and the Data Domain system.

Figure 6 SCSI Commands between Media Server and Data Domain System



Most DD Boost features are independent of 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 below, 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

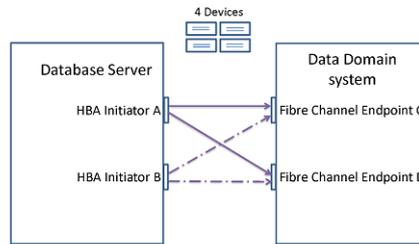
DD Boost-over-Fibre Channel Path Management

The IFGROUP-based mechanism described in [Advanced Load Balancing and Link Failover on page 16](#) 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 Processor-type SCSI devices to the database 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 case where:

- Database 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 containing:
 - 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

In this case, the database 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 desired 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. But 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 8 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)

```

and 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` will fail 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 may 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 IO requests which can be

conducted simultaneously over a given generic SCSI device. For example, 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 (e.g. backup jobs) are trying to use the same generic SCSI device.

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

We use the term “queue-depth” to describe the system-imposed limit on the number of simultaneous SCSI requests on a single device. Client systems (like 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 54](#) for guidance regarding 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 pre-configured DD Boost user name and password. If the encrypted file replication feature is enabled, the session between the source and destination Data Domain systems is encrypted using Secure Sockets Layer (SSL), ensuring that all 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 knowledge of RMAN. Encrypted file replication uses the ADH-AES256-SHA cipher suite. There is no ability to configure a different suite in DD OS.

Encrypted file replication is available to a Data Domain system with an installed Replicator license, and applies to all 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 will fail.

Encrypted file replication can be used in conjunction with the encryption of data-at-rest feature available on 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 using SSL for sending over a WAN.

Note

- Both the source and the destination Data Domain systems must be running 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 in conjunction with each other.

For more information, see [Enabling Encrypted File Replication on page 32](#) and the discussion of this topic in the *EMC Data Domain Operating System Administration Guide*.

Low-Bandwidth Optimization

The low-bandwidth Replicator option reduces the WAN bandwidth utilization. This option is useful if file replication is being 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 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 will participate in managed file replication must have this option enabled.

For more information, see [Enabling Low-Bandwidth Optimization on page 31](#) and the discussion of this topic in the *EMC Data Domain Operating System Administration Guide*.

Client Access Validation

Configuring client access validation for DD Boost limits access to the Data Domain system for DD Boost clients by requiring a one-time DD Boost authentication per connection. Connection authentication against the hostname is needed only until credentials are available. The list of clients can be updated at anytime without a restart requirement, thus eliminating access validation impact on jobs in progress.

In-flight Encryption

In-flight encryption allows applications to encrypt in-flight backup or restore data over 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

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. For details, see the *EMC DD OS Administration Guide*.

CHAPTER 3

Configuring the Data Domain System

This chapter includes:

- [RMAN Plugin and DD OS Upgrades](#)..... 26
- [Firewalls and Ports](#)..... 26
- [Enabling DD Boost on a Data Domain System](#)..... 26
- [Configuring the DD Boost Server](#).....27

RMAN Plugin and DD OS Upgrades

The RMAN plugin and the DD OS maintain compatibility as defined in the *EMC Data Domain Boost Compatibility Guide*. Consult that publication before upgrading either the RMAN plugin 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 for DD Boost backups and file replication to work:

- 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 Data Domain 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 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 user name and password are case-sensitive, and they must match the username and password provided in the task [Registering Each Data Domain System on page 44](#).

```
# 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 using the `ddbboost access del clients` command will not stop a file backup that is already in progress. All operations in progress will continue. The change in access rights will 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 will terminate any existing operations, so be careful that no RMAN operations are in progress when disabling DD Boost. After changing access, `ddbboost enable` will re-enable DD Boost and allow RMAN operations that will now be subject to the changed access rights.

5. Enable DD Boost:

```
# ddbboost enable
DD Boost enabled
```

Configuring the DD Boost Server

The tasks in this section 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 performing RMAN backup and restore operations.

Storage unit names are case-sensitive.

The storage unit name must be provided when performing RMAN backup and restore operations.

For more information, see [AIX, HPUX, Linux, or UNIX Environment on page 37](#) or [Windows Environment on page 43](#).

Procedure

1. On the Data Domain system, enter:

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

2. Repeat the above step 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 RMAN plugins 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 15](#). If distributed segment processing is enabled, then the data flows as shown in [Figure 3 on page 15](#).

The option can be configured using the following command:

```
# ddboost 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 RMAN plugin 1.0 or later communicating with a Data Domain system that is running 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. See [Advanced Load Balancing and Link Failover on page 16](#) for 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 below.

Procedure

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

```
# ddboost ifgroup default add interface 192.168.1.1
# ddboost ifgroup default add interface 192.168.1.2
# ddboost ifgroup default add interface 192.168.1.3
# ddboost 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. See [Modifying an Interface Group on page 53](#).

Note

You can manage Advanced Load Balancing and Link Failover via the `ddbboost ifgroup` command (see [ifgroup on page 69](#)) or from the EMC Data Domain System Manager (DD System Manager) **Data Management** > **DD Boost** page (see the *EMC Data Domain Operating System Administration Guide*).

Configuring DD Boost-over-FC Service

In order to support the DD Boost-over-FC service, it is necessary to install supported Fibre Channel Target HBAs into the system. (See also the *EMC Data Domain Operating System Command Reference Guide* and *Administration Guide* for 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. Optional: 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`.

Note

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>
```

Example:

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

4. To display the available list of `scitarget` endpoint, enter:

```
# 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
```

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
```

Example:

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

DD Boost-over-Fibre Channel Path Management

The IFGROUP-based mechanism described in [Advanced Load Balancing and Link Failover on page 16](#) 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.

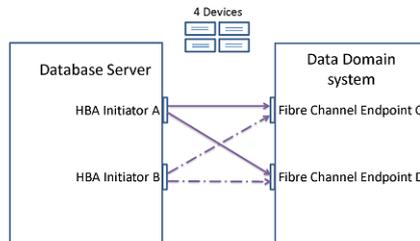
The Data Domain system advertises one or more Processor-type SCSI devices to the database server, over one or more physical paths. The operating system discovers all devices through all available paths, and creates a generic SCSI device for each discovered device and path.

For example, consider the case where:

- Database 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 containing:
 - Both FC target endpoints on the Data Domain System
 - Both initiator HBA ports
 - 4 devices (0, 1, 2, and 3)

Figure 9 DD Boost-over-FC Path Management Scenario



In this case, the database 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 desired 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.

Enabling Low-Bandwidth Optimization

To enable the low-bandwidth option, enter:

```
# 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 it to be effective.

Low-bandwidth optimization can be monitored and managed with the `ddbboost file-replication` command (see [file-replication on page 66](#)) or from the DD System Manager **Data Management** > **DD Boost** page (see the *EMC Data Domain Operating System Administration Guide*).

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

Note

- Enabling this feature takes additional resources (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.
 - For more information on this feature, see the *EMC Data Domain Operating System Administration Guide*.
-

Enabling Encrypted File Replication

To enable the encrypted file replication option, enter:

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

The output indicates that the encryption you requested was enabled.

See [file-replication on page 66](#) for more details on how to monitor and manage the encryption option for file replication.

Turning on this feature takes additional resources (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 discussion of this topic in the *EMC 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 client access control for DD Boost limits access to the Data Domain system for DD Boost clients and removes dependency on the DNS. By default, if no clients are added to the clients list when DD Boost is enabled, all clients will be automatically included in the clients list. By default a * wildcard is used.

To restrict access, remove the * wildcard from the list and add individual clients.

The database server client list may contain both fully qualified domain names or short names. The client name must match the “hostname” on the media host and is case sensitive.

To delete all clients from the DD Boost clients list, enter:

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

Optionally, to delete all clients previously added and reset the DD Boost clients list, enter:

```
# 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, enter:

```
# ddbboost clients add client-list [encryption-strength {medium | high}]
```

Example:

```
# ddbboost clients add ddbboost-dl.emc.com ddbboost-dlddbboost-dl.emc.com
: Addedddbboost-dl : Added
```

To view the DD Boost clients list, enter:

```
# 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 is not found, the client will be denied access.

Enabling In-flight Encryption

Run the following command to enable in-flight encryption for backup and restore operations over LAN:

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

This command can enable encryption for a single client, or a set of clients

Note

In scenarios where medium and high encryption are in use and the ciphers are not the same, the more secure encryption option is selected.

Note

This option is 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.

CHAPTER 4

Configuring DD Boost for RMAN

This chapter includes:

- [Install and Configure the RMAN Plug-in AIX, HPUX, Linux or Solaris Environments](#) 36
- [Install and Configure the RMAN Plug-in in Windows Environments](#) 38
- [Registering Each Data Domain System](#) 44
- [Performing File Replication](#) 47
- [Restoring From Multiple Copies](#) 49

Install and Configure the RMAN Plug-in AIX, HPUX, Linux or Solaris Environments

Installing the RMAN Plugin on Database Servers

Install the RMAN plugin on all database servers that need to access the Data Domain system.

Installation must be performed for each `ORACLE_HOME` directory used with RMAN. Set the `ORACLE_HOME` directory and the `ORACLE_SID` appropriately for each installation. Installing the plugin anywhere other than the `ORACLE_HOME` directory will prevent the plugin from working.

Perform the download and installation steps for your environment described in the following sections.

Installing the RMAN Plugin (Linux or UNIX)

Procedure

1. Stop all RMAN and Oracle processes before installing the RMAN plugin.
2. Download the RMAN plugin (a set of libraries in a `tar.gz` package) for your operating system from the Support portal (<https://support.emc.com>).
3. Use the `gunzip` and `tar` commands to uncompress the file:

```
# gunzip filename.gz
# tar -vxf filename
```

4. Ensure that `ORACLE_HOME` and `ORACLE_SID` are set properly. To verify the value of `ORACLE_HOME` or `ORACLE_SID` environment variables, use the `echo` command, shown here with an example value:

```
# echo $ORACLE_HOME
/u01/app/oracle/product/11.2.0/dbhome_1
```

To set the `ORACLE_HOME` or `ORACLE_SID` environment variables, use the `export` command:

```
# export ORACLE_HOME=/u01/app/oracle/product/11.2.0/dbhome_1
```

5. The package contains an installation script called `install.sh`. As an Oracle user, enter:

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

If the plugin already exists, the following message displays:

```
A plugin already exists, do you want to proceed with installation?
(y or n)
```

6. Enter `y` to overwrite an existing library (if the plugin already exists) or `n` to cancel.

Note

For examples see [Installing and Configuring the RMAN Plugin for Multiple SIDs Under One ORACLE_HOME on page 83](#).

Uninstalling the UNIX Plugin

Before you begin

As with the `install.sh` macro, `ORACLE_HOME` must be set and defined before running the `uninstall.sh` script, and `ORACLE_HOME` should be set to the Oracle home directory from which RMAN is to be uninstalled.

Note

It is not necessary to uninstall the UNIX plugin for an upgrade. An existing plugin is overwritten during an upgrade.

Procedure

1. Stop all RMAN and Oracle services.
2. Run the uninstall script:

```
# ./uninstall.sh
Do you want to proceed with uninstalling the Data Domain Plugin
for RMAN? (y or n)
```

3. Enter **y** to continue.

When the uninstall is complete, the following message displays:

```
The Data Domain RMAN Plugin was uninstalled.
```

AIX, HPUX, Linux, or UNIX Environment

From within RMAN, configure a channel.

Note

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

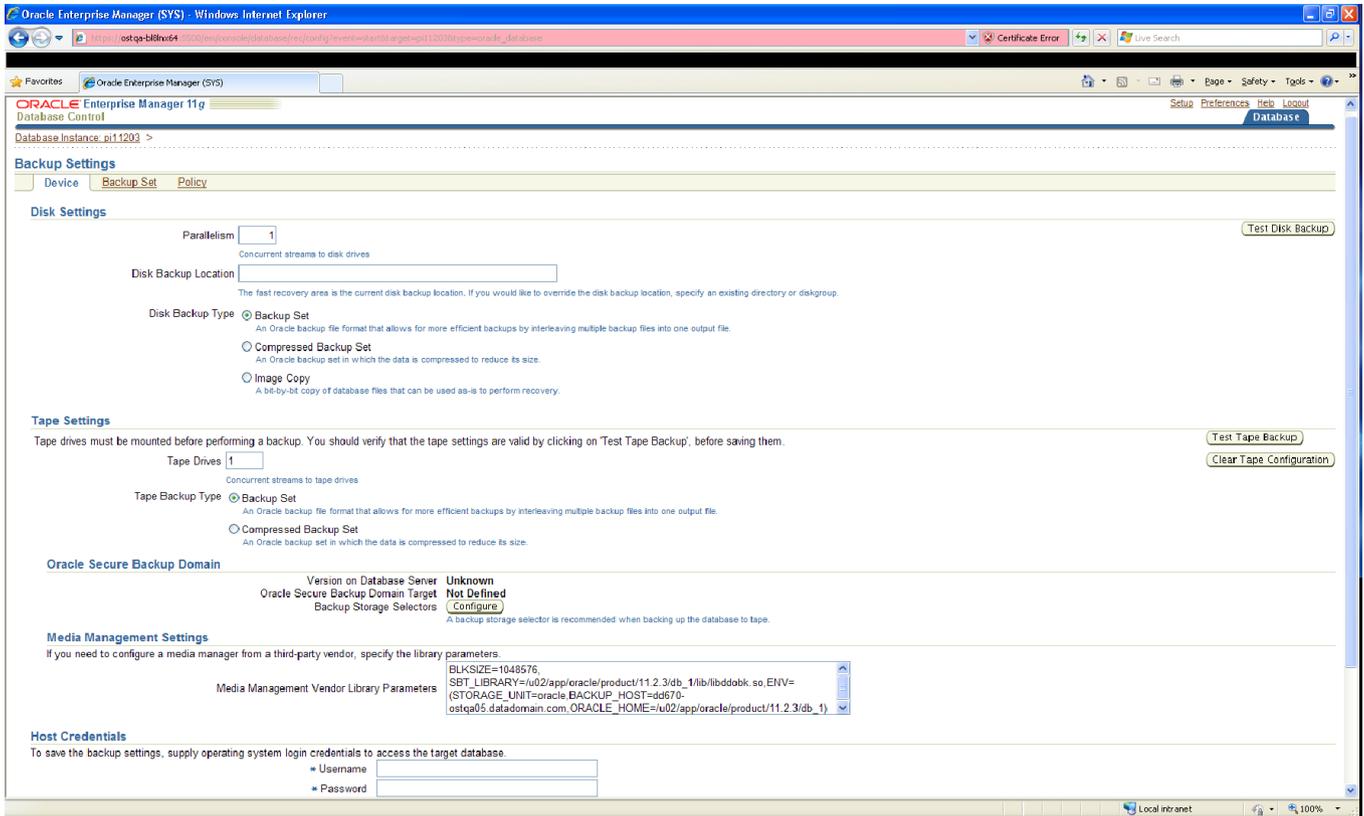
The following text shows the command syntax for configuring a channel in a Linux or UNIX environment.

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

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

An example of a Linux or UNIX configuration as viewed in the OEM is shown in the following illustration.

Figure 10 Linux or UNIX Configuration Details



Install and Configure the RMAN Plug-in in Windows Environments

Installing the RMAN Plugin on Database Servers

Note

Before installing the RMAN plugin, make sure that no backup or restore processes are running.

If this is an upgrade of a preexisting Windows DD Boost RMAN plugin, the installation may fail unless all Oracle instances are stopped. It is recommended that you stop all Oracle instances before proceeding with a Windows installation.

Installation should be performed for each `ORACLE_HOME` directory that will be used with RMAN, setting `ORACLE_HOME` and `ORACLE_SID` appropriately for each installation. Installing other than to the `ORACLE_HOME` directory will prevent the plugin from working.

The RMAN plugin software must be installed on database servers that need to access the Data Domain system.

Perform the download and installation steps for your environment described in the following sections.

Installing the RMAN Plugin (Windows)

Procedure

1. Stop all RMAN and Oracle processes before installing the RMAN plugin.
2. Download the Windows RMAN plugin installer `libDDobkSetup.exe` from the Support portal (<https://support.emc.com/>). Navigate to **Download Software > DD Boost Oracle RMAN**.

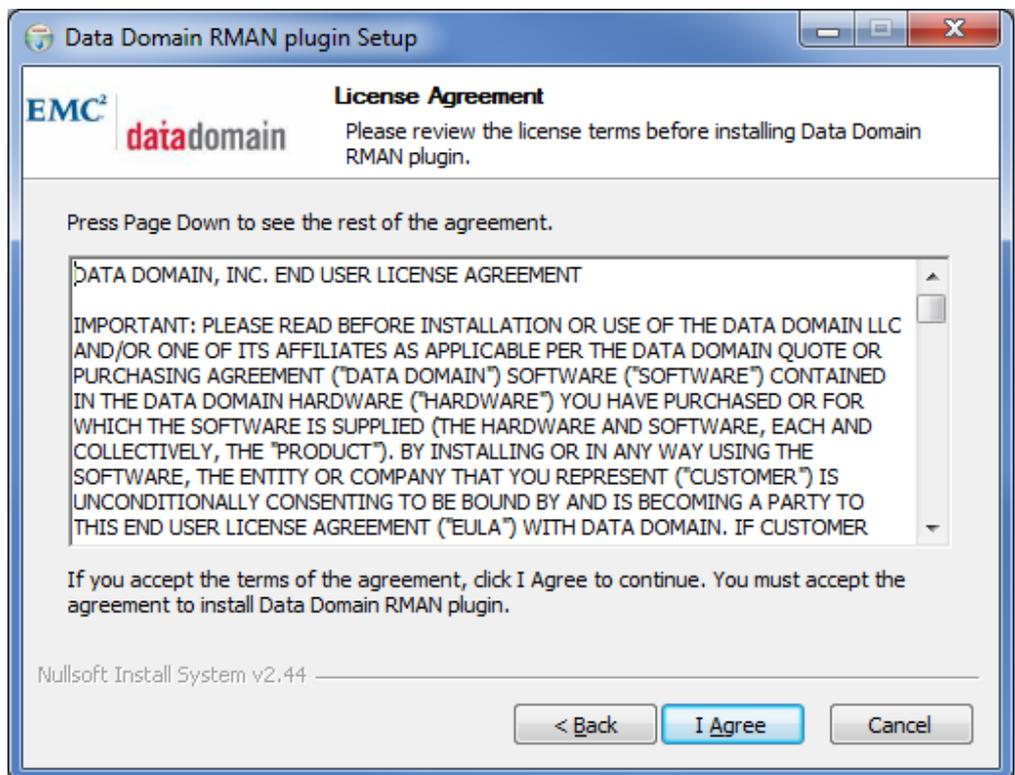
3. Double-click `libDDobkSetup.exe` to launch the installer.

A warning message displays.

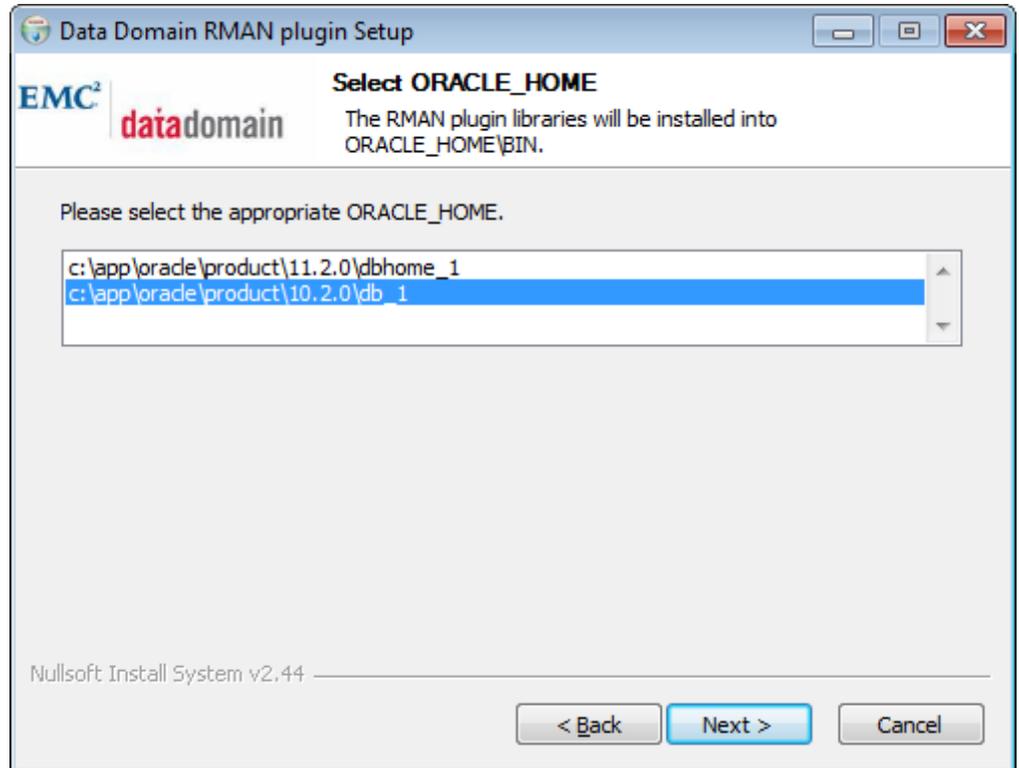
4. Click **Yes** to continue (if you stopped all RMAN and Oracle processes as instructed in Step 1).



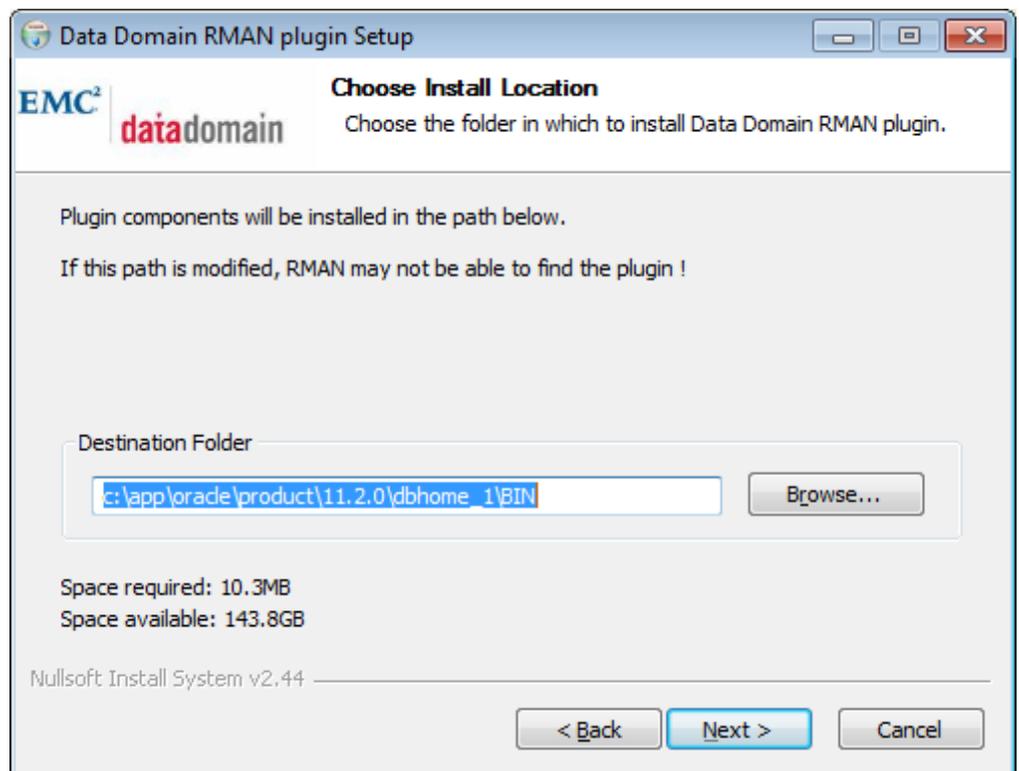
5. Read the terms of the license agreement and click **I Agree** to continue.



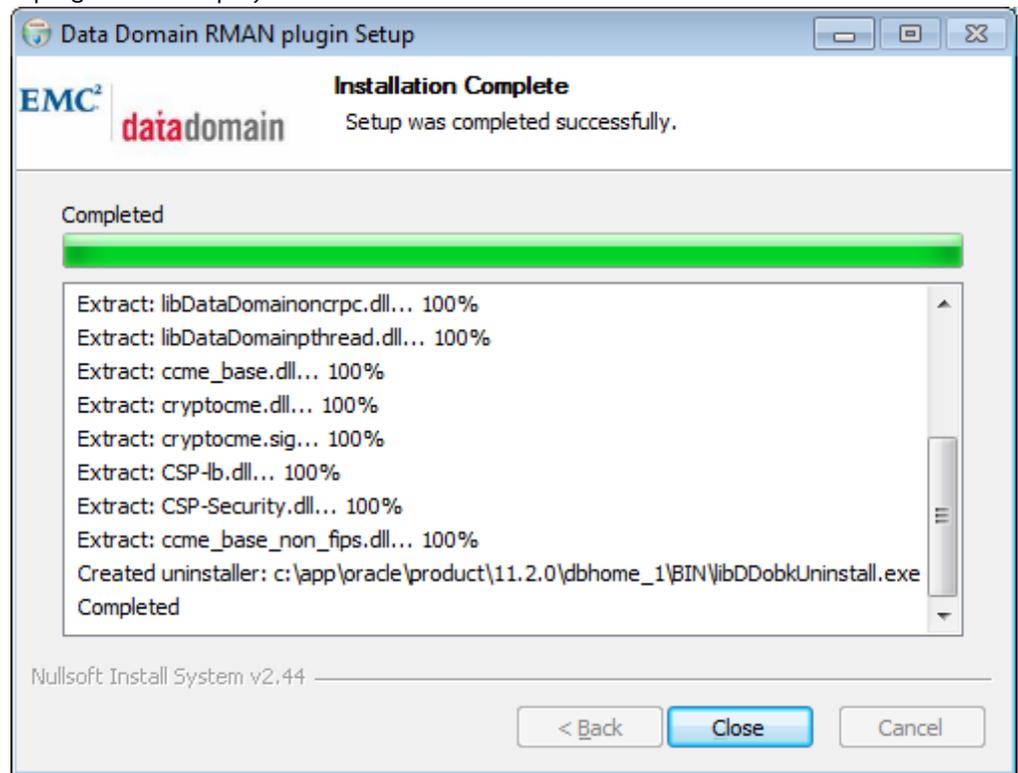
6. In the **Select ORACLE_HOME** panel, select from the detected list of `ORACLE_HOME` directories and click **Next**.



7. In the **Choose Install Location** panel, the destination folder (*ORACLE_HOME*) is shown. Plugin components are installed in folders relative to this destination folder. Typically, this should not be changed. Click **Install** to start the installation.



A progress bar displays.



8. When the Installation is complete, click the **Show Details** button to view the files installed and their location.

Uninstalling the Windows Plugin

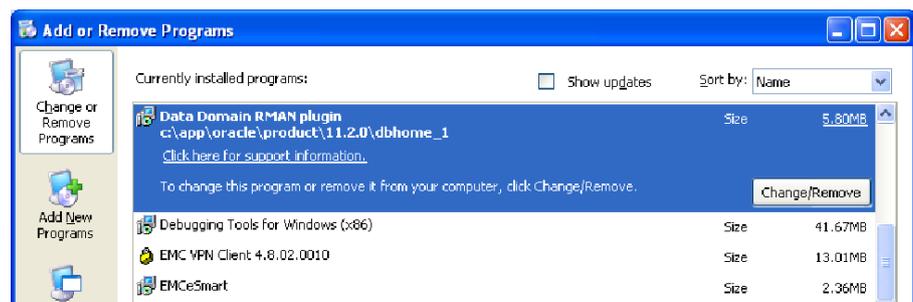
Note

It is not necessary to uninstall the Windows plugin for an upgrade. An existing plugin is overwritten during an upgrade.

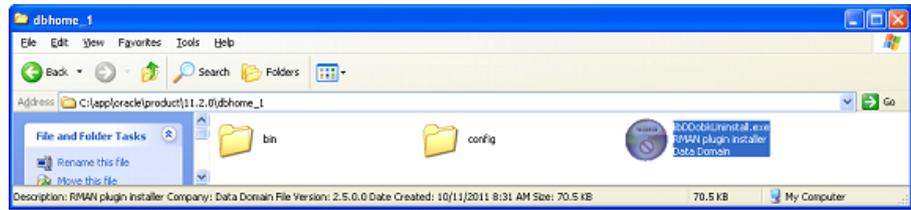
Procedure

1. Stop all RMAN and Oracle services.
2. Perform one of the following tasks:
 - a. Uninstall Using the Windows Control Panel

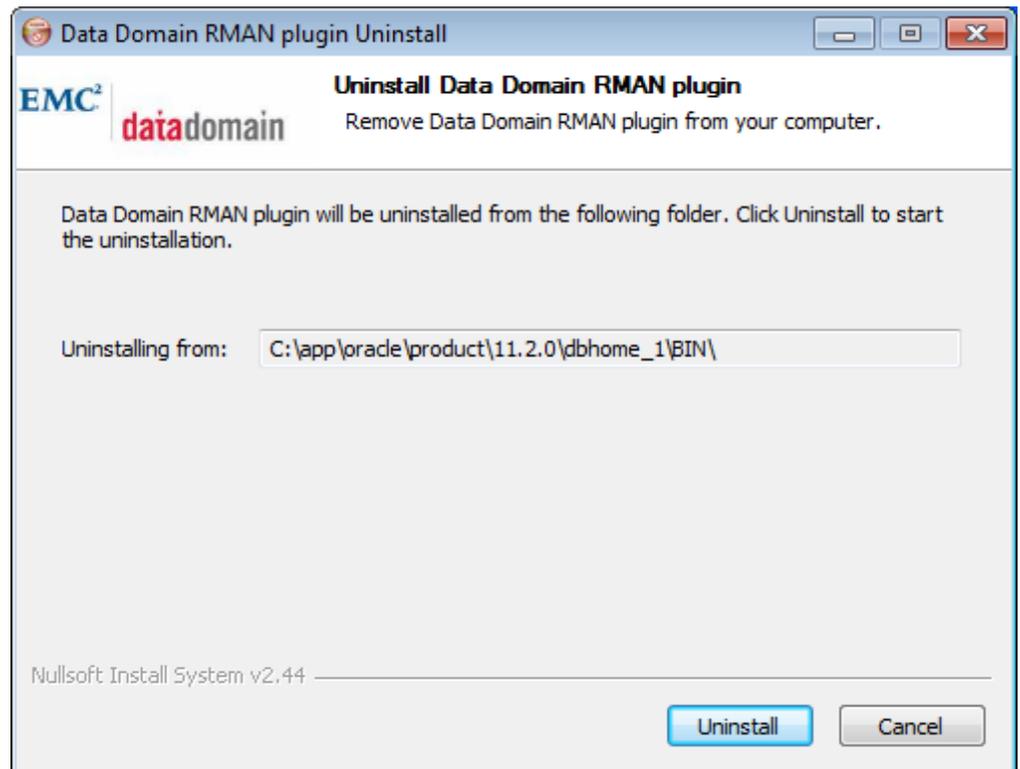
Uninstall the Data Domain RMAN plugin using the Windows Control Panel uninstall/remove program feature (as you would uninstall a typical Windows program). Be careful to select and uninstall the correct item. In the following illustration the RMAN plugin is highlighted.



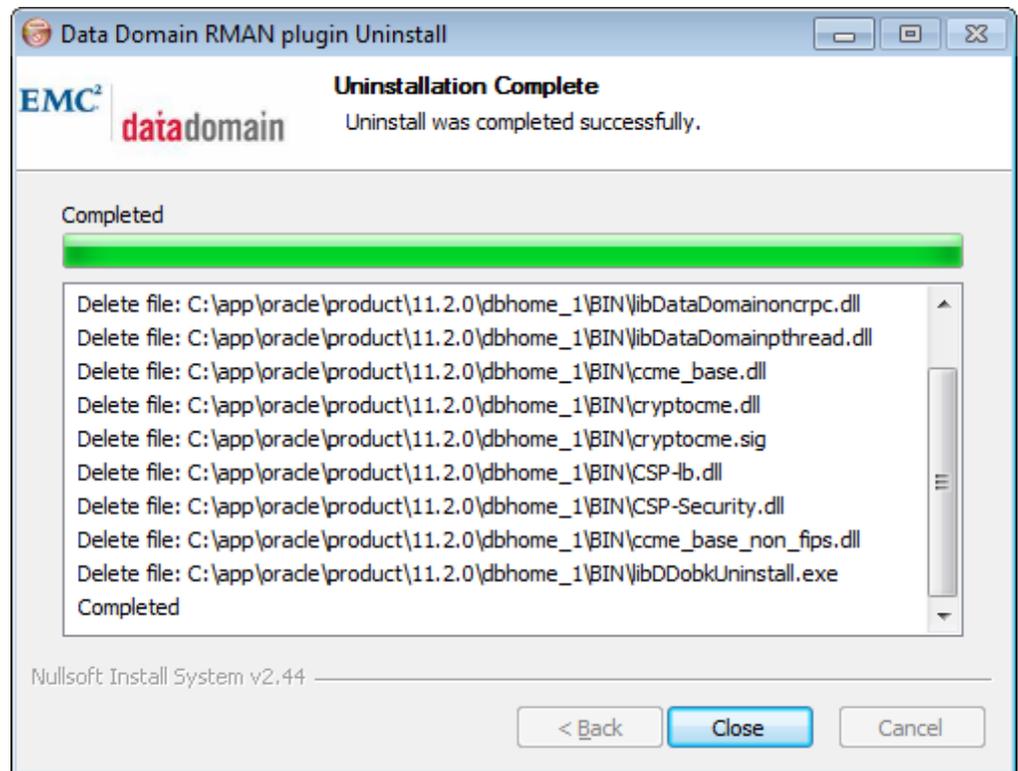
- b. Uninstall Using the Uninstaller. Double-click `libDDobkUninstall.exe`, which was installed in the `ORACLE_HOME` directory.



- 3. In the **Uninstall Data Domain RMAN plugin** panel, click **Uninstall**.



- 4. When the uninstall completes, click **Show Details** to view which files were removed.



Windows Environment

In a Windows environment the configuration syntax differs between Oracle Database 11g and 10g.

Oracle Database 11g

From within RMAN, configure a channel.

Note

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

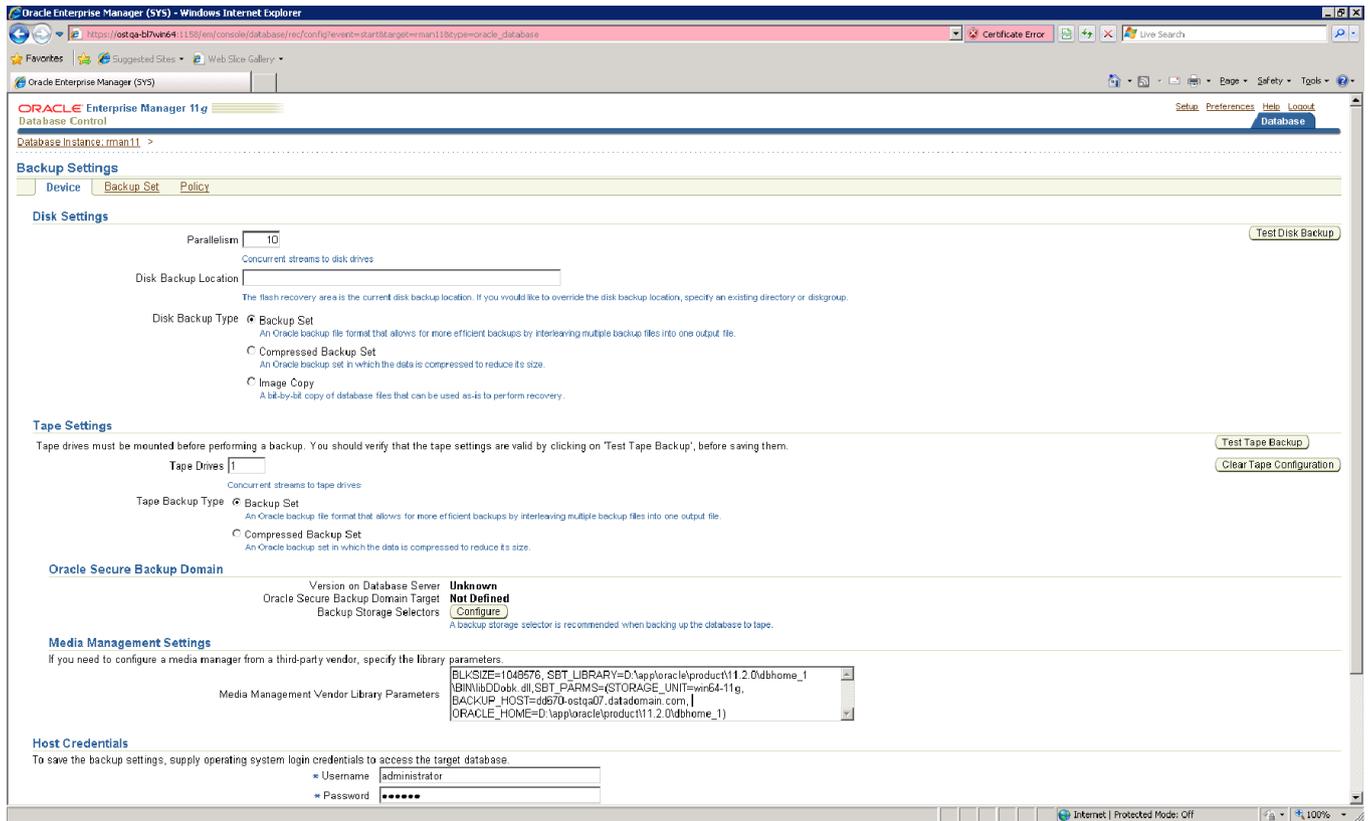
The following text shows the command syntax for configuring a channel in a Windows environment for Oracle Database 11g:

```
RMAN> CONFIGURE CHANNEL DEVICE TYPE 'SBT_TAPE'
PARMS 'BLKSIZE=1048576,
SBT_LIBRARY=path-to- $\$$ ORACLE_HOME\BIN\libDDobk.dll,
SBT_PARMS=(STORAGE_UNIT=storage-unit-on-Data-Domain-system,
BACKUP_HOST=hostname-of-Data-Domain-system,
ORACLE_HOME=path-to- $\$$ ORACLE_HOME) ';
```

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

An example of a Windows configuration as viewed in the OEM is shown in the following illustration.

Figure 11 Windows Configuration Summary Screen



Oracle Database 10g

From within RMAN, configure a channel.

Note

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

The following text shows the command syntax for configuring a channel in a Windows environment for Oracle Database 10g:

```
RMAN> CONFIGURE CHANNEL DEVICE TYPE 'SBT_TAPE'
PARMS 'BLKSIZE=1048576,
SBT_LIBRARY=path-to-$ORACLE_HOME\BIN\libDDobk.dll,
SBT_PARMS=(STORAGE_UNIT=storage-unit-on-Data-Domain-system,
BACKUP_HOST=hostname-of-Data-Domain-system,
ORACLE_HOME=path-to-$ORACLE_HOME)';
```

The configuration that you specify is also displayed in the Media Management Settings field of the 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 DD Boost RMAN plugin. This registration provides the RMAN plugin with the necessary username and password for connecting to the Data Domain system. The RMAN plugin stores the username and password securely in an encrypted lockbox file and retrieves them when connecting to a Data Domain system.

The first time a Data Domain system is registered, the encrypted lockbox file is created in the `/config` directory in 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 attempts to access the lockbox, access to the lockbox file is not permitted. This prevents other systems from connecting to the Data Domain system. Normally, these other systems will have their own Oracle instance and Oracle home directory with RMAN installed and therefore will have their own lockbox file.

In these typical cases when there is only a single Oracle host system in use, or when multiple host systems exist with each having 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 45](#). Each client system will have its own lockbox file accessible only to that system.

When there is a Data Domain system with 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 46](#).

In some cases, such as in an Oracle RAC environment where multiple system nodes are sharing an Oracle home directory and thus sharing a lockbox file, it is necessary to create a lockbox file that can be accessed by multiple systems. In this case, follow the instructions in [Registering a Data Domain System for Shared Lockbox Access on page 46](#).

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 such as shown in the example below. Edit the command for your environment. This registration task is required only once for each Data Domain system; the registration command should not be included in each backup script. This registration must be performed before any backups can be made.

```
RUN {
ALLOCATE CHANNEL C1 TYPE SBT_TAPE PARMS
 'SBT_LIBRARY=path-to- $\$$ ORACLE_HOME/lib/libddobk.so';
send 'set username user password word servername server';
RELEASE CHANNEL C1;
}
```

The username is the `ddboost` user that you set with the `user add` command and `ddboost set username` command in Step 4 of the task [Enabling DD Boost on a Data Domain System on page 26](#). The password is the one that you set up with the DD OS CLI `user` command referred to in the Notes at the end of that task section. Remember that both username and password are case-sensitive and both must match exactly as set on the Data Domain system. The `servername` server is the name of the Data Domain system exactly as specified by the `BACKUP_HOST` value in the `CONFIGURE CHANNEL` command. The `servername` 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 `servername`. 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 to and from your Data Domain system using `'device type sbt_tape'`. If you want all your backups to go to the Data Domain system, configure `SBT_TAPE` as your default device 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 PARMS
  'SBT_LIBRARY=path-to- $\$$ ORACLE_HOME/lib/libddobk.so';
send 'set username user password word servername DFC-server';
RELEASE CHANNEL C1;
}
```

For the *DFC-server* value, enter the configured Fibre Channel server name.

Registering a Data Domain System for Shared Lockbox Access

If multiple database host systems need to use the same lockbox file as might be the case in an RAC environment, there is an additional step in the procedure for creating the lockbox file. This step is needed so that host database systems other than the one that created the lockbox can also access the lockbox. This step must be done before any Data Domain systems are registered. Once this step is done, then each Data Domain system can be registered.

The `hostnames` command specifies all of the hosts that need to access the lockbox file. When you use the `hostnames` command to create the lockbox, all of 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 of the hosts sharing 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 situation, you must delete the existing lockbox file and then create a new lockbox file with all of the host names.

The `hostnames` command format is:

```
RUN {
  ALLOCATE CHANNEL C1 TYPE SBT TAPE PARMS
    'SBT_LIBRARY=path-to- $\$$ ORACLE_HOME/lib/libddobk.so';
  send 'hostnames hostname-1 hostname-2 ... hostname-n';
  RELEASE CHANNEL C1;
}
```

This command must be done as the first RMAN command, before the registration command. All of the hosts that will access the shared lockbox file must be listed in the list *hostname-1 ... hostname-n*. These *hostname* values should be entered as the fully-qualified host names (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- $\$$ ORACLE_HOME/lib/libddobk.so';
  send 'hostnames dbclient1.datadomain.com dbclient2.datadomain.com
dbclient3.datadomain.com';
  RELEASE CHANNEL C1;
}
```

When the `hostnames` command is received by the RMAN plugin, it checks to see 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 desired and repeat the command. The RMAN plugin will not delete an existing lockbox file.

If no lockbox file exists, the RMAN plugin creates a lockbox file and all of the specified hosts are granted access to the new lockbox. This means that the lockbox can then be used and shared by all of these hosts. Next, register each Data Domain system that these hosts will be accessing. You can register additional Data Domain systems at any time

after the lockbox has been created. Each Data Domain system is registered as described in [Registering a Standard Data Domain System on page 45](#). Since all of the specified hosts have access to the lockbox, the registering of the Data Domain systems can be done on any of the listed host systems. Each Data Domain system need only be registered once.

It is not possible to add additional host names after the `hostnames` command has been executed. The only way to add additional host names is to delete the current lockbox file and then recreate the file by running the `hostnames` command again, this time with all of the desired host names. This means the lockbox contents are lost, and all of 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 using RMAN 1.2 and the instructions in this section. It is not possible to upgrade a lockbox created with earlier versions of RMAN so that it can be accessed by multiple client database systems.

Note

Lockbox files created with RMAN 1.0 or 1.1 will continue to work with RMAN 1.2 when accessed from the client database system that created them; it is not necessary to modify the lockbox file.

Performing File Replication

File replication allows you to make multiple copies of an RMAN backup. Prior to RMAN 1.2.0.2, two copies of a backup could be made by specifying `COPIES 2` in the backup command. RMAN 1.3.0.0 provides the ability to 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 below. Edit the script command for your environment, specifying the Data Domain system that is the destination system for the file replication. Specifying `COPIES 2` requires an installed Replicator license on both the source and destination Data Domain systems. Making a backup by specifying `COPIES 2` requires Oracle Enterprise Edition.

For example, specifying `COPIES 2` instructs the DD Boost RMAN plugin to perform Managed File Replication to the second Data Domain system. The second Data Domain system name is specified as part of the file name of the second copy. This is the `'destination-system'` in the command below.

The two file names specified in this command must be identical except for the inclusion of `'destination-system'` in the second name. The RMAN plugin requires that the file names of the two copies be identical. The underscore character `"_"` is also required. Omitting the underscore in the file names will result in an error and the backup will fail. Any changes made to the file name formats shown must include an underscore character and must guarantee that the file names are identical, except for `'destination-system'` in the name for the second copy. This means that using Oracle formats such as `"%U"` and `"%c"` is not allowed, as these Oracle formats include the copy number in the generated file name and thus result in the file name generated for the second copy being different from the file name generated for the first copy.

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

Note

Enter a small letter "u" in '%u%p'. A capital "U," while a valid Oracle format, is invalid for DD Boost RMAN and causes the replication to fail.

To use **COPIES 2**, the value of *BACKUP_TAPE_IO_SLAVES* must be **TRUE**, which you can configure by entering:

```
SQL> show parameter backup_tape

NAME                                TYPE                                VALUE
-----                                -
backup_tape_io_slaves                boolean                             FALSE
SQL>

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
SQL>
```

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 normal 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:

```
backup copies <n> database filesperset=1 format '%u%p',
'destinationssystem-2/%u%p', 'destinationssystem-3/%u%p',
'destinationssystem-4/%u%p'
```

. Specifying two, three, or four copies creates the additional copies on the second, third, and fourth destination systems respectively. All of the formats must be identical except for the initial Data Domain system name so that all of 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. This means that 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 will be 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 performing backups. The network configuration must allow backups to the source Data Domain system and to all of the destination Data Domain systems. The destination systems must all be accessible from the source system. When more than two copies are being made the destination systems do not need to be accessible to each other, only to the source system. Replication will not work in configurations where the destination Data Domain system is only connected to the source Data Domain system via a private network or link.

The following example shows the output of the RMAN `list` command, which shows 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 Copies

If there are multiple copies of a backup (as described in the previous section), a restore can be performed from any Data Domain system that contains a copy of the backup. Typically, a restore is performed 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 DDR is not available for any reason, the backup can be restored 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 desired backup is available, choose one based on factors such as the available bandwidth, the speed of the connection, or other performance or logistical consideration. Specify that desired system as the backup host in the restore command.

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

For example, if three copies of a backup were made to ddr-host-1, ddr-host-2, and ddr-host-3, and ddr-host-1 is down when a restore is required, 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 attempts to get the file from ddr-host-3. If that does not succeed, the backup fails.

If you choose 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 attempts to get the file from ddr-host-2.

Note

If ddr-host-2 is down, but ddr-host-1 & ddr-host-3 are available, specifying ddr-host-3 as the backup host performs the restore from ddr-host-3, but if a file is not found on ddr-host-3, RMAN will 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 not available the restore will always fail even if the backup is available from other copies.

CHAPTER 5

Data Domain System Administration

This chapter includes:

- [Restricting DD Boost Access to Specific Database Servers](#)..... 52
- [File Replication Version Compatibility](#).....52
- [Modifying an Interface Group](#)..... 53
- [Sizing DD Boost-over FC Device-Set](#)..... 54

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.

This example shows how to 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
```

This set of steps establishes a set of access controls that enables DD Boost access only to the two database servers specified (**`databaseserver1.datadomain.com` and `databaseserver2.datadomain.com`**).

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

Note

- Ensure that no backup jobs are running to the Data Domain system when changing access control. Use the `ddboost disable` command to prevent operations while access is changed. Disabling DD Boost disables data access to all database servers.
 - Specify only a fully-qualified domain name, IP address, or resolvable DNS name for the client when modifying 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 using the `ddboost access del clients` command will not stop a file backup that is already in progress. All operations in progress will continue. The change in access rights will not cause existing operations to fail.
 - Enter the `ddboost enable` command to re-enable DD Boost and allow RMAN operations after changing 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 using the DD Boost RMAN plug-in. 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 follows:

- All maintenance and patch versions within a *family*, are backwardly 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. Make sure 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 enter 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 would no longer be 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 using the `net` command that modify the interfaces, such as enabling an interface that is configured for ifgroup, execute the `ddboost show connections` command to update the load balancing view. This updating allows the ifgroup to use the interface. See [show connections on page 72](#).

Removing Advanced Load Balancing and Link Failover Configuration

The following example illustrates removing a configured interface group on the Data Domain system.

Procedure

1. Make sure 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 using the following command:

```
# ddboost show connections
```

See [show connections on page 72](#) for sample output.

2. Make sure 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:

```
# ddbboost ifgroup disable default
```

4. Reset the interface group:

```
# ddbboost ifgroup default reset
```

Results

All the interfaces are released from the group.

Sizing DD Boost-over FC Device-Set

As described in [DD Boost-over-Fibre Channel Path Management on page 20](#), 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 an initiator, by adding the initiator to a ddbboost-type scsitarget access group:

```
# ddbboost 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:

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

The maximum number of supported DFC devices per Data Domain system is 64.

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

1. Is the database server queue-depth constrained?
As described in [DD Boost-over-Fibre Channel Path Management on page 20](#), 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 (e.g. 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 database server and Data Domain system
For each advertised DFC device, the database server operating system will create n generic SCSI devices, one for each physical path through which the database server OS can access the device.

For example, if:

- database 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 database 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 database 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.

CHAPTER 6

Troubleshooting

This chapter includes:

- [About Troubleshooting](#)..... 58
- [Investigating Problems](#)..... 58
- [Oracle Limitations](#)..... 58
- [Reporting Problems](#)..... 58
- [Data Domain System Settings for File-Replication](#)..... 59
- [Lockbox Troubleshooting Information](#)..... 59
- [Failure to Load Media Management Library](#)..... 59

About Troubleshooting

This chapter provides basic troubleshooting tips that might enable customers to resolve issues on their own. For issues that cannot be resolved, customers should contact their contracted support providers.

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

Investigating Problems

When investigating 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 *Data Domain Boost Compatibility Guide*, which is available at the Data Domain Support portal (<https://support.emc.com>). Navigate to **Documentation > Compatibility Matrices** and select **Data Domain Boost Compatibility List**. A supported configuration can become unsupported if any component changes.
- **Authorization Failures**
If you encounter authorization failures, ensure that all of 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 RMAN plugin 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 Data Domain always include:

- The `sbtio.log` file (This is critical for analysis.)
- Output of the RMAN command: `show all`
- All RMAN commands that you used.

If possible, delete the `sbtio.log` file, recreate the issue in as few steps as possible, then capture the `sbtio.log` file to send with your problem report. Set TRACE level to 5 (with the `CONFIGURE CHANNEL` command) when attempting to recreate your problem.

RMAN determines where the `sbtio.log` file is located. The `sbtio.log` files are 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 try a `find` command to locate it.

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. See the *EMC Data Domain Operating System Administration Guide* for information on using the `replication throttle` command to display or control the throttle settings.

Lockbox Troubleshooting Information

The `ddboost.config` file is the DD Boost RMAN configuration file. It contains the saved username/password pairs for the DD systems you are using. If you are connecting to system `ddr-system` as user `ddr-user` with password `user-password`, then there is an entry in this file that says when making connections to `ddr-system` make the connection using username `ddr-user` with password `user-password`. This file is encrypted and maintained with the RSA lockbox libraries (which are installed as part of the DD Boost RMAN plugin) hence some of the error strings returned in `sbtio.log` say `lockbox`.

If you change the hostname of the Data Domain system without changing it in all the scripts etc., then attempts to connect to `ddr-system` fail because there is no host with that name. 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. For instructions, see [Registering Each Data Domain System on page 44](#).

The `ddboost.config` file is located in the `ORACLE_HOME/config` directory. When installing the RMAN DD Boost plugin, make sure that you have a writable `/config` directory in your specified `ORACLE_HOME` directory. If the `ddboost.config` file is not found, it is created. When you specify the `ORACLE_HOME` directory, it must contain a `/config` directory. Not having this directory present can cause failure to create the `ddboost.config` file.

You cannot reference a lockbox except from the database server where it was created, unless at the time you created the lockbox you set it up to be accessible from multiple servers. See [Registering a Data Domain System for Shared Lockbox Access on page 46](#).

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

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

Deleting the `ddboost.config` file(s) is never 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 Media Management Library

After installation and configuration of the RMAN plugin is completed, attempts to use the RMAN plugin 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 a Unix environment (Linux, HP/UX, or AIX). The error message indicates that RMAN has not been able to locate one of the required libraries, either `libddobk.so`, `libDDBoost.so`, or another library used by these libraries. In different Linux environments, the search rules used for locating shared libraries, such as the RMAN plugin libraries, are complex. The search rules can depend on the settings of various environment variables, such as `LD_LIBRARY_PATH`, `LD_LIBRARY_PATH_64`, or `LD_RUN_PATH`. These rules and environment variables may be set in your environment by other applications and may cause RMAN to be unable to locate the needed libraries.

If you encounter this error in a Unix environment, the following steps can help find the source of the problem and correct it.

Procedure

1. Check that the pathname of the `libddobk.so` library was correctly specified in the RMAN configuration command.
 (See [AIX, HPUX, Linux, or UNIX Environment on page 37](#) or [Windows Environment on page 43](#).) If this is not correct, the `libddobk.so` library cannot be found. Check that the `libddobk.so` exists in the specified directory and the directory and the library file have appropriate access for the user.
2. Check that the Oracle home directory was correctly specified in the RMAN configuration command.
 (See [AIX, HPUX, Linux, or UNIX Environment on page 37](#) or [Windows Environment on page 43](#).) If this is not correct, the `libDDBoost.so` library cannot be found. Check that the `/lib` subdirectory exists in the Oracle home directory and that the `libDDBoost.so` library exists in this directory and that the directory and library file have appropriate access for the user.
3. Use the `ldd` command on the `libddobk.so` and the `libDDBoost.so` libraries to see if the libraries they use can be found, as shown in the following example:

```

# ldd libddobk.so
libpthread.so.0 => /lib64/libpthread.so.0 (0x00002aaaaabf1000)
librt.so.1 => /lib64/librt.so.1 (0x00002aaaaad06000)
libm.so.6 => /lib64/libm.so.6 (0x00002aaaaae20000)
libdl.so.2 => /lib64/libdl.so.2 (0x00002aaaaafa7000)
libDDBoost.so => not found
libc.so.6 => /lib64/libc.so.6 (0x00002aaaab420000)
/lib64/ld-linux-x86-64.so.2 (0x0000555555554000)
    
```

All the libraries should show as resolving. If `libDDBoost.so` is shown as not found as in this example, then that is the missing library. Note that `libDDBoost.so` may be shown as not found even though there is a `libDDBoost.so` file in the same directory as `libddobk.so`.

4. Correcting the problem depends on your environment and your other applications. One possible solution is to place a link to the `libDDBoost.so` library in a place where other libraries are being found. In this example, placing a link to `libDDBoost.so` in `/lib64` can solve the problem:

```
cd /lib64
ln -s /u01/app/oracle/product/11.2.0/db_home1/lib/libDDBoost.so
libDDBoost.so
```

It may be possible to solve the problem by adding the directory containing `libDDBoost.so` to an environment variable such as `LD_LIBRARY_PATH`, `LD_LIBRARY_PATH_64`, or `LD_RUN_PATH`. Use caution when using or modifying environment variables, as their use varies from system to system and changing them might impact other applications.

5. After making your change, use `ldd` to verify that the library is resolved:

```
# ldd libddobk.so
libpthread.so.0 => /lib64/libpthread.so.0 (0x00002aaaaabf1000)
librt.so.1 => /lib64/librt.so.1 (0x00002aaaaad06000)
libm.so.6 => /lib64/libm.so.6 (0x00002aaaaae20000)
libdl.so.2 => /lib64/libdl.so.2 (0x00002aaaaafa7000)
libDDBoost.so => /u01/app/oracle/product/11.2.0/db_home1/lib/libDDBoost.so
(0x00002aaaab0ab000)
libc.so.6 => /lib64/libc.so.6 (0x00002aaaab420000)
/lib64/ld-linux-x86-64.so.2 (0x0000555555554000)
```

Note

Even if `ldd libddobk.so` report looks good, it is still possible that Oracle did not find the library and you get the following error message:

```
ORA-27211: Failed to load Media Management Library
Additional information: XXXX
```

If you get the error message, you should create a Symlink so that Oracle finds the library:

- **SUSE platform:** Symlink `/lib64/libDDBoost.so -> $ORACLE_HOME/lib/libDDBoost.so`
 - **HPUX platform:** Symlink `/usr/lib/hpux64/libDDBoost.so -> $ORACLE_HOME/lib/libDDBoost.so`
 - **Solaris platform:** Symlink `/usr/lib/64/libDDBoost.so -> $ORACLE_HOME/lib/libDDBoost.so`
-

CHAPTER 7

Using DD Boost Commands

This chapter includes:

- [Command Summary](#).....64
- [ddboost Command Options](#).....64

Command Summary

The `ddboost` command includes the following options.

Table 3 Summary of DD Boost Commands

<code>access</code>	Manage DD Boost access for clients.	access on page 64
<code>destroy</code>	Delete all storage units and their contents from the Data Domain system.	destroy on page 65
<code>disable</code>	Disable DD Boost.	disable on page 65
<code>enable</code>	Enable DD Boost.	enable on page 65
<code>fc</code>	Add or modify DD Boost over Fibre Channel devices.	fc on page 65
<code>file-replication</code>	Set, reset, or show the status of the low- bandwidth optimization and/or encryption option for file replication. Show or reset statistics, history, and performance.	file-replication on page 66
<code>ifgroup</code>	For Advanced Load Balancing and Link Failover: Manage interface group (IP addresses) on a Data Domain private network that processes data transfer.	ifgroup on page 69
<code>option</code>	Set, reset, or show the status of the distributed segment processing feature.	option on page 70
<code>reset</code>	Delete the DD Boost user and reset statistics. Clear all job connection counters shown in the output of the <code>ddboost showconnections</code> command in the event that network connection is lost.	reset on page 71
<code>set user-name</code>	Set the DD Boost user name when DD boost is enabled	set user-name on page 71
<code>show</code>	Show the DD Boost connections, DD Boost histogram, or statistics.	show on page 72
<code>status</code>	Show whether DD Boost is enabled or disabled.	status on page 73
<code>storage-unit</code>	Create and delete a storage unit, or list the storage units and files in a storage unit.	storage-unit on page 73

ddboost Command Options

access

Manage DD Boost access for clients.

ddboost access add clients

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

Add clients to DD Boost access list.

ddboost access delete clients

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

Delete clients from DD Boost access list.

ddboost access reset

```
ddboost access reset
```

Reset DD Boost client access list to factory default.

ddboost access show

```
ddboost access show
```

Show 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 storage units and their contents from the Data Domain system. The command permanently removes all of 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.

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

fc

For a full list of `ddboost fc` commands consult the *EMC Data Domain Operating System Command Reference Guide*.

fc group add

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

Administrators can use this command to grant access to these DD Boost over Fibre Channel devices via 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 can 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 can use this command to reset low-bandwidth optimization and/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 entered on both Data Domain systems—the source and destination (target) systems. Only an administrator can set these two 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 maximum filesystem 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. Enter this CLI command on the source and destination systems:

```
# filesys clean start
```

The amount of time the cleaning cycle takes 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 can use this command to reset file-replication statistics when DD Boost is enabled.

file-replication show active

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

Use this command to show 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. On a 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}]
```

This command shows 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 no error occurred. A sample line with a description of each column follows:

- 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 hr]
```

This command shows the file-replication history over time. A cumulative history is shown, 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 show 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 is shown for 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 4 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 hr]
```

Use this command to show 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 5 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                Inbound
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

ifgroup is an interface group.

ifgroup add interface

```
ddboost ifgroup add interface group-name {interface ipaddr | client host}
```

Add an interface, client, or both to *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.

See the *EMC Data Domain Operating System Administration Guide* for more information on interface groups.

ifgroup create *group-name*

```
ddboost ifgroup create group-name
```

Create a group with the name *group-name* for the interface. This command option may be run on a single node and on a master node in a Global Deduplication Array (GDA).

The *group-name* may contain characters ^, [0-9, a-z, A-Z],* \$, underscore(_), and hyphen (-). Hostnames and fully qualified hostnames may be used. Wildcard hostnames, indicated by an asterisk (*) may be used. Reserved group names that cannot be used are: “default,” “all,” and “none.”

ifgroup del interface

```
ddboost ifgroup del group-name {interface ipaddr | client host}
```

Remove an interface, client, or both from *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
```

Destroy the group name. Only empty groups can be destroyed. Interfaces or clients cannot be destroyed but may be removed sequentially or by running 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
```

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
```

Enable a specific group by entering the *group-name*. If *group-name* is not specified, the command applies to the default group.

ddboost ifgroup rename

```
ddboost ifgroup rename source-group-name destination-group-name
```

Rename the ifgroup *source-group-name* to *destination-group-name*. This command option does not require disabling the group.

ifgroup reset group-name

```
ddboost ifgroup reset group-name
```

Reset a specific group by entering the *group-name*. If *group-name* is not specified, the command applies to the default group.

ifgroup show config

```
ddboost ifgroup show config {interfaces | clients | groups | all}
group-name
```

Display selected configuration options. If no selection is made, all information about the specified *group-name* is shown.

If *group-name* is not specified, all information about the default group is displayed. Select the all argument to view configuration options of all groups. All users may run this command option.

ifgroup status group-name

```
ddboost ifgroup status group-name
```

Show status of link aggregation: enabled or disabled. Status is displayed for all groups unless *group-name* is specified. All users may run this command option.

option

option reset distributed-segment processing

```
ddboost option reset distributed-segment processing
```

Administrators use this command to reset distributed segment processing to the default option, which is enabled.

option set distributed-segment processing

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

Administrators can 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
```

Reset virtual synthetics to the default option (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}
```

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]
```

Show status of virtual synthetics: enabled or disabled. All users may run this command.

reset

reset stats

```
ddboost reset stats
```

Administrators use this command as follows:

- to reset all statistics when DD Boost is enabled, or
- as a network recovery procedure to clear job connections after the network connection has been lost.

reset user-name

```
ddboost reset user-name
```

Delete and reset the DD Boost user name.

set user-name

```
ddboost set user-name user-name
```

Set the DD Boost user name when DD Boost is enabled.

show

show connections

```
ddboost show connections
```

Use this command to show the following information:

- the number of active clients, and for each active client:
 - the client's name
 - whether or not the client is idle (yes/no)
 - the number of installed CPUs
 - the amount of memory in MiB
 - the installed RMAN plugin version number
 - the name and version of the operating system
 - the RMAN version number.
- the number of connections that a system uses for DD Boost, and whether these connections are spread across interfaces using Advanced Load Balancing and Link Failover. (See sample output for client connections below.)
- the number of connections used for a given group, even when the interface has been removed from the ifgroup.
- an overview of interfaces available for DD Boost.

show histogram

```
ddboost show histogram [op]
```

Use this command to display a DD Boost histogram for the Data Domain system.

Table 6 DD Boost Histogram

<i>op</i>	The name of the operation.
mean-ms	The mathematical mean time for completion of the operations.
stddev	The standard deviation for time to complete operations, derived from the mean time.
max-s	The maximum time taken for a single operation.
2, 4, 6, 8, or 10ms	The number of operations that took less than the specified number of milliseconds (ms).
100ms	The number of operations that took between 10 ms and 100 ms.
1s	The number of operations that took between 100 ms and one second.
10s	The number of operations that took between 1 second and 10 seconds.
>10s	The number of operations that took over 10 seconds.

show stats

```
ddboost show stats [interval seconds] [count count]
```

Use this command to show the output of 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 shown. Optionally, you can show 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 via distributed segment processing and the amount of compression achieved is displayed.

show user-name

```
ddboost show user-name
```

Display the current DD Boost user.

status

```
ddboost status
```

Enter 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}]
```

Create a storage unit and set limits. See `mtree` in the *Data Domain Operating System Command Reference Guide* for details on quota limits. Note that if the quota feature is not enabled, the quota is created but a message appears stating the feature is disabled and limits are not enforced, as shown in Example 2. See `quota` in the *Data Domain Operating System Command Reference Guide* for details.

Quotas may cause OpenStorage backup applications to report non-intuitive sizes and capacities. See Knowledge Base article 85210, available on the Support portal, for 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 storage units or, optionally, the names of all files in a specified storage unit.

To display the compression for all storage units (the original byte size, global compression, and local compression for all storage units), enter the compression option.

Note

To interrupt the output of this command, press Ctrl+C.

APPENDIX A

Checklists and Notes

This appendix includes:

- [Installation and Configuration Checklist](#)..... 76

Installation and Configuration Checklist

Configuring the Data Domain system

For details about these steps, refer to [Configuring the Data Domain System on page 25](#).

Procedure

1. Add the DD Boost License.
2. Establish the server name, username and password for the Data Domain system.

Note

The server name, user name, and password must exactly match what is used in the `send` command of Step 2 in [Configuring RMAN for Your Environment on page 76](#).

3. Enable DD Boost.
4. Create a logical storage unit. (See [Creating Storage Units on page 27](#).)
5. Enable Distributed Segment Processing.
6. Check connectivity. (Can I ping the DDR server name from the backup server?) Ensure firewall is not blocking ports.

Downloading and Installing the RMAN Plugin

For details about these steps, refer to [Configuring DD Boost for RMAN on page 35](#).

Procedure

1. Download RMAN plugin.
2. Run install/setup.

Configuring RMAN for Your Environment

For details about these steps, refer to [Configuring DD Boost for RMAN on page 35](#).

Procedure

1. `CONFIGURE CHANNEL DEVICE TYPE 'SBT_TAPE'` (See [AIX, HPUX, Linux, or UNIX Environment on page 37](#) or [Windows Environment on page 43](#). Setting the correct `PARMS` is essential here.)
2. `ALLOCATE CHANNEL` (See [Registering Each Data Domain System on page 44](#).) Set the library and enter the `send` command for server name, user name, and password). The server name, user name, and password must match exactly what is used in Step 2 in [Configuring the Data Domain system on page 76](#).
3. `CONFIGURE DEFAULT DEVICE TYPE 'SBT_TAPE'` (See [Registering Each Data Domain System on page 44](#).)

Testing the Installation by Using `sbtttest` (Optional)

`sbtttest` is a diagnostic tool provided by Oracle. Before running `sbtttest` with the DD Boost plugin, you must set the following system environment variables, using the same values that are required in the `PARMS` fields of step 1 in [Configuring RMAN for Your Environment on page 76](#).

```
ORACLE_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 <ddboost directory>/lib/libddobk.so
```

Backup Examples

```
RMAN> backup copies 2 database format '%u_%p', 'destinationssystem/%u_%p' filesperset=1;  
RMAN> backup database format '%u_%p';
```


APPENDIX B

Performing a Backup with Oracle Enterprise Manager

This appendix includes:

- [Example Screens](#).....80

Example Screens

Figure 12 Login Window

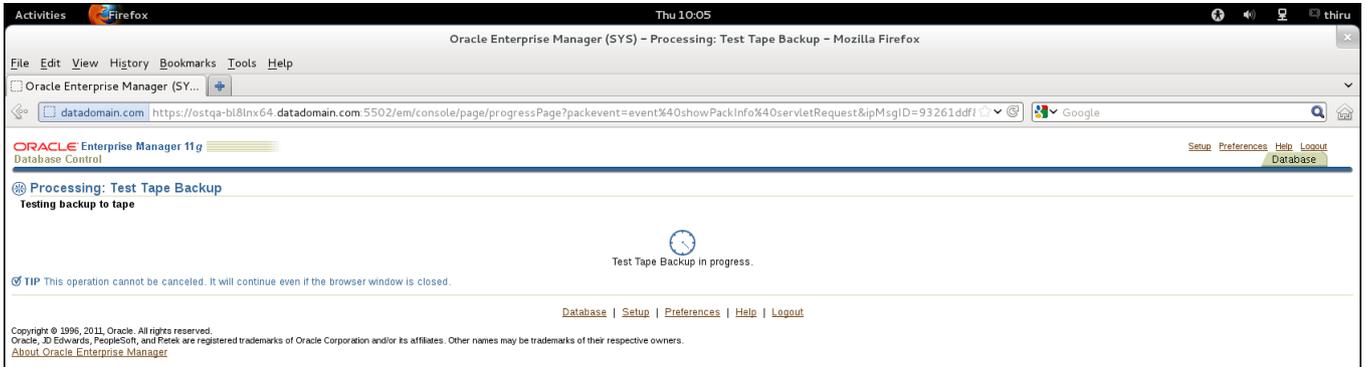


Figure 13 Database Instance Page



Figure 14 Processing Page

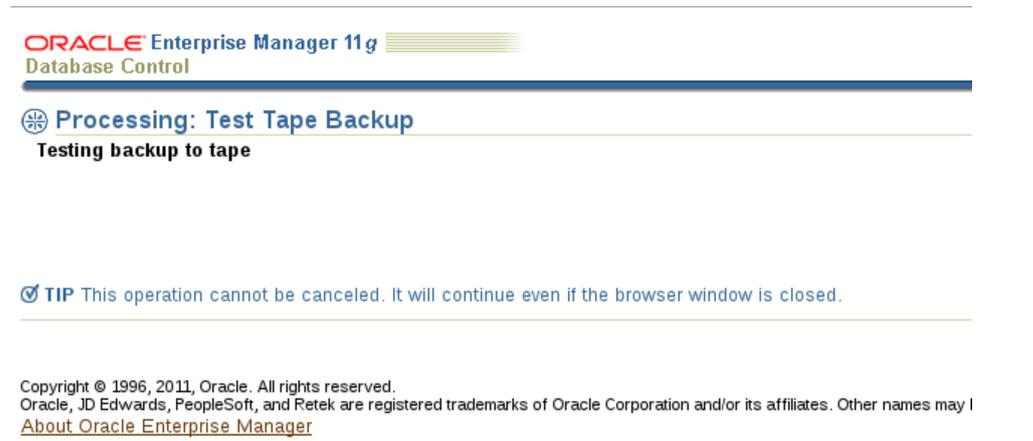


Figure 15 Schedule Customized Backup Options Page

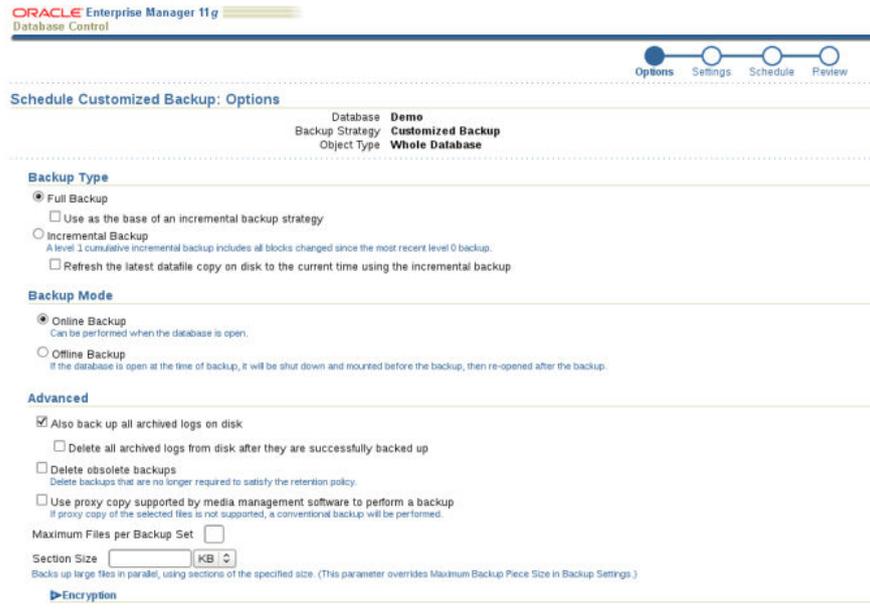


Figure 16 Schedule Customized Backup Schedule Page



Figure 17 Schedule Customized Backup Review Page

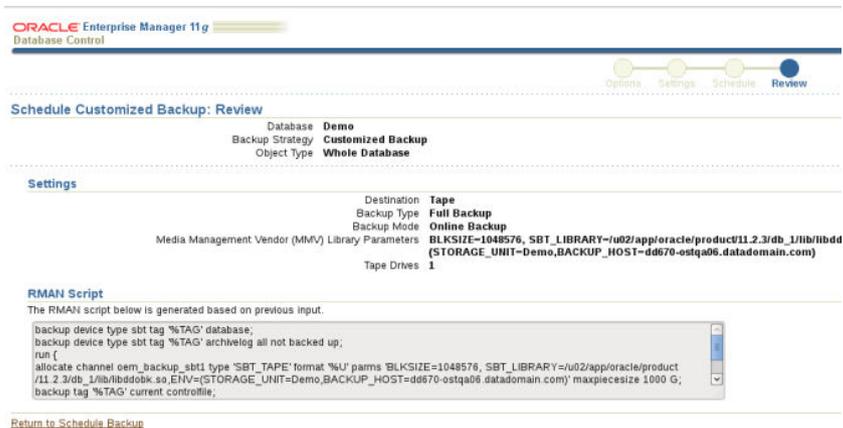


Figure 18 Execution Page

ORACLE Enterprise Manager 11g
Database Control

Execution: Demo Page Refreshed Apr 26, 2012 4:20:25 PM EDT

Summary

Status	Succeeded	Type	Database Backup
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 RMAN Plugin for Multiple SIDs Under One ORACLE_HOME

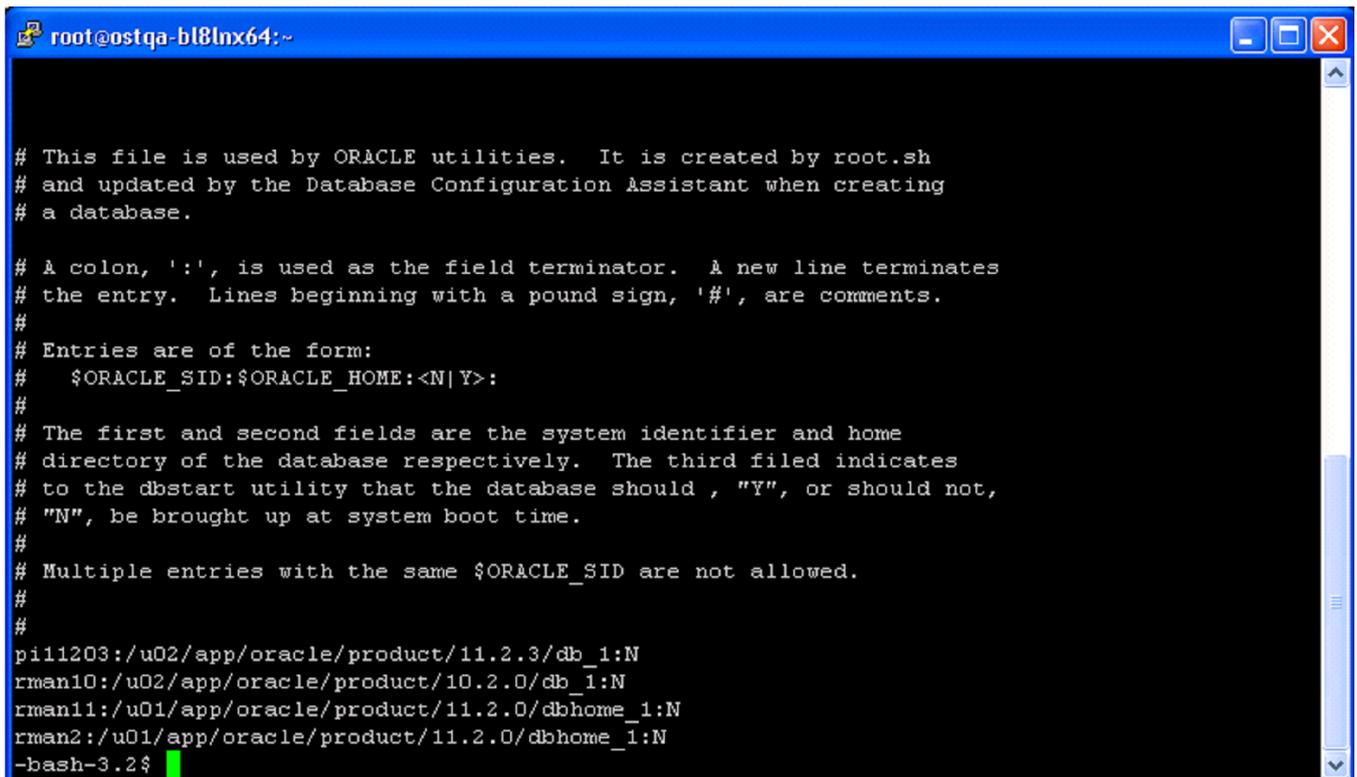
This appendix includes:

- [Installation and Configuration Examples](#)..... 84

Installation and Configuration Examples

Figure 19 on page 84 shows two SIDs, rman11 and rman2, under the same ORACLE_HOME: /u01/app/oracle/product/11.2.0/dbhome_1

Figure 19 Two SIDs Under One ORACLE_HOME



```

root@ostqa-bl8lnx64:~
# This file is used by ORACLE utilities.  It is created by root.sh
# and updated by the Database Configuration Assistant when creating
# a database.

# A colon, ':', is used as the field terminator.  A new line terminates
# the entry.  Lines beginning with a pound sign, '#', are comments.
#
# Entries are of the form:
#   $ORACLE_SID:$ORACLE_HOME:<N|Y>:
#
# The first and second fields are the system identifier and home
# directory of the database respectively.  The third field indicates
# to the dbstart utility that the database should, "Y", or should not,
# "N", be brought up at system boot time.
#
# Multiple entries with the same $ORACLE_SID are not allowed.
#
#
p11203:/u02/app/oracle/product/11.2.3/db_1:N
rman10:/u02/app/oracle/product/10.2.0/db_1:N
rman11:/u01/app/oracle/product/11.2.0/dbhome_1:N
rman2:/u01/app/oracle/product/11.2.0/dbhome_1:N
-bash-3.2$

```

Figure 20 on page 84 shows using the tar command to extract the plugin files from the tar file and running the install.sh script to install the plug-in files under ORACLE_HOME:/u01/app/oracle/product/11.2.0/dbhome_1 which serves rman11 and rman2 SIDs.

The installation completes with the message: Successfully installed the Data Domain plugin for RMAN. This serves all SIDs under this ORACLE_HOME.

Note

In the case of SIDs in multiple ORACLE_HOMES, installation would be required for each ORACLE_HOME.

Figure 20 RMAN Plugin Installed in ORACLE_HOME

```

$ tar -xvf RMAN_1.1.0.0-356444_RMAN_linux_64.tar
libddobk.so
install.sh
libccme_base.so
libcryptocme2.so
libcryptocme2.sig
libCSP-lb.so
libCSP-Security.so
license.txt

```

Figure 20 RMAN Plugin Installed in ORACLE_HOME (continued)

```

readme.txt
libDDBoost.so
$ ./install.sh
a plugin already exists, do you want to proceed with installation? (y or n)
y
Installing the Data Domain plugin for RMAN ...
Copying libraries to /u01/app/oracle/product/11.2.0/dbhome_1/lib
cp libddobk.so /u01/app/oracle/product/11.2.0/dbhome_1/lib/libddobk.so
cp libDDBoost.so /u01/app/oracle/product/11.2.0/dbhome_1/lib/libDDBoost.so
Successfully installed the Data Domain plugin for RMAN

```

[Figure 21 on page 85](#) shows the contents of the directory `ORACLE_HOME/lib`, where the RMAN plugin libraries are installed. The files in the list are the installed plugin files.

Figure 21 Plugin Libraries

```

->ls -slpd *
320 -rwxr-xr-x 1 oracle oinstall 320599 May 22 09:11 libccme_base_non_fips.so
568 -rwxr-xr-x 1 oracle oinstall 574678 May 22 09:11 libccme_base.so
  4 -rwxr-xr-x 1 oracle oinstall  2638 May 22 09:11 libcryptocme.sig
424 -rwxr-xr-x 1 oracle oinstall 426114 May 22 09:11 libcryptocme.so
152 -rwxr-xr-x 1 oracle oinstall 150760 May 22 09:11 libCSP-lb.so
 872 -rwxr-xr-x 1 oracle oinstall 888088 May 22 09:11 libCSP-Security.so
1756 -rwxr-xr-x 1 oracle oinstall 1791336 May 22 09:11 libDDBoost.so
  72 -rwxr-xr-x 1 oracle oinstall  66296 May 22 09:11 libddobk.so

```

[Figure 22 on page 85](#) shows an example of setting `ORACLE_SID`, running Recovery Manager, and registering the Data Domain system.

Set `ORACLE_SID` by executing either the `. oraenv` script or the `export` command and logging in to RMAN to register with the Data Domain systems.

In the example `ORACLE_SID` is set for `rman2`.

In this example, five different Data Domain systems are being registered for access from this client system, all with the same user name and password. For a description of this task see the section [Registering Each Data Domain System on page 44](#).

Figure 22 Registering Data Domain Systems

```

$ . oraenv
ORACLE_SID = [oracle] ? rman2
$ rman target /

Recovery Manager: Release 11.2.0.1.0 - Production on Fri Feb 24 16:54:50 2012
Copyright (c) 1982, 2009, Oracle and/or its affiliates. All rights reserved.
Connected to target database: RMAN2 (DBID=83534869)

RMAN> @setuser.rman.11g

RMAN> run

2> {allocate channel c1 device type SBT_TAPE parms 'SBT_LIBRARY=/u01/app/oracle/product
/11.2.0/dbhome_1/lib/libddobk.so'
3> send 'set username sysadmin password xxxxxxxxx servername host1.datadomain.com';
4> send 'set username sysadmin password xxxxxxxxx servername host2.datadomain.com';
5> send 'set username sysadmin password xxxxxxxxx servername host3.datadomain.com';
6> send 'set username sysadmin password xxxxxxxxx servername host4.datadomain.com';
7> send 'set username sysadmin password xxxxxxxxx servername host5.datadomain.com';
8> release channel c1;
9> }
using target database control file instead of recovery catalog
allocated channel: c1
channel c1: SID=101 device type=SBT_TAPE

```

Figure 22 Registering Data Domain Systems (continued)

```

channel c1: Data Domain Boost API
sent command to channel: c1
released channel: c1

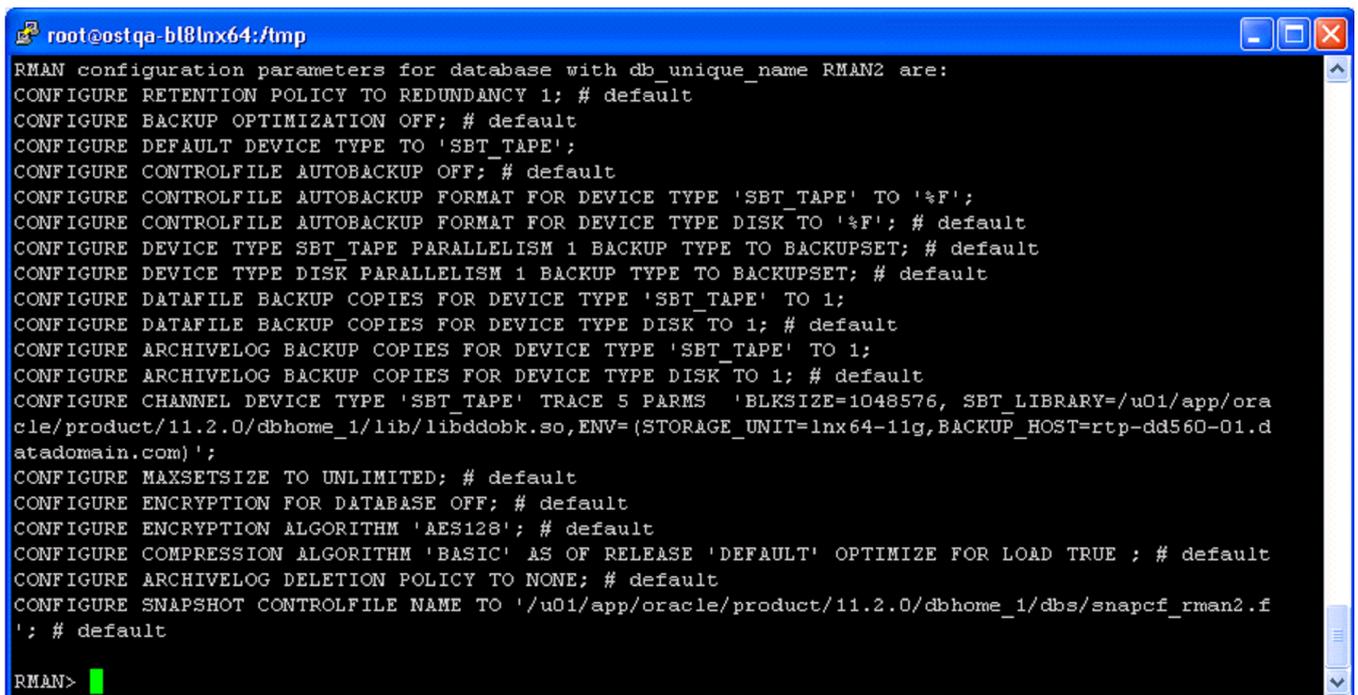
RMAN>
RMAN> **end-of-file**
RMAN>

```

[Figure 23 on page 86](#) shows the configurations for Oracle SIDs rman2 and rman11.

Figure 23 Configure RMAN for each SID (rman2 and rman11)

First configuration:



```

root@ostqa-bl8lnx64:/tmp
RMAN configuration parameters for database with db_unique_name RMAN2 are:
CONFIGURE RETENTION POLICY TO REDUNDANCY 1; # default
CONFIGURE BACKUP OPTIMIZATION OFF; # default
CONFIGURE DEFAULT DEVICE TYPE TO 'SBT_TAPE';
CONFIGURE CONTROLFILE AUTOBACKUP OFF; # default
CONFIGURE CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE 'SBT_TAPE' TO '%F';
CONFIGURE CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE DISK TO '%F'; # default
CONFIGURE DEVICE TYPE SBT TAPE PARALLELISM 1 BACKUP TYPE TO BACKUPSET; # default
CONFIGURE DEVICE TYPE DISK PARALLELISM 1 BACKUP TYPE TO BACKUPSET; # default
CONFIGURE DATAFILE BACKUP COPIES FOR DEVICE TYPE 'SBT_TAPE' TO 1;
CONFIGURE DATAFILE BACKUP COPIES FOR DEVICE TYPE DISK TO 1; # default
CONFIGURE ARCHIVELOG BACKUP COPIES FOR DEVICE TYPE 'SBT_TAPE' TO 1;
CONFIGURE ARCHIVELOG BACKUP COPIES FOR DEVICE TYPE DISK TO 1; # default
CONFIGURE CHANNEL DEVICE TYPE 'SBT_TAPE' TRACE 5 PARMS 'BLKSIZE=1048576, SBT_LIBRARY=/u01/app/oracle/product/11.2.0/dbhome_1/lib/libddobk.so, ENV=(STORAGE_UNIT=lnx64-11g, BACKUP_HOST=rtp-dd560-01.d
atadomain.com)';
CONFIGURE MAXSETSIZE TO UNLIMITED; # default
CONFIGURE ENCRYPTION FOR DATABASE OFF; # default
CONFIGURE ENCRYPTION ALGORITHM 'AES128'; # default
CONFIGURE COMPRESSION ALGORITHM 'BASIC' AS OF RELEASE 'DEFAULT' OPTIMIZE FOR LOAD TRUE ; # default
CONFIGURE ARCHIVELOG DELETION POLICY TO NONE; # default
CONFIGURE SNAPSHOT CONTROLFILE NAME TO '/u01/app/oracle/product/11.2.0/dbhome_1/dbs/snapcf_rman2.f
'; # default

RMAN>

```

Second configuration:

Figure 23 Configure RMAN for each SID (rman2 and rman11) (continued)

```

root@ostqa-bl8lnx64:~
RMAN> show all;

RMAN configuration parameters for database with db_unique_name RMAN11 are:
CONFIGURE RETENTION POLICY TO REDUNDANCY 1; # default
CONFIGURE BACKUP OPTIMIZATION OFF; # default
CONFIGURE DEFAULT DEVICE TYPE TO 'SBT_TAPE';
CONFIGURE CONTROLFILE AUTOBACKUP OFF; # default
CONFIGURE CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE 'SBT_TAPE' TO '%F';
CONFIGURE CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE DISK TO '%F'; # default
CONFIGURE DEVICE TYPE 'SBT_TAPE' PARALLELISM 1 BACKUP TYPE TO BACKUPSET;
CONFIGURE DEVICE TYPE DISK PARALLELISM 1 BACKUP TYPE TO BACKUPSET; # default
CONFIGURE DATAFILE BACKUP COPIES FOR DEVICE TYPE 'SBT_TAPE' TO 1;
CONFIGURE DATAFILE BACKUP COPIES FOR DEVICE TYPE DISK TO 1; # default
CONFIGURE ARCHIVELOG BACKUP COPIES FOR DEVICE TYPE 'SBT_TAPE' TO 1;
CONFIGURE ARCHIVELOG BACKUP COPIES FOR DEVICE TYPE DISK TO 1; # default
CONFIGURE CHANNEL DEVICE TYPE 'SBT TAPE' TRACE 5 PARMS 'BLKSIZE=1048576, SBT_LIBRARY=/u01/app/oracle/pr
oduct/11.2.0/dbhome_1/lib/libddobk.so, ENV=(STORAGE_UNIT=lnx64-11g, BACKUP_HOST=rtp-dd560-01.datadomain.co
m)';
CONFIGURE MAXSETSIZE TO UNLIMITED; # default
CONFIGURE ENCRYPTION FOR DATABASE OFF; # default
CONFIGURE ENCRYPTION ALGORITHM 'AES128'; # default
CONFIGURE COMPRESSION ALGORITHM 'BASIC' AS OF RELEASE 'DEFAULT' OPTIMIZE FOR LOAD TRUE ; # default
CONFIGURE ARCHIVELOG DELETION POLICY TO NONE; # default
CONFIGURE SNAPSHOT CONTROLFILE NAME TO '/u01/app/oracle/product/11.2.0/dbhome_1/dbs/snapcf_rman11.f'; #
default

RMAN>

```

Figure 24 Taking Backups from two SIDs, rman2 and rman11

First backup:

```

ostqa-ai7-1.datadomain.com - PuTTY
-bash-3.2$ echo $ORACLE_SID
rman11g
-bash-3.2$ rman target /

Recovery Manager: Release 11.2.0.3.0 - Production on Tue May 21 12:28:16 2013

Copyright (c) 1982, 2011, Oracle and/or its affiliates. All rights reserved.

connected to target database: RMAN11G (DBID=2497026946)

RMAN> show all;

using target database control file instead of recovery catalog
RMAN configuration parameters for database with db_unique_name RMAN11G are:
CONFIGURE RETENTION POLICY TO REDUNDANCY 1;
CONFIGURE BACKUP OPTIMIZATION OFF; # default
CONFIGURE DEFAULT DEVICE TYPE TO 'SBT_TAPE';
CONFIGURE CONTROLFILE AUTOBACKUP OFF; # default
CONFIGURE CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE 'SBT_TAPE' TO '%F';
CONFIGURE CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE DISK TO '%F'; # default
CONFIGURE DEVICE TYPE 'SBT_TAPE' PARALLELISM 1 BACKUP TYPE TO BACKUPSET;
CONFIGURE DEVICE TYPE DISK PARALLELISM 1 BACKUP TYPE TO BACKUPSET; # default
CONFIGURE DATAFILE BACKUP COPIES FOR DEVICE TYPE 'SBT_TAPE' TO 1;
CONFIGURE DATAFILE BACKUP COPIES FOR DEVICE TYPE DISK TO 2;
CONFIGURE ARCHIVELOG BACKUP COPIES FOR DEVICE TYPE 'SBT_TAPE' TO 1;
CONFIGURE ARCHIVELOG BACKUP COPIES FOR DEVICE TYPE DISK TO 2;
CONFIGURE CHANNEL DEVICE TYPE 'SBT TAPE' TRACE 5 PARMS 'BLKSIZE=1048576, SBT_LIBRARY=/u01/app/oracle/product/11.2.0/dbhome_1/lib/libdd
obk.so, ENV=(STORAGE_UNIT=oracle, BACKUP_HOST=ostqa-dd890-23.datadomain.com, ORACLE_HOME=/u01/app/oracle/product/11.2.0/dbhome_1)';
CONFIGURE MAXSETSIZE TO UNLIMITED;
CONFIGURE ENCRYPTION FOR DATABASE OFF;
CONFIGURE ENCRYPTION ALGORITHM 'AES128'; # default
CONFIGURE COMPRESSION ALGORITHM 'BASIC' AS OF RELEASE 'DEFAULT' OPTIMIZE FOR LOAD TRUE ; # default
CONFIGURE ARCHIVELOG DELETION POLICY TO NONE; # default
CONFIGURE SNAPSHOT CONTROLFILE NAME TO '/u01/app/oracle/product/11.2.0/dbhome_3/dbs/snapcf_rman11g.f'; # default

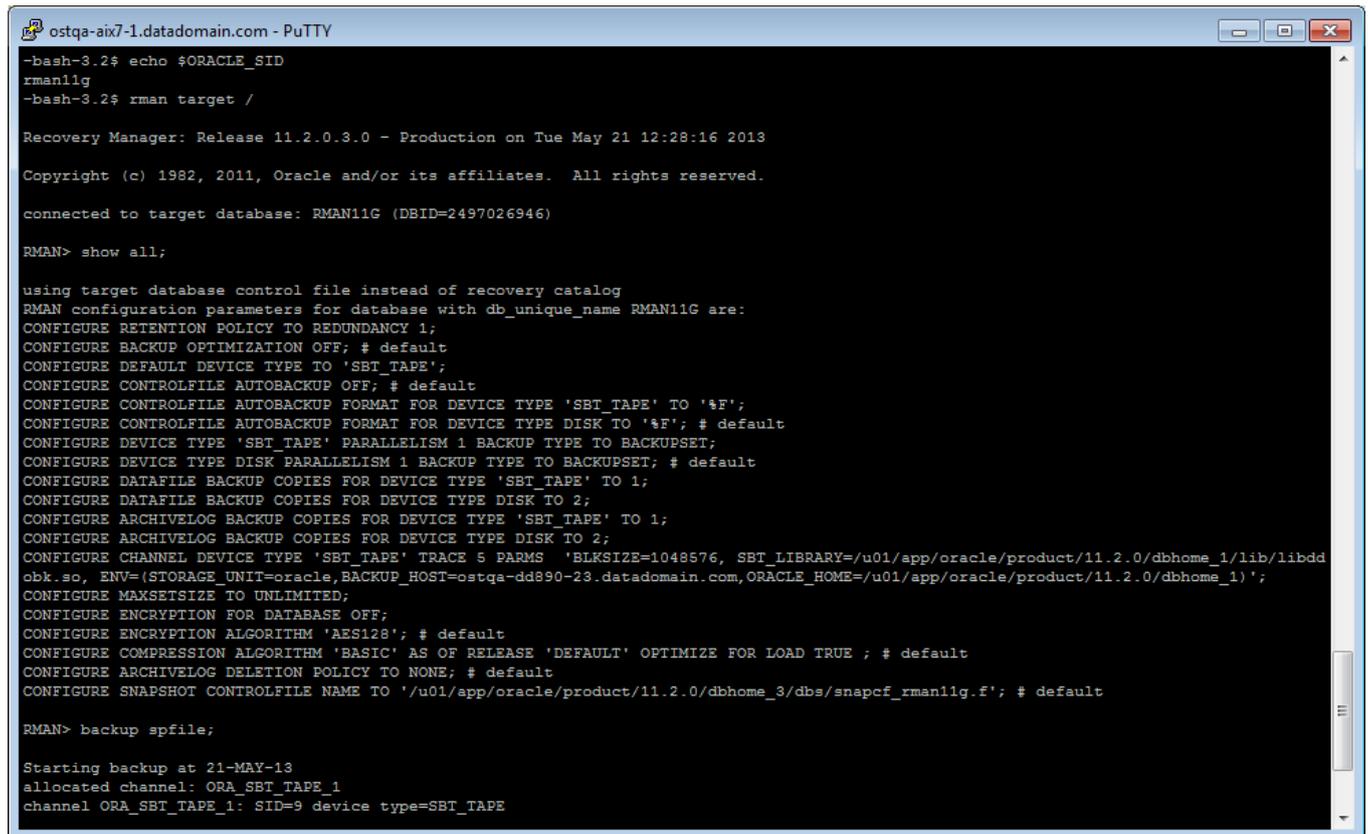
RMAN> backup spfile;

Starting backup at 21-MAY-13
allocated channel: ORA_SBT_TAPE_1
channel ORA_SBT_TAPE_1: SID=9 device type=SBT_TAPE

```

Figure 24 Taking Backups from two SIDs, rman2 and rman11 (continued)

Second backup:



```

ostqa-ai7-1.datadomain.com - PuTTY
-bash-3.2$ echo $ORACLE_SID
rman11g
-bash-3.2$ rman target /

Recovery Manager: Release 11.2.0.3.0 - Production on Tue May 21 12:28:16 2013

Copyright (c) 1982, 2011, Oracle and/or its affiliates. All rights reserved.

connected to target database: RMAN11G (DBID=2497026946)

RMAN> show all;

using target database control file instead of recovery catalog
RMAN configuration parameters for database with db_unique_name RMAN11G are:
CONFIGURE RETENTION POLICY TO REDUNDANCY 1;
CONFIGURE BACKUP OPTIMIZATION OFF; # default
CONFIGURE DEFAULT DEVICE TYPE TO 'SBT_TAPE';
CONFIGURE CONTROLFILE AUTOBACKUP OFF; # default
CONFIGURE CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE 'SBT_TAPE' TO '%F';
CONFIGURE CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE DISK TO '%F'; # default
CONFIGURE DEVICE TYPE 'SBT_TAPE' PARALLELISM 1 BACKUP TYPE TO BACKUPSET;
CONFIGURE DEVICE TYPE DISK PARALLELISM 1 BACKUP TYPE TO BACKUPSET; # default
CONFIGURE DATAFILE BACKUP COPIES FOR DEVICE TYPE 'SBT_TAPE' TO 1;
CONFIGURE DATAFILE BACKUP COPIES FOR DEVICE TYPE DISK TO 2;
CONFIGURE ARCHIVELOG BACKUP COPIES FOR DEVICE TYPE 'SBT_TAPE' TO 1;
CONFIGURE ARCHIVELOG BACKUP COPIES FOR DEVICE TYPE DISK TO 2;
CONFIGURE CHANNEL DEVICE TYPE 'SBT_TAPE' TRACE 5 PARMS 'BLKSIZE=1048576, SBT_LIBRARY=/u01/app/oracle/product/11.2.0/dbhome_1/lib/libdd
obk.so, ENV=(STORAGE_UNIT=oracle,BACKUP_HOST=ostqa-dd890-23.datadomain.com,ORACLE_HOME=/u01/app/oracle/product/11.2.0/dbhome_1)';
CONFIGURE MAXSETSIZE TO UNLIMITED;
CONFIGURE ENCRYPTION FOR DATABASE OFF;
CONFIGURE ENCRYPTION ALGORITHM 'AES128'; # default
CONFIGURE COMPRESSION ALGORITHM 'BASIC' AS OF RELEASE 'DEFAULT' OPTIMIZE FOR LOAD TRUE ; # default
CONFIGURE ARCHIVELOG DELETION POLICY TO NONE; # default
CONFIGURE SNAPSHOT CONTROLFILE NAME TO '/u01/app/oracle/product/11.2.0/dbhome_3/dbs/snapcf_rman11g.f'; # default

RMAN> backup spfile;

Starting backup at 21-MAY-13
allocated channel: ORA_SBT_TAPE_1
channel ORA_SBT_TAPE_1: SID=9 device type=SBT_TAPE

```