

Dell EMC Data Domain Boost for Enterprise Applications and ProtectPoint Database Application Agent

Version 4.0

Installation and Administration Guide

302-003-659

REV 06

Copyright © 2013-2018 Dell Inc. or its subsidiaries. All rights reserved.

Published July 2018

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

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS-IS." DELL MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. USE, COPYING, AND DISTRIBUTION OF ANY DELL SOFTWARE DESCRIBED IN THIS PUBLICATION REQUIRES AN APPLICABLE SOFTWARE LICENSE.

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

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

CONTENTS

Figures		11
Tables		13
Preface		15
Chapter 1	Product Overview	21
	Terminology used in this guide.....	22
	Introduction to the product.....	22
	DD Boost backups and restores.....	23
	ProtectPoint backups and restores on a VMAX system.....	24
	ProtectPoint with RecoverPoint backups and restores on an XtremIO system.....	28
	Supported configurations.....	34
	High-availability configurations.....	35
	Data Domain High Availability (HA).....	36
	Virtualization support.....	37
	Data Domain replication.....	37
	VMAX replication.....	37
	Usage limits on Data Domain resources.....	42
	Database application agent operations with Data Domain usage limits on capacity.....	43
	Database application agent operations with Data Domain usage limits on streams.....	44
	Database application agent ProtectPoint operations with Data Domain usage limits.....	44
	Road map for product operations.....	46
Chapter 2	Data Domain System Configuration for DD Boost Operations	49
	Licensing the Data Domain system.....	50
	Enabling DD Boost on a Data Domain system.....	50
	Changing the DD Boost access rights.....	51
	Enabling encryption over a WAN connection.....	52
	Enabling the DD Boost operations through a firewall.....	52
	Setting up the storage units.....	53
	Enabling the distributed segment processing.....	53
	Enabling the advanced load balancing and link failover.....	54
	Enabling the DD Boost-over-FC service.....	55
	DD Boost-over-FC path management.....	57
	Validating and troubleshooting the database and Data Domain system connection.....	57
Chapter 3	Product Installation	59
	Road map to install or update the software.....	60
	AIX installation.....	62

Installing the software on AIX.....	62
Uninstalling the software on AIX.....	65
HP-UX installation.....	65
Installing the software on HP-UX.....	65
Uninstalling the software on HP-UX.....	67
Linux installation.....	67
Installing the software on Linux.....	68
Uninstalling the software on Linux.....	69
Solaris installation.....	69
Installing the software on Solaris.....	70
Uninstalling the software on Solaris.....	71
Microsoft Windows installation.....	71
Installing the software on Windows.....	72
Uninstalling the software on Windows.....	73
Software components.....	74
Software links created during installation.....	75

Chapter 4	Product Configuration	77
	Road map for configuration.....	78
	Setting up the configuration file.....	78
	Syntax rules for the configuration file.....	79
	Common parameters.....	80
	Common ProtectPoint parameters for VMAX.....	83
	Common ProtectPoint with RecoverPoint parameters for XtremIO..	85
	Configuring product operations over FC and IP networks.....	88
	Configuring the optimization of ProtectPoint backups for third-party	
	multipathing software.....	89
	Configuring restores of replicated backups.....	90
	Configuring restores of replicated ProtectPoint backups.....	94
	Configuring ProtectPoint VMAX restores directly from Data Domain.....	96
	Configuring ProtectPoint VMAX restores from local snapshots.....	96
	Configuring rollback restores of ProtectPoint backups.....	97
	Configuring usage limits on Data Domain resources.....	101
	Configuring usage quota on Data Domain capacity.....	101
	Configuring usage limits on Data Domain streams.....	102
	Configuring the lockbox.....	103
	Lockbox requirements.....	103
	Configuring the lockbox with the ddbmadmin command.....	104
	Configuring the display and deletion of save set information.....	117
	Using the ddbmadmin command to display and delete save sets...118	
	Using the ddbmadmin command to display clients for a device path	
	121
	Using the ddbmadmin command to display save set information...121	
	Using the ddbmadmin command to display save file information..	122
	Using the ddbmadmin command to delete save sets.....	124
	Using the ddbmadmin command to upgrade the backup index.....	125
	Configuring the use of Data Domain Cloud Tier for data movement to the	
	cloud.....	126
	Setting up the DD Cloud Tier policy for data movement to the cloud	
	127
	Performing the data movement to the cloud.....	130
	Performing the data recall from the cloud.....	130
	General troubleshooting tips.....	132
	Debug log settings.....	132

	Debug log files.....	133
	Backup or restore fails due to an inaccessible lockbox.....	134
	Command ddbmadmin -P encounters a conflict with an installed application.....	135
	Lockbox creation might fail on an NFS/CIFS share.....	135
	Lockbox creation procedure when UAC is enabled on Windows...	135
	Major system update can produce an error about lockbox stable value threshold.....	136
	Restore fails if the CLIENT parameter setting does not match the backup client name.....	136
	ProtectPoint specific troubleshooting tips.....	137
	ProtectPoint operation might fail due to configuration issues.....	137
	Restore might fail due to an incompatible file system or volume manager version.....	137
	Troubleshooting the ProtectPoint for VMAX restores.....	138
	Manual cleanup of FAST.X restore devices after a failed restore of a ProtectPoint for VMAX backup.....	138
Chapter 5	DD Boost Operations on DB2 Systems	141
	Overview of DD Boost operations in a DB2 environment.....	142
	DD Boost DB2 backup processes.....	142
	DD Boost DB2 restore processes.....	143
	DB2 backups of transaction logs.....	143
	Configuration of DD Boost operations in a DB2 environment.....	144
	Integrating the product into the DB2 environment.....	144
	Configuring the DB2 parameters for DD Boost operations.....	144
	Configuring DB2 transaction log archiving.....	146
	Configuring DB2 backup deletion.....	147
	Preventing deletion of DB2 backup images.....	149
	Estimating the Data Domain resource usage on DB2 systems.....	149
	Performing DD Boost backups and recovery with the DB2 CLP.....	150
	Performing DB2 backups with the DB2 CLP.....	150
	Performing DB2 restores with the DB2 CLP.....	154
	Performing DB2 recovery with the DB2 CLP.....	156
	Performing DD Boost backups and restores with the DB2 GUI.....	158
	Performing DD Boost backup data recovery with utility programs.....	158
	Retrieving DB2 database backups and log backups with the ddbmdb2adutil utility.....	159
	Recovering DB2 backup data with the IBM HPU utility.....	164
	Preparing for DB2 disaster recovery.....	167
	DB2 DPF requirements for DD Boost operations.....	168
	DB2 HADR requirements for DD Boost operations.....	168
	DB2 pureScale requirements for DD Boost operations.....	170
	Performing DD Boost backups in a DB2 pureScale environment...	171
	Performing DD Boost restores in a DB2 pureScale environment...	171
	Restoring between a DB2 pureScale instance and Enterprise Server Edition.....	171
	Backups and restores of transaction logs in a DB2 pureScale environment.....	173
	Deleting DD Boost backups in a DB2 pureScale environment.....	173
	DB2 troubleshooting tips for DD Boost operations.....	174
	DB2 multistream restore and rollforward might fail on AIX with DFC.....	174
	DB2 issue with local hostname resolution.....	174
	DB2 issues with logarchopt n setting.....	175

DB2 pruning issues with AUTO_DEL_REC_OBJ.....177
 DB2 issues due to incorrect log retrieval..... 178
 Database backup might fail when run concurrently with backups of
 a high number of archived logs.....179
 DB2 operation might generate empty debug logs on Windows..... 179

Chapter 6 ProtectPoint Operations on DB2 Systems 181

Overview of ProtectPoint operations in a DB2 environment..... 182
 ProtectPoint DB2 backup processes..... 182
 ProtectPoint DB2 restore processes..... 183
 DB2 backups of transaction logs..... 183
 Configuration of ProtectPoint operations in a DB2 environment..... 183
 Configuring the DB2 parameters for ProtectPoint operations..... 184
 Configuring DB2 transaction log archiving..... 185
 Preparing for DB2 redirected rollback restores of ProtectPoint for
 VMAX backups..... 185
 Preparing for DB2 ProtectPoint with RecoverPoint backups and
 rollback restores.....186
 Performing ProtectPoint backups and recovery with the DB2 CLP.....187
 Performing ProtectPoint backups with the DB2 CLP..... 187
 Performing ProtectPoint restores with the DB2 CLP..... 188
 Performing DB2 recovery with the DB2 CLP..... 189
 Managing and deleting ProtectPoint DB2 backups..... 189
 Querying ProtectPoint DB2 backups.....190
 Deleting ProtectPoint DB2 backups..... 190
 Preparing for DB2 disaster recovery..... 190
 DB2 DPF requirements for ProtectPoint operations..... 192
 File system requirements for ProtectPoint operations in a DPF
 environment..... 192
 Configuration requirements for ProtectPoint operations in a DPF
 environment..... 194
 Performing ProtectPoint backups in a DPF environment..... 194
 Performing ProtectPoint restores in a DPF environment..... 195
 Performing query and deletion operations in a DPF environment....
 196
 DB2 HADR requirements for ProtectPoint operations.....197
 DB2 pureScale requirements for ProtectPoint operations.....199
 Overview of ProtectPoint backups and restores of a DB2 pureScale
 database..... 199
 Configuration requirements for ProtectPoint operations in a DB2
 pureScale environment..... 200
 Performing ProtectPoint backups in a DB2 pureScale environment..
 202
 Performing ProtectPoint restores in a DB2 pureScale environment..
 203
 Redirected rollback restores of ProtectPoint for VMAX backups to
 alternate LUNs in a different cluster.....205
 Backups and restores of transaction logs in a DB2 pureScale
 environment..... 210
 Deleting ProtectPoint backups in a DB2 pureScale environment..210
 DB2 troubleshooting tips for ProtectPoint operations..... 210
 DB2 ProtectPoint restore might fail with DB2 error code SQL2081N
 211

Chapter 7	DD Boost Operations on Oracle Systems	213
	Overview of DD Boost operations in an Oracle environment.....	214
	Oracle backup processes.....	214
	Oracle restore processes.....	215
	Oracle backups of archived redo logs.....	216
	Configuration of DD Boost operations in an Oracle environment.....	216
	Setting up the configuration file in an Oracle environment.....	216
	Creating the RMAN scripts for DD Boost Oracle operations.....	217
	Configuring operations in an Oracle Data Guard environment.....	218
	Setting up Oracle Optimized Deduplication	219
	Estimating the Data Domain resource usage on Oracle systems...219	
	Migrating an Oracle configuration from DD Boost for RMAN 1.x or later..	220
	Updating the RMAN scripts used with DD Boost for RMAN 1.x or later.....	221
	Using the correct RMAN script for restore operations.....	222
	Performing DD Boost backups and restores with Oracle RMAN.....	224
	Performing DD Boost backups and restores with Oracle Enterprise Manager	224
	Performing backups and restores of Oracle CDBs and PDBs.....	225
	Performing Oracle backup deletion and maintenance operations.....	225
	Preparing for Oracle disaster recovery.....	226
	Oracle RAC and active-passive cluster requirements for DD Boost operations.....	227
	Oracle troubleshooting tips for DD Boost operations.....	227
Chapter 8	ProtectPoint Operations on Oracle Systems	229
	Overview of ProtectPoint operations in an Oracle environment.....	230
	ProtectPoint Oracle backup processes.....	230
	ProtectPoint Oracle restore processes.....	231
	ProtectPoint Oracle backups of archived redo logs.....	231
	Configuration of ProtectPoint operations in an Oracle environment.....	231
	Setting up the configuration file in an Oracle environment.....	232
	Creating the RMAN scripts for ProtectPoint operations.....	233
	Preparing for restore of archived logs.....	236
	Preparing the Data Domain device for restore on Windows.....	237
	Preparing for Oracle ProtectPoint with RecoverPoint backups and rollback restores that use RecoverPoint pre-5.0.....	237
	Configuring operations in an Oracle Data Guard environment.....	238
	Performing ProtectPoint backups and restores with Oracle RMAN.....	238
	Performing ProtectPoint backups and restores with Oracle Enterprise Manager.....	239
	Performing backups and restores of Oracle CDBs and PDBs.....	240
	Performing Oracle backup deletion and maintenance operations.....	240
	Preparing for Oracle disaster recovery.....	241
	Oracle RAC and active-passive cluster requirements for ProtectPoint operations.....	242
	ProtectPoint restore and rollback for VCS on Solaris.....	242
	Performing a ProtectPoint VCS restore.....	243
	Performing a ProtectPoint VCS rollback.....	245
	Oracle troubleshooting tips for ProtectPoint operations.....	248
	Oracle rollback restore to a new database might fail when OMF is enabled.....	249
Chapter 9	DD Boost Operations on SAP HANA Systems	251

Overview of DD Boost operations in an SAP HANA environment.....	252
SAP HANA backup processes.....	253
SAP HANA restore processes.....	253
SAP HANA backups of redo logs.....	253
Configuration of DD Boost operations in an SAP HANA environment.....	254
Integrating the product into the SAP HANA environment.....	254
Configuring the SAP HANA parameters.....	254
Configuring support of SAP HANA 2.0 SPS 00.....	255
Enabling the configuration file in SAP HANA Studio.....	256
Configuring automatic backups of SAP HANA redo logs.....	257
Estimating the Data Domain resource usage on SAP HANA systems	257
Performing DD Boost backups, recovery, and deletion with SAP HANA Studio.....	259
Performing DD Boost backups by using SAP HANA Studio.....	259
Performing DD Boost restore and recovery by using SAP HANA Studio.....	260
Deleting DD Boost backups by using SAP HANA Studio.....	263
Performing DD Boost backups and recovery with SAP HANA CLI.....	265
Performing DD Boost backups with the SAP HANA CLI.....	265
Canceling DD Boost backups with the SAP HANA CLI.....	266
Checking DD Boost backups with the SAP HANA CLI.....	266
Performing DD Boost recovery with the SAP HANA CLI.....	267
Preparing for SAP HANA disaster recovery.....	267
SAP HANA scale-out requirements for DD Boost operations.....	268
SAP HANA troubleshooting tips for DD Boost operations.....	269
Limitation in dynamic tiering support with SAP HANA.....	269
Limitations in support of SAP HANA 1.0 SPS 09.....	269
Limitations in support of SAP HANA 2.0 SPS 00.....	269
Chapter 10	DD Boost Operations on SAP with Oracle Systems
	271
Overview of DD Boost operations in an SAP with Oracle environment.....	272
SAP with Oracle backup processes.....	273
SAP with Oracle restore processes.....	273
Configuration of DD Boost operations in an SAP with Oracle environment.... 274	
Confirming the environment and file permissions.....	274
Enabling administrator privileges for SAP with Oracle restores on Windows.....	274
Configuring the DD Boost operations with the backint utility.....	275
Configuring the DD Boost operations with Oracle RMAN.....	279
Estimating the Data Domain resource usage on SAP with Oracle systems.....	282
Performing DD Boost backups and recovery with SAP BR*Tools.....	283
Performing DD Boost backups with BR*Tools.....	283
Performing DD Boost restore and recovery with BR*Tools.....	283
Preparing for SAP with Oracle disaster recovery.....	284
Restoring the required Oracle and SAP BR*Tools files.....	285
Recovering an SAP Oracle database after disaster.....	286
SAP with Oracle RAC and cluster requirements for DD Boost operations.... 286	
Active-passive cluster requirements.....	286
Oracle RAC requirements.....	286
SAP with Oracle troubleshooting tips for DD Boost operations.....	287

Chapter 11	ProtectPoint Operations on SAP with Oracle Systems	289
	Overview of ProtectPoint operations in an SAP with Oracle environment....	290
	SAP with Oracle backup processes.....	290
	SAP with Oracle restore processes.....	291
	Configuration of ProtectPoint operations in an SAP with Oracle environment.....	291
	Integrating the product into the BR*Tools environment.....	292
	Confirming the environment and file permissions.....	294
	Enabling administrator privileges for SAP with Oracle restores on Windows.....	294
	Configuring the SAP with Oracle parameters.....	294
	Preparing for restore of archived logs.....	298
	Preparing the Data Domain device for restore on Windows.....	298
	Preparing for rollback restores of SAP with Oracle ProtectPoint backups.....	298
	Preparing for SAP with Oracle ProtectPoint with RecoverPoint backups and rollback restores that use RecoverPoint pre-5.0.....	299
	Performing ProtectPoint backups and recovery with SAP BR*Tools.....	300
	Performing ProtectPoint backups with BR*Tools.....	300
	Performing ProtectPoint restore and recovery with BR*Tools....	301
	Preparing for SAP with Oracle disaster recovery.....	302
	Restoring the required Oracle and SAP BR*Tools files.....	302
	Recovering an SAP Oracle database after disaster.....	303
	SAP with Oracle RAC and cluster requirements for ProtectPoint operations.....	303
	Active-passive cluster requirements.....	304
	Oracle RAC requirements.....	304
	ProtectPoint restore and rollback for VCS on Solaris.....	304
	Performing a ProtectPoint VCS restore.....	304
	Performing a ProtectPoint VCS rollback.....	307
	SAP with Oracle troubleshooting tips for ProtectPoint operations.....	310
Appendix A	Performance Optimization	311
	Backup and recovery performance optimization.....	312
	Hardware component 70 percent rule.....	312
	Impact of software components on performance.....	312
	Performance optimization in DB2 systems.....	313
	Performance optimization in Oracle systems.....	314
	Performance optimization in SAP HANA systems.....	314
	Performance optimization in SAP with Oracle systems.....	315
Glossary		317

CONTENTS

FIGURES

1	ProtectPoint database application agent environment.....	25
2	ProtectPoint backup workflow.....	27
3	ProtectPoint with RecoverPoint environment.....	30
4	ProtectPoint with RecoverPoint backup workflow.....	32
5	ProtectPoint with RecoverPoint restore workflow.....	34
6	DD Boost for Enterprise Applications in a stand-alone configuration.....	35
7	ProtectPoint backup to a secondary Data Domain in an SRDF configuration.....	38
8	ProtectPoint backup to a primary or secondary Data Domain in an SRDF configuration	39
9	SRDF/Metro supported topology.....	41
10	Database file system layout in a ProtectPoint DPF environment.....	193
11	Target file system layout requirements for a redirected rollback restore to an alternate pureScale cluster.....	210
12	Scheduled backup settings in Oracle Enterprise Manager.....	225
13	Scheduled backup settings in Oracle Enterprise Manager.....	240
14	Specifying the configuration file in SAP HANA Studio.....	256
15	Configuring automatic log backups in SAP HANA Studio.....	257
16	Specifying backup settings in SAP HANA Studio.....	260
17	Specifying the recovery type in SAP HANA Studio.....	261
18	Locating the log backups in SAP HANA Studio.....	262
19	Selecting the data backup in SAP HANA Studio.....	263
20	Specifying settings to delete a backup in SAP HANA Studio.....	264
21	Specifying settings to delete the older backups of a backup in SAP HANA Studio...	265

FIGURES

TABLES

1	Revision history.....	15
2	Style conventions.....	18
3	Network connection types in a ProtectPoint environment.....	26
4	Network connection types in a ProtectPoint with RecoverPoint environment.....	31
5	Software installation directories on AIX.....	63
6	Software installation directories on HP-UX.....	66
7	Software installation directories on Linux.....	68
8	Software installation directories on Solaris.....	70
9	Product software components.....	74
10	Common parameters.....	80
11	Common ProtectPoint parameters for VMAX.....	83
12	Common ProtectPoint with RecoverPoint parameters for XtremIO.....	86
13	Parameters for operations over FC networks.....	89
14	Parameters for restores from a secondary Data Domain system.....	90
15	Options of the ddbmadmin command for lockbox operations.....	105
16	Examples of noninteractive ddbmadmin commands.....	105
17	The ddbmadmin command options for save set display and deletion.....	119
18	Parameters for debugging.....	132
19	DB2 parameters for DD Boost operations.....	144
20	Options of the ddbmdb2adutil utility for backup image retrieval.....	161
21	Example values for host entry in system configuration file.....	175
22	DB2 parameter for ProtectPoint operations.....	185
23	SAP HANA parallelism parameter.....	254
24	SAP with Oracle parameters for DD Boost operations with backint.....	277
25	SAP with Oracle parameters for ProtectPoint operations.....	295

TABLES

Preface

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

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

Note

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

Purpose

This document describes how to install, configure, and use the Data Domain Boost for Enterprise Applications and ProtectPoint database application agent version 4.0.

Audience

This document is intended for database administrators (DBAs) or system administrators who are responsible for installing and maintaining backup and recovery systems for databases or applications.

Users of this guide must be familiar with the following topics:

- Backup, recovery, database, applications, and network terminology
- Backup and recovery procedures
- Disaster recovery procedures

Revision history

The following table presents the revision history of this document.

Table 1 Revision history

Revision	Date	Description
06	July 18, 2018	Updated the following topics: <ul style="list-style-type: none">• Manual cleanup of FAST.X restore devices after a failed restore of a ProtectPoint for VMAX backup —Added this topic at the end of Chapter 4.• Overview of ProtectPoint operations in a DB2 environment—Deleted the Note about nonsupport of the cancellation of an ongoing ProtectPoint for VMAX restore that uses FAST.X restore LUNs.• Overview of ProtectPoint operations in an Oracle environment—Deleted the Note about nonsupport of the cancellation of an ongoing ProtectPoint for VMAX restore that uses FAST.X restore LUNs.• Overview of ProtectPoint operations in an SAP with Oracle environment—Deleted the Note about

Table 1 Revision history (continued)

Revision	Date	Description
		nonsupport of the cancellation of an ongoing ProtectPoint for VMAX restore that uses FAST.X restore LUNs.
05	April 23, 2018	<p>Updated the following topics:</p> <ul style="list-style-type: none"> • Introduction to the product—Updated the paragraph about support of in-flight encryption for DD Boost clients. • Enable encryption over a WAN connection—Updated the details about how to enable encryption over a WAN connection. • Road map to install or update the software—Before step 1, added the paragraph about the support of coexistence with the NetWorker client. • Overview of ProtectPoint operations in a DB2 environment—Added the Note about nonsupport of the cancellation of an ongoing ProtectPoint for VMAX restore that uses FAST.X restore LUNs. • Overview of ProtectPoint operations in an Oracle environment—Added the Note about nonsupport of the cancellation of an ongoing ProtectPoint for VMAX restore that uses FAST.X restore LUNs. • Overview of ProtectPoint operations in an SAP with Oracle environment—Added the Note about nonsupport of the cancellation of an ongoing ProtectPoint for VMAX restore that uses FAST.X restore LUNs.
04	December 15, 2017	<p>Added the following topic in Chapters 7 and 8:</p> <ul style="list-style-type: none"> • Performing backups and restores of Oracle CDBs and PDBs
03	November 17, 2017	<p>Updated the following information:</p> <ul style="list-style-type: none"> • "Road map to install or update the software"—In the second paragraph of the Note, added details about DB2 archived log backup configuration. • "DB2 backups of transaction logs"—In the final Note, added the second paragraph about restarting or reactivating the DB2 database after a product software update or deployment of a new vendor library. • "Configuring DB2 transaction log archiving"—Before the procedural steps, updated the first paragraph and added details on how to deactivate, activate, and list the active DB2 databases. • "Configuring DB2 backup deletion"—Before the procedural steps, added details to ensure that the

Table 1 Revision history (continued)

Revision	Date	Description
		<p>DB2 backup history and backup configuration and storage are synchronized and the backup object removal succeeds.</p> <ul style="list-style-type: none"> • "Preparing for rollback restores of SAP with Oracle ProtectPoint backups "—Added this topic about specifying a new location for the SAPBACKUP directory prior to a rollback restore. • Limitation in dynamic tiering support with SAP HANA—Updated this troubleshooting topic based on the new support of SAP HANA 2.0 SPS 02.
02	August 14, 2017	<p>Updated the following information:</p> <ul style="list-style-type: none"> • Example Linux script to create a policy on page 129 —In the Note, added the second paragraph about copy and paste operations. In the example script, changed <code>printf \$2;</code> to <code>print \$2;</code>. • Performing the data movement to the cloud on page 130—Added the <code>-c</code> option in the first <code>ddbadmin</code> command line. In the Note, added the first paragraph about optimal Data Domain performance. Corrected the <code>ddbadmin</code> spelling in the last command line. • Performing the data recall from the cloud on page 130—Added the <code>-c</code> option in the first <code>ddbadmin</code> command line. Corrected the <code>ddbadmin</code> spelling in the last command line.
01	June 30, 2017	Initial release of this guide for the Data Domain Boost for Enterprise Applications and ProtectPoint database application agent 4.0.

Related documentation

You can find additional publications for this product release and related products at the Support website.

The online software compatibility guide, available at <http://compatibilityguide.emc.com:8080/CompGuideApp/>, provides details about supported environments and platforms.

Special notice conventions that are used in this document

The following conventions are used for special notices:

NOTICE

Identifies content that warns of potential business or data loss.

Note

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

Typographical conventions

The following type style conventions are used in this document:

Table 2 Style conventions

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

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

Where to find product documentation

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

Where to get support

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

To access a product specific Support page:

1. Go to <https://support.emc.com/products>.
2. In the **Find a Product by Name** box, type a product name, and then select the product from the list that appears.
3. Click the following button:



4. (Optional) To add the product to **My Saved Products**, in the product specific page, click **Add to My Saved Products**.

Knowledgebase

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

To search the Knowledgebase:

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



Live chat

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

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

Service requests

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

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

Note

To create a service request, you must have a valid support agreement. Contact a sales representative for details about obtaining a valid support agreement or with questions about an account.

To review an open service request:

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

Online communities

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

How to provide feedback

Feedback helps to improve the accuracy, organization, and overall quality of publications. You can send feedback to DPAD.Doc.Feedback@emc.com.

CHAPTER 1

Product Overview

This chapter includes the following topics:

- [Terminology used in this guide](#)..... 22
- [Introduction to the product](#)..... 22
- [Supported configurations](#)..... 34
- [Usage limits on Data Domain resources](#)..... 42
- [Road map for product operations](#)..... 46

Terminology used in this guide

The terms *database application agent*, *product*, and *software* in this guide refer to the database agent software that enables the Data Domain Boost for Enterprise Applications and ProtectPoint workflows and functionality.

The generic sections of this guide use the term *transaction logs* for the logs that are required to recover data that the database application agent backed up. The different applications that the product supports use application-specific terms for the logs, such as archived logs.

The UNIX references in this guide apply to both UNIX and Linux operating systems, unless specified otherwise. The Windows references apply to all the supported Microsoft Windows operating systems, unless specified otherwise.

The database application agent processes distinguish between the *restore* and *recovery* of a database:

- *Restore* means to retrieve data from backup and store the data on disk.
- *Recover* means to apply the transaction logs to make the database consistent.

The term *point-in-time restore* is also known as *object level restore*.

The glossary provides details about terms used in this guide. The terms include Data Domain, VMAX, and XtremIO specific terms related to the supported Data Domain Boost (DD Boost), ProtectPoint, and ProtectPoint with RecoverPoint operations.

Introduction to the product

The database application agent enables you to perform backups and restores of DB2, Oracle, SAP HANA, or SAP with Oracle database data with a Data Domain system. You can use the database-specific backup and recovery tools to perform the product operations.

The database application agent performs the following types of backups:

- DD Boost backups to a Data Domain system.
- ProtectPoint backups from VMAX primary storage to a Data Domain system.
- ProtectPoint with RecoverPoint backups from XtremIO primary storage to a Data Domain system.

You can use the database application agent to perform DD Boost backups and restores of DB2, Oracle, SAP HANA, or SAP with Oracle database data.

You can perform the DD Boost backups and restores over either an Ethernet (IP) or Fibre Channel (FC) network connection.

The database application agent supports in-flight encryption for DD Boost clients with a Data Domain system over a WAN connection. To use this feature, you can configure the Data Domain system with either medium-strength or high-strength encryption and set the authentication mode to anonymous. The configuration is transparent to the database application agent. The latest *Data Domain Boost Administration Guide* provides details.

You can also use the database application agent to protect specific types of database data through the supported ProtectPoint operations:

- You can protect DB2, Oracle, and SAP with Oracle database data that resides on a VMAX array. The database application agent uses the ProtectPoint technology to

protect the data on the VMAX system. [ProtectPoint backups and restores on a VMAX system](#) on page 24 provides more details.

- You can protect DB2, Oracle, and SAP with Oracle database data that resides on an XtremIO array. The database application agent uses the ProtectPoint and RecoverPoint technologies to protect the data on the XtremIO system. [ProtectPoint with RecoverPoint backups and restores on an XtremIO system](#) on page 28 provides more details.

The database application agent supports the Data Domain IPv6, IPv4, and mixed IPv4 and IPv6 networks for both DD Boost and ProtectPoint backups and restores.

Note

The database application agent does not distinguish TCP/IP network types (LAN, WAN, or MAN) and can successfully operate where packet loss is strictly 0% and latency is less than 20 ms.

DD Boost backups and restores

A DD Boost backup to a Data Domain system takes advantage of the DD Boost feature by using the following two components:

- The DD Boost library API enables the backup software to communicate with the Data Domain system.

The online software compatibility guide, available at <http://compatibilityguide.emc.com:8080/CompGuideApp/>, provides details about the supported versions of the DD Boost library and DD OS.

- The distributed segment processing (DSP) component reviews the data that is already stored on the Data Domain system and sends only unique data for storage. The DSP component enables the backup data to be deduplicated on the database or application host to reduce the amount of data transferred over the network. [DD Boost distributed segment processing](#) on page 23 provides more details.

During the restore of a backup to the client, the Data Domain system converts the stored data to its original nondeduplicated state before sending the data over the network.

DD Boost distributed segment processing

There are two modes of operation for sending backup data to a Data Domain system through DD Boost, one with distributed segment processing (DSP) enabled and the other with DSP disabled. The operation mode is set on the Data Domain system.

When DSP is enabled, the deduplication process is distributed between the DD Boost library and the Data Domain system. Parts of the deduplication process are run on the database or application host so that the DD Boost library sends only unique data to the Data Domain system over the network.

Distributed segment processing provides the following benefits:

- Throughput is potentially greater because the DD Boost library sends only unique data instead of all the data to the Data Domain system. The throughput improvements depend on the level of redundancy in the data being backed up, the overall workload on the database server, and the database server capability. In general, greater throughput is attained with higher redundancy, greater database server workload, and greater database server capability.
- The network bandwidth requirements are significantly reduced because only the unique data is sent to the Data Domain system over the network.

- Recovery from failed backups can be potentially much faster. If a large backup fails in the middle or toward the end and a user restarts the backup, the data that was already sent to the Data Domain system does not need to be resent. The backup completes more quickly on retry.

When distributed segment processing is enabled, the DD Boost library, which is part of the product, performs the following tasks:

1. Segments the data.
2. Computes IDs for the data segments.
3. Checks with the Data Domain system for duplicate segments.
4. Compresses unique segments that are not found on the Data Domain system.
5. Sends the compressed data to the Data Domain system, which writes the unique data to disk.

The local compression algorithm that is used by the DD Boost library must be configured on the Data Domain system. The *Data Domain Operating System Administration Guide* provides more information about local compression and its configuration.

When distributed segment processing is disabled, the DD Boost library sends the data directly to the Data Domain system over the network. The Data Domain system then segments, deduplicates, and compresses the data before writing the data to the disk.

Note

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

ProtectPoint backups and restores on a VMAX system

The database application agent uses the ProtectPoint technology. This technology enables snapshot backups of database data from primary storage on a VMAX system to protection storage on a Data Domain system.

In addition to storing the backups on the Data Domain system, the database application agent keeps the last SnapVX snapshot, also known as a local snapshot, on the VMAX system for a faster restore.

The online software compatibility guide, available at <http://compatibilityguide.emc.com:8080/CompGuideApp/>, provides details about the platforms, file systems, and volume managers supported for ProtectPoint operations.

The database application agent also protects files that are required for the database recovery and do not reside on VMAX or cannot be backed up through snapshots due to database vendor restrictions. The database application agent protects these files through a DD Boost backup. As a result, the database application agent provides overall protection of the database, regardless of where the data resides.

You use the database-specific backup and recovery tools to perform a ProtectPoint backup and recovery.

ProtectPoint technology uses the following features on the Data Domain system and VMAX array to provide the VMAX to Data Domain protection:

- On the Data Domain system:
 - vdisk and scsitarget services
 - FastCopy
- On the VMAX array:

- FAST.X, which can encapsulate external devices on Data Domain to VMAX
- SnapVX

A ProtectPoint backup takes a SnapVX snapshot on the VMAX system and moves the blocks to the Data Domain system over a storage area network (SAN), without going through the application host. The Data Domain protection storage device (vdisk) appears as an internal device to VMAX while the data itself is actually stored on the Data Domain system. The VMAX system tracks the data that has changed since the last update to the Data Domain protection device. Therefore, the VMAX system only sends the changed data to the Data Domain system during a ProtectPoint backup, instead of all the data.

A ProtectPoint backup of a database is a full backup with the cost of an incremental backup. The ProtectPoint backup also has minimum overhead on the application host because all the changed blocks are moved directly from VMAX to Data Domain over SAN.

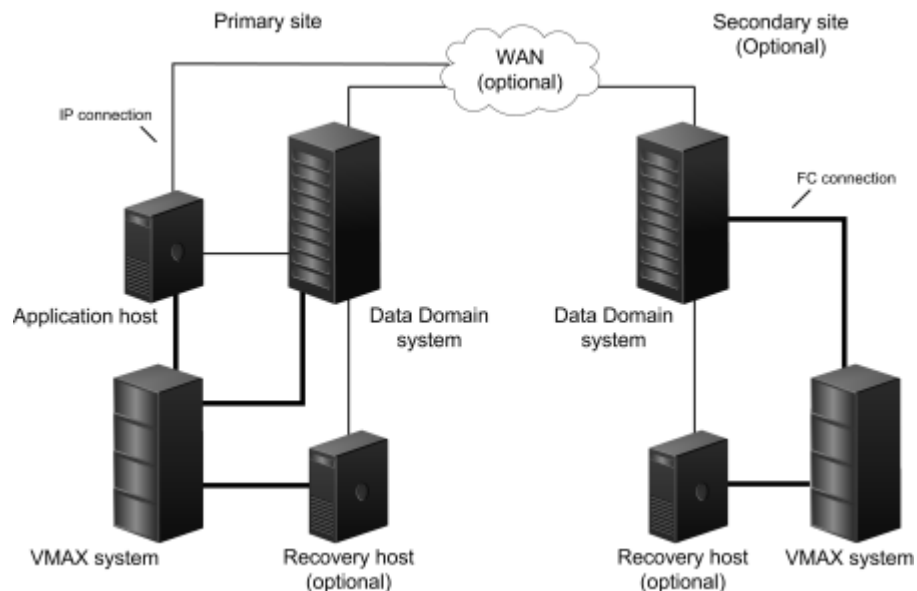
Note

With a VMAX system, when the database or archived logs are not on a logical volume manager (LVM), the database application agent serializes the ProtectPoint backup of each LUN. When an LVM is used, the database application agent performs a multistream backup that backs up each LUN that is part of the volume group in parallel. For example, if the database is on a volume group, db_vg, that contains 10 LUNs, the backup is performed by using 10 streams.

The following figure shows a sample ProtectPoint topology with a primary site and an optional secondary site. At the primary site, the application host accesses the database data that is stored on the VMAX system, and the backup data is transferred to the Data Domain system. A separate recovery host is optional. If the recovery is performed to the original application host, then the application host is also the recovery host.

The backup data can be replicated from the Data Domain system at the primary site to the Data Domain system at the secondary site. You can also restore the data to an optional recovery host at the secondary site.

Figure 1 ProtectPoint database application agent environment



ProtectPoint operations require both IP network (LAN or WAN) and Fibre Channel (FC) SAN connections. The following table lists the required types of network connections.

Table 3 Network connection types in a ProtectPoint environment

Connected components	Connection type
Primary site:	
Primary application host to primary VMAX system	FC
Primary application host to primary Data Domain system	IP
Primary VMAX system to primary Data Domain system	FC
(Optional) Primary recovery host to primary VMAX system	FC
(Optional) Primary recovery host to primary Data Domain system	IP
Secondary site (optional):	
Secondary recovery host to secondary VMAX system	FC
Secondary recovery host to secondary Data Domain system	IP
Secondary VMAX system to secondary Data Domain system	FC
Cross-site connections (optional):	
Primary application host to secondary Data Domain system	IP
Primary Data Domain system to secondary Data Domain system	IP
Primary VMAX system to secondary VMAX system VMAX replication on page 37 describes the SRDF/S support.	All supported by SRDF/S
Secondary VMAX system to primary Data Domain system	FC, if distance permits
Primary VMAX system to secondary Data Domain system	FC, if distance permits

ProtectPoint backup workflow with VMAX

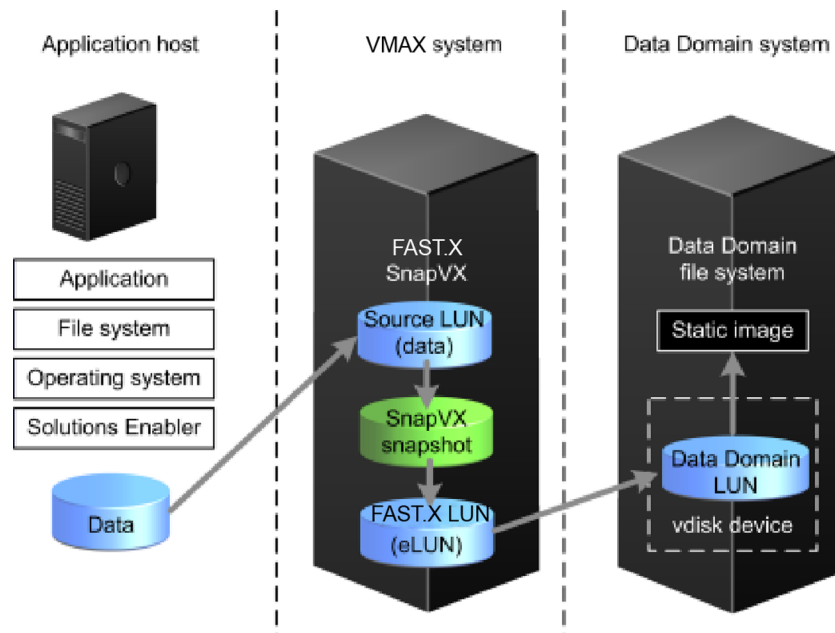
The ProtectPoint backup workflow with a VMAX system includes the following steps.

1. The database administrator starts a ProtectPoint backup by running the database native backup tool and specifying the backup type.
2. The database application agent maps the files in the backup to a list of VMAX source devices (source LUNs) and checks if those devices can be protected by using ProtectPoint.
3. The database application agent notifies the application or database that those files can be quiesced or placed in hot backup mode.
4. The database application agent creates a SnapVX snapshot on the VMAX system.
5. The database application agent notifies the application or database that the files can be unquiesced or taken out of backup mode, for minimum impact on the application or database.
6. The VMAX system copies the changed data on each source LUN to a corresponding Data Domain vdisk device, which is a VMAX FAST.X encapsulated LUN.

7. For each DD vdisk device, the Data Domain system creates and stores a DD vdisk static image, which is a permanent backup.
8. The database backup tool records the successful backup in its backup catalog.
9. The database application agent records the backup in its own catalog in a predefined DD Boost storage unit on the Data Domain system.
10. The database application agent backs up the files that ProtectPoint cannot protect, such as the Oracle control file or DB2 archive logs. The database application agent uses the DD Boost workflow to back up the files to the Data Domain system, which provides full data protection.

The following figure shows the ProtectPoint backup workflow in a ProtectPoint database application agent environment.

Figure 2 ProtectPoint backup workflow



ProtectPoint restore workflow with VMAX

The ProtectPoint restore workflow with a VMAX system includes the following steps.

1. The database administrator starts a ProtectPoint restore and recovery by running the database native recovery tool on the recovery host.
2. The database native recovery tool requests that the database application agent restore the required data and provides a backup handle.
3. The database application agent performs lookups in its own catalog to find the requested backup, which is a static image on the Data Domain system.
4. If the restore is performed from the last backup and the last SnapVX snapshot is in a valid state:
 - a. The database application agent links the snapshot to the VMAX LUN, which is known as the database application agent's restore LUN.
 - b. The restore continues with step 6.
5. The database application agent instantiates and overwrites the corresponding static image to a DD vdisk device, which is an encapsulated FAST.X LUN on a VMAX system. The FAST.X LUN is known as the database application agent's restore LUN.

6. By default, the database application agent mounts the restore LUN back to the recovery host and copies the required files to the requested locations.

If the DBA selects a rollback restore to the original host, then the database application agent performs a VMAX LUN-level restore to the original source device.

If the DBA selects a rollback restore to an alternate host, then the database application agent performs a VMAX LUN-level restore to the alternate target device.

NOTICE

A rollback restore is performed as part of the workflow that the database native recovery tool starts. If there are other files on the LUNs than the files requested for the particular database or database objects, the database application agent by default fails the restore as a safety measure.

A file that is named `psrollback.res` should be created in the required directory if the LUN selected for restore has other partitions or data that are not specified for the rollback restore. [Configuring rollback restores of ProtectPoint backups](#) on page 97 provides details.

Depending on the type of database, the database recovery tool might apply the transaction logs to make the database consistent or to roll forward to a time after the backup occurred. If the logs are not on the system, the database application agent restores and applies the logs through either a DD Boost restore or ProtectPoint restore, depending on how the logs were originally backed up.

ProtectPoint with RecoverPoint backups and restores on an XtremIO system

The database application agent uses the ProtectPoint and RecoverPoint technologies. These technologies enable snapshot backups of database data from primary storage on an XtremIO system to protection storage on a Data Domain system. The ProtectPoint and RecoverPoint technologies provide block movement of data from the XtremIO system source LUNs (managed by RecoverPoint consistency groups) to the Data Domain system. The database application agent also enables the restore of ProtectPoint with RecoverPoint backups from the Data Domain system.

A ProtectPoint with RecoverPoint backup of a database is a full backup with the cost of an incremental backup. The backup also has minimum overhead on the application host because all the changed blocks are moved directly from XtremIO to Data Domain through a RecoverPoint appliance (RPA).

The online software compatibility guide, available at <http://compatibilityguide.emc.com:8080/CompGuideApp/>, provides details about the platforms, file systems, and volume managers supported for the ProtectPoint with RecoverPoint operations.

The database application agent also protects files that are required for the database recovery and do not reside on XtremIO or cannot be backed up through snapshots due to database vendor restrictions. The database application agent protects these files through the DD Boost backup. As a result, the database application agent provides overall protection of the database, regardless of where the data resides.

You use the database-specific backup and recovery tools to perform a ProtectPoint with RecoverPoint backup and recovery.

The ProtectPoint and RecoverPoint technologies use the following features on the Data Domain system, RecoverPoint cluster, and XtremIO array to provide the XtremIO to Data Domain protection:

- On the Data Domain system:
 - vdisk and scsitarget services
 - FastCopy
 - DD Boost
- On the RecoverPoint cluster:
 - RecoverPoint consistency groups
- On the XtremIO array:
 - XtremIO Initiator Groups

A ProtectPoint with RecoverPoint backup takes a point-in-time snapshot on the XtremIO system and moves the blocks to the Data Domain system through the RPA, without going through the application host. The RecoverPoint system tracks the data that has changed since the last update to the Data Domain protection device. Therefore, the RecoverPoint system only sends the changed data to the Data Domain system during a ProtectPoint with RecoverPoint backup, instead of all the data.

In RecoverPoint, source LUNs (volumes) are protected by consistency groups. If two data sets are dependent on each other, such as a database and a database log, they should be part of the same consistency group. Logical components of a consistency group include copies, replication sets, and journals:

- Copies are all the volumes of a consistency group that are either a source or a target of replication at a specific RPA cluster. The copies include production copies, local copies, remote copies, and their journal volumes.
- A consistency group consists of one or more replication sets that include a production volume and any local or remote volumes to which the production volume is replicating. The number of replication sets in the system is equal to the number of production volumes being replicated.

A RecoverPoint group set is a user-defined set of consistency groups that is used to perform operational and recovery activities. The RecoverPoint documentation provides complete details about consistency groups and their components and setup procedures.

For ProtectPoint with RecoverPoint operations:

- The local copy in a consistency group exists on the Data Domain system, and there is no journal volume for that local copy. The consistency group can have a maximum of one local copy that is on a Data Domain system.
- You cannot enable parallel bookmarking for a group set.
- If a logical volume manager (LVM) controls the volumes on the application host, then all the LVM physical volumes (disks) that belong to one LVM volume group must be added to one RecoverPoint consistency group. LVM2 on Linux is an example of an LVM type.

Do not add the physical volumes belonging to one LVM volume group to multiple consistency groups. If you add the physical volumes to multiple consistency groups, the ProtectPoint with RecoverPoint backup fails with the following message:

```
Consistency group is already running.
```

You can create a consistency group that contains the physical volumes from two LVM volume groups. For example, the LVM volume group VG1 has x number of physical volumes, and VG2 has y number of physical volumes. You can create a consistency group that contains all the $x+y$ physical volumes from both volume groups.

Note

Whether or not an LVM is used, all the LUNs in the same consistency group are backed up in parallel through the RecoverPoint software. The RecoverPoint documentation provides details.


- If you delete a consistency group, then any associated static images (created by backups) on the Data Domain system are not deleted automatically. You can restore from these backups, but you cannot delete these backups with the database-specific backup and recovery tools.

If any of these backups become obsolete, delete the associated static images manually from the Data Domain system according to the Data Domain documentation.

The following figure shows a sample ProtectPoint with RecoverPoint topology with a primary site and an optional secondary site. At the primary site, the application host accesses the database data that is stored on the XtremIO system, and the backup data is transferred to the Data Domain system. A separate recovery host is optional. If the recovery is performed to the original application host, then the application host is also the recovery host.

If you have an optional secondary site, the backup data can be replicated from the Data Domain system at the primary site to the Data Domain system at the secondary site. At the secondary site, you can also recover the data to an optional recovery host.

Figure 3 ProtectPoint with RecoverPoint environment



ProtectPoint with RecoverPoint operations require both IP network (LAN or WAN) and Fibre Channel (FC) SAN connections. The following table lists the required types of network connections.

Table 4 Network connection types in a ProtectPoint with RecoverPoint environment

Connected components	Connection type
Primary site:	
Primary application host to primary XtremIO system	FC
Primary application host to RPA	IP
Primary application host to primary Data Domain system	IP or (FC and IP)
Primary XtremIO system to RPA	FC and IP
RPA to primary Data Domain system	IP and (optional) FC
(Optional) Primary recovery host to primary XtremIO system	FC
(Optional) Primary recovery host to primary Data Domain system	IP or (FC and IP)
(Optional) Primary recovery host to RPA	IP
Secondary site (optional):	
Secondary recovery host to secondary XtremIO system	FC
Secondary recovery host to secondary Data Domain system	FC and IP
Cross-site connections (optional):	
Primary application host to secondary Data Domain system	IP
Primary Data Domain system to secondary Data Domain system	IP

ProtectPoint with RecoverPoint backup workflow with XtremIO

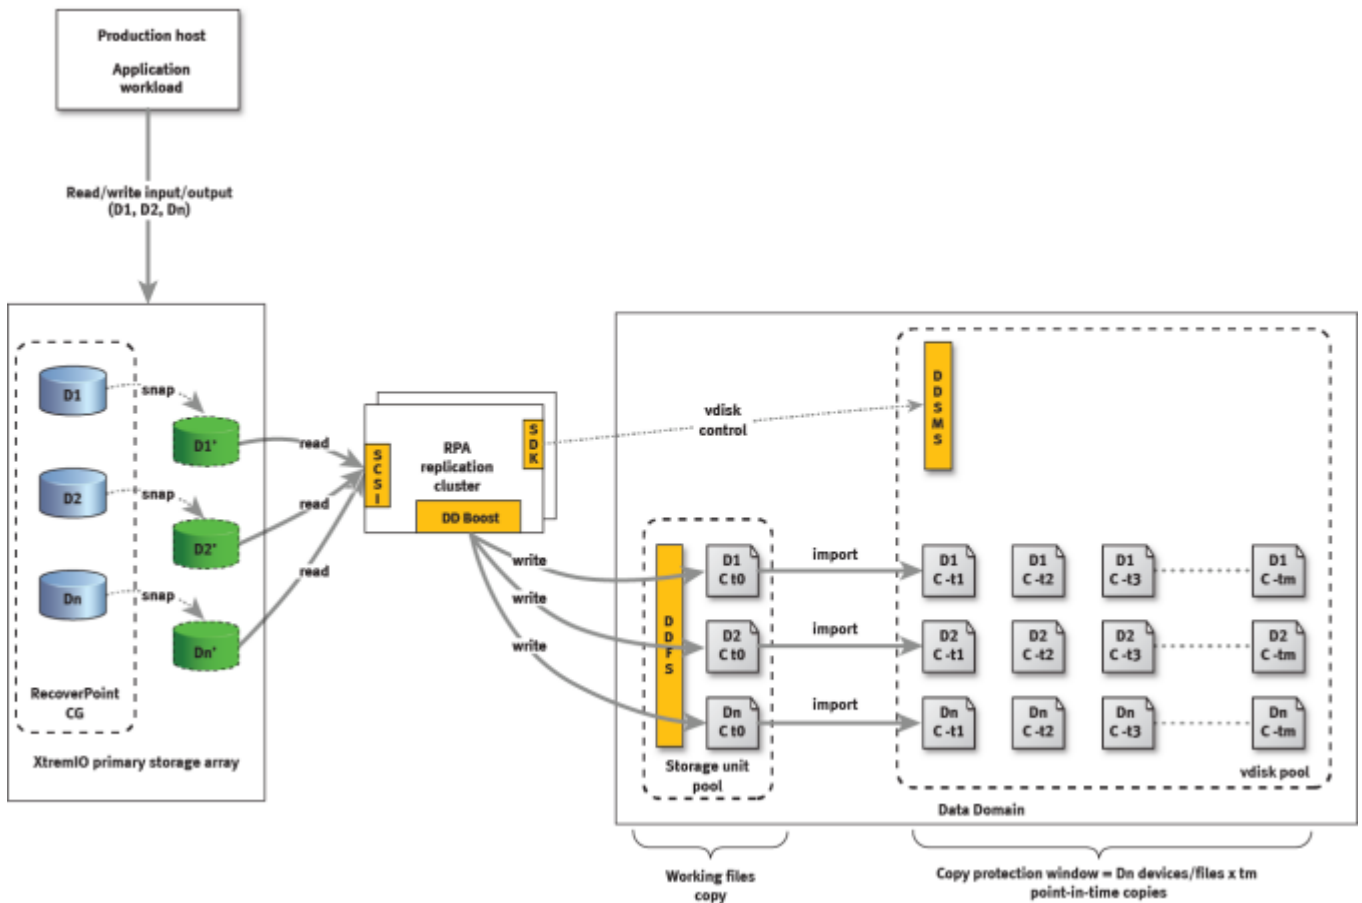
The ProtectPoint with RecoverPoint backup workflow with an XtremIO system includes the following steps.

1. The database administrator starts a ProtectPoint with RecoverPoint backup by running the database native backup tool and specifying the backup type.
2. The database application agent maps the files in the backup to a list of XtremIO source LUNs, and obtains information about the relevant consistency groups from RecoverPoint. The agent checks if the source LUNs can be protected by using ProtectPoint with RecoverPoint.
3. The database application agent notifies the application or database that the files can be quiesced or placed in hot backup mode.
4. The database application agent notifies RecoverPoint to create a point-in-time snapshot (bookmark in RecoverPoint terminology) of the consistency groups that contain the source LUNs.
5. RecoverPoint creates a snapshot of all the required consistency groups on the XtremIO system.
6. The database application agent notifies the application or database that the files can be unquiesced or taken out of backup mode, for minimum impact on the application or database.
7. RecoverPoint uses DD Boost to write all the changed blocks (changed since the previous snapshot) to working files on the Data Domain system.
8. RecoverPoint uses the FastCopy service to create and store a DD vdisk static image from each DD Boost working file. The vdisk static images form the permanent backup.

9. The database backup tool records the successful backup in its backup catalog.
10. The database application agent records the backup in its own catalog in a predefined DD Boost storage unit on the Data Domain system.
11. The database application agent backs up the files that ProtectPoint cannot protect, such as the Oracle control file or DB2 archive logs. The database application agent uses the DD Boost workflow to back up the files to the Data Domain system, which provides full data protection.

The following figure shows the ProtectPoint with RecoverPoint backup workflow in a ProtectPoint database application agent environment.

Figure 4 ProtectPoint with RecoverPoint backup workflow



ProtectPoint with RecoverPoint restore workflow with XtremIO

The ProtectPoint with RecoverPoint restore workflow with an XtremIO system includes the following steps.

1. The database administrator starts a ProtectPoint with RecoverPoint restore and recovery by running the database native recovery tool on the recovery host.
2. The database native recovery tool requests that the database application agent restore the required data and provides a backup handle.
3. The database application agent performs lookups in its own catalog to find the requested backup, which consists of static images on the Data Domain system.
4. The database application agent instantiates the corresponding static images on restore LUNs on the Data Domain system through the vdisk service.

5. By default, the database application agent mounts the restore LUNs directly to the recovery host, which can be either the original backup host or a different host, and copies the required files to the requested locations.
6. If the DBA selects a rollback restore to the original host, then the database application agent requires the RecoverPoint cluster to perform a LUN-level restore to the original source LUNs.

With RecoverPoint pre-5.0, if the DBA selects a rollback restore, then the database application agent requires the RecoverPoint cluster to perform a restore of the entire consistency group to the original source LUNs. If a consistency group being restored contains multiple LUNs, then all those LUNs are overwritten and inaccessible during the rollback restore, even when the backed-up objects reside on only certain LUNs.

NOTICE

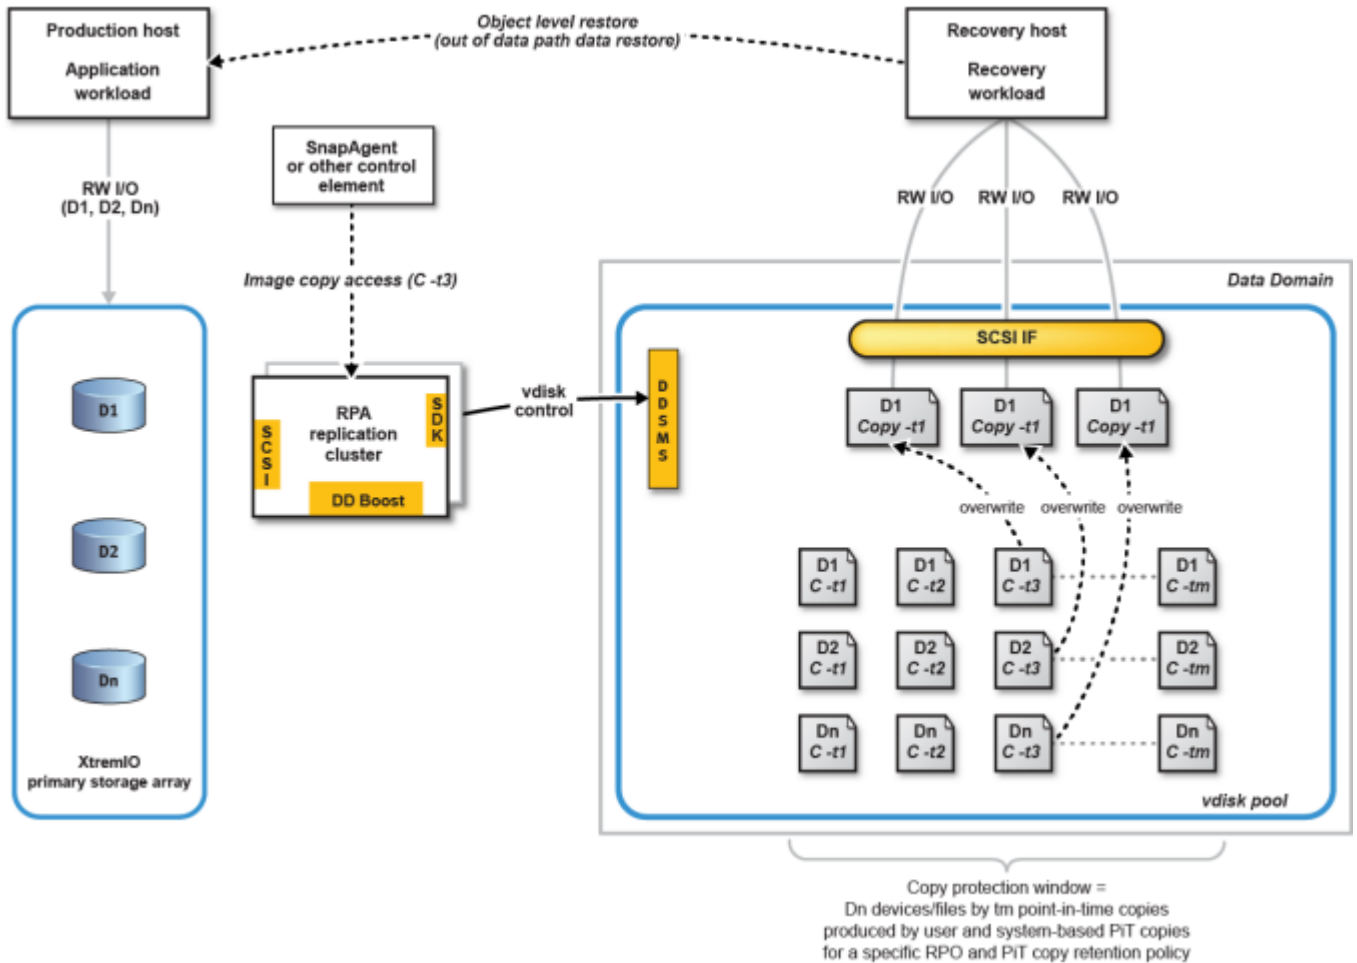
A rollback restore is performed as part of the workflow that the database native recovery tool starts. If there are other files on the LUNs than the files requested for the particular database or database objects, the database application agent by default fails the restore as a safety measure.

A file that is named `psrollback.res` should be created in the required directory if any LUN that is to be restored has other partitions or data that are not specified for the rollback restore. [Configuring rollback restores of ProtectPoint backups](#) on page 97 provides details.

Depending on the type of database, the database recovery tool might apply the transaction logs to make the database consistent or to roll forward to a time after the backup occurred. If the logs are not on the application host, the database application agent restores and applies the logs through either a DD Boost restore or ProtectPoint restore. The type of restore depends on how the logs were originally backed up.

The following figure shows the ProtectPoint with RecoverPoint point-in-time restore workflow, which is the default restore workflow in the ProtectPoint database application agent environment.


Figure 5 ProtectPoint with RecoverPoint restore workflow



Supported configurations

You can deploy the database application agent in a database stand-alone or high-availability configuration in a supported environment.

The following figure shows a stand-alone configuration that uses the product to back up or restore the data on the database server host to or from the Data Domain system.

Figure 6 DD Boost for Enterprise Applications in a stand-alone configuration

The DD Boost for Enterprise Applications version must be compatible with the Data Domain systems, VMAX systems, and XtremIO systems used. Data Domain does not support combinations other than the ones that are detailed in the online software compatibility guide, which is available at <http://compatibilityguide.emc.com:8080/CompGuideApp/>.

High-availability configurations

The database application agent supports backups and restores in the following high-availability environments:

- DB2 Database Partitioning Feature (DPF) environment
- DB2 High Availability Disaster Recovery (HADR) environment
- DB2 pureScale environment
- Oracle Real Application Clusters (RAC) environment
- SAP HANA replication environment (DD Boost operations only)
- SAP HANA scale-out multinode environment (DD Boost operations only)
- SAP with Oracle RAC environment
- Operating system active-passive cluster

A DB2 DPF system offers an environment where a single database is divided into multiple partitions, either on the same host or on multiple hosts.

A DB2 HADR system consists of a primary host and multiple standby hosts. HADR replicates all the data changes from the primary host to the standby hosts, and provides fast failover to a standby host when the primary host fails. You can perform a backup on the primary host only.

A DB2 pureScale system is an active-active application cluster with a shared-disk architecture that includes a single database partition shared by the group of cluster member nodes. The term node as used in other active-active application clusters is referred to as a member in a DB2 pureScale system.

An Oracle RAC system is an active-active application cluster environment where a node is a physical and virtual host. In an Oracle RAC system, an Oracle instance is a memory structure and a group of Oracle Server processes running on a node. The RAC system enables multiple Oracle instances across multiples nodes to access the same Oracle database simultaneously. Oracle RAC is a cluster software infrastructure that provides concurrent access to the same storage and the same set of datafiles from all nodes in the cluster. All the database files reside on shared disks.

SAP HANA software supports high availability environments for DD Boost operations, including SAP HANA replication environments and SAP HANA scale-out appliances that contain multiple nodes.

An operating system active-passive cluster includes multiple hosts (nodes) connected by a shared SCSI bus with common storage attached. A user can define cluster services, such as Microsoft cluster services or Veritas cluster services, and assign the services their own IP addresses and names (virtual cluster hosts). The services and their associated storage can migrate for failover between the hosts in the cluster.

The online software compatibility guide, available at <http://compatibilityguide.emc.com:8080/CompGuideApp/>, provides details about the supported versions and high-availability environments.

Note

You must install the database application agent on each node in a high-availability environment.

The application-specific chapters in this guide provide details about the configuration procedures and the backup and restore operations in the supported high-availability environments.

Data Domain High Availability (HA)

The database application agent supports the Data Domain High Availability (HA) for DD Boost operations and for improved resilience with ProtectPoint operations. Data Domain HA enables you to configure two Data Domain systems as an Active-Standby pair, which provides redundancy in the event of a system failure. HA keeps the active and standby systems in sync, so that if the active node fails due to hardware or software issues, the standby node can take over the services and continue where the failed node left off.

Data Domain HA includes the following additional features:

- Supports failover of backup, restore, replication, and management services in the two-node system. Automatic failover requires no user intervention.
- Provides a fully redundant design with no single point of failure when the system is configured as recommended.
- Provides an Active-Standby system with no loss of performance on failover.
- Provides failover within 10 minutes for most operations. CIFS, VTL, and NDMP must be restarted manually.
- Supports both IP and Fibre Channel (FC) connections. Both nodes must have access to the same IP networks, FC SANs, and hosts.

The latest version of the *Data Domain Operating System Administration Guide* provides complete details about all the supported HA features.

The deployment of the database application agent with Data Domain HA improves the resilience in the ProtectPoint workflows, in terms of the data paths involved in the operations. However, if a failover occurs when the host sends vdisk commands to the Data Domain system (control path) in a ProtectPoint workflow, the database application agent fails its operation.

Virtualization support

The database application agent supports several types of virtualization software, such as VMware, Solaris zones, and Microsoft Hyper-V.

The online software compatibility guide, available at <http://compatibilityguide.emc.com:8080/CompGuideApp/>, provides details about the supported environments and platforms. [Solaris installation](#) on page 69 provides details about Solaris zones.

Note

You must install the database application agent in the guest operating system.

Data Domain replication

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

The database application agent does not initiate or monitor a replication. However, the product can restore from the replicated copy on a secondary Data Domain system. You must have used the product to create the backup on a primary Data Domain system. A Data Domain administrator performs the backup replication from the primary system to the secondary system.

Note

The replication process must not change the names of the directories and files created by the database application agent.

To enable the backup replication and subsequent restore from a secondary Data Domain system, the user ID or primary group ID of the DD Boost users on the primary and secondary systems must be identical.

You must meet specific configuration requirements to enable the restore of replicated backups from a secondary Data Domain system. [Configuring restores of replicated backups](#) on page 90 provides details.

The Knowledgebase Article number 182294, titled *Configuration of DDBoost Users on Source and Destination DDRs for MTree Replication*, provides more details. The article is available on the Support website at <https://support.emc.com>.

VMAX replication

The database application agent supports ProtectPoint protection that uses a primary or secondary VMAX system in a VMAX replication environment. In this environment, the primary and secondary VMAX storage arrays are connected by a Symmetrix Remote Data Facility (SRDF).

The database application agent for ProtectPoint supports both the SRDF synchronous mode, SRDF/S, and SRDF/Metro. The following topics provide details about the database application agent support of SRDF/S and SRDF/Metro.

Support of VMAX SRDF/S

In a VMAX replication environment, the database application agent for ProtectPoint supports the Symmetrix Remote Data Facility (SRDF) in synchronous mode, SRDF/S.

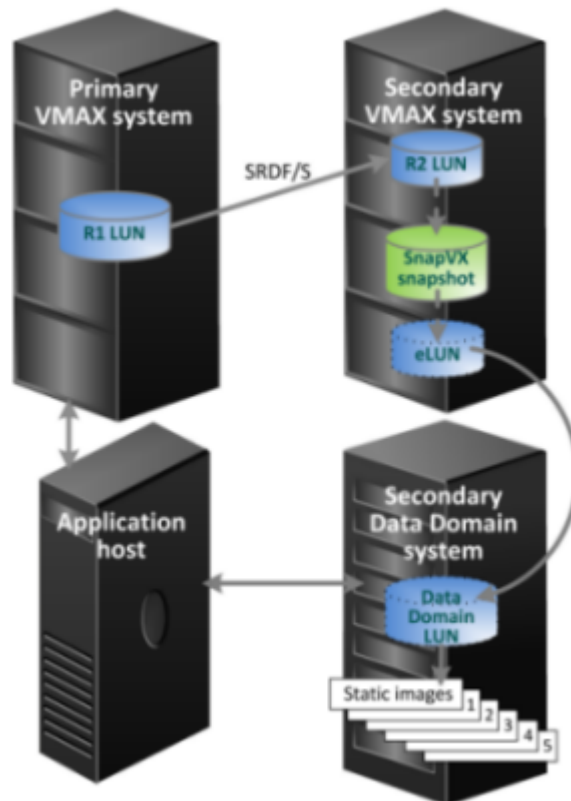
SRDF/S is a VMAX feature that maintains a synchronous, real-time copy of data at the LUN level between the primary and secondary VMAX arrays. A source LUN referred to as R1 on the primary array is associated with a source LUN referred to as R2 on the secondary array. The SRDF/S software maintains continuous synchronization of the two sources by copying all changes on one LUN device to the other. The VMAX documentation provides more details about VMAX replication and the SRDF/S functionality.

The following figure shows VMAX arrays with an SRDF/S link, where the secondary VMAX system is attached to a secondary Data Domain system. In this SRDF configuration, you can use the database application agent to perform a ProtectPoint backup to the secondary Data Domain system, which backs up the R2 LUN.

Note

The database application agent also uses the DD Boost workflow to back up any nonsnapshotable files and create catalog entries.

Figure 7 ProtectPoint backup to a secondary Data Domain in an SRDF configuration



The following figure shows VMAX arrays with an SRDF/S link, where both the primary and secondary VMAX systems are attached to a Data Domain system. In this SRDF configuration, you can use the database application agent to perform a ProtectPoint backup to either the primary or secondary Data Domain system. The primary ProtectPoint backup backs up the R1 LUN to the primary Data Domain system. The

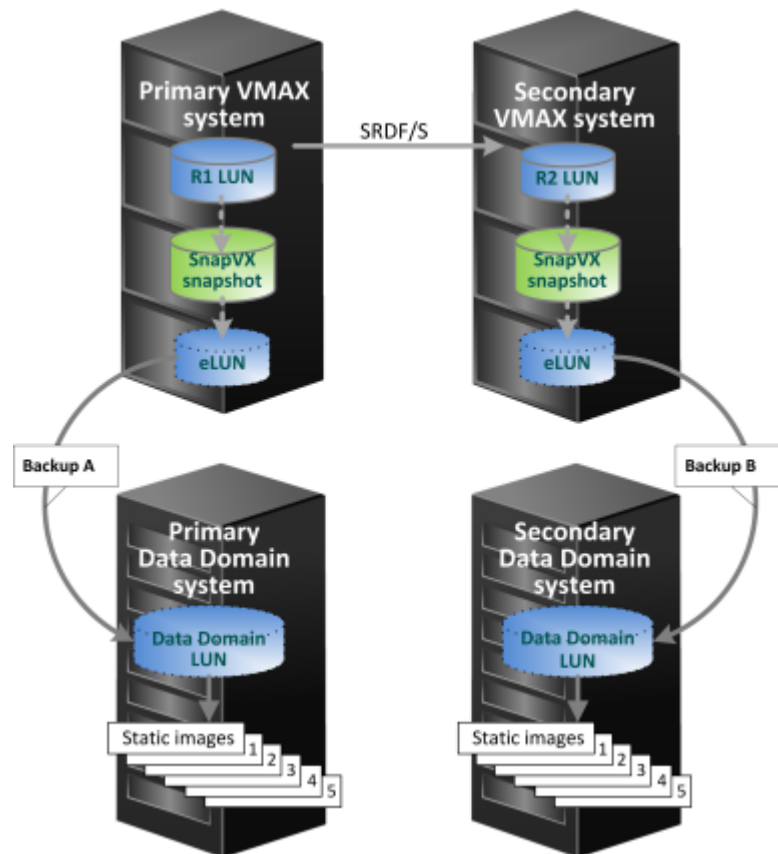
secondary ProtectPoint backup backs up the R2 LUN to the secondary Data Domain system.

Note

The database application agent cannot perform backups to both Data Domain systems in the same backup session.

In these SRDF configurations, the database application agent validates the synchronization of the R1 and R2 LUNs. The database application agent then creates a SnapVX snapshot of the R2 LUN to transfer the backup data to the secondary Data Domain system.

Figure 8 ProtectPoint backup to a primary or secondary Data Domain in an SRDF configuration



SRDF/S requirements and configuration support for the ProtectPoint database application agent are as follows:

- The database application agent automatically determines the state of the SRDF/S link at runtime.
- If there is no SRDF/S link at the start of an operation, then the backup or restore operation fails.
- The database application agent does not support any changes to the SRDF/S link mode made during a backup or restore operation.
- If the SRDF link is in a failed over or failed back state, then the database application agent operations fail.
- SRDF replication cannot transition between asynchronous and synchronous modes during any VMAX operation. The mode must remain constant.

- The database application agent does not support the creation of snapshots of file systems or volume groups that cross SRDF groups.
- The database application agent supports only single-hop remote connections. The database application agent does not support cascaded VMAX configurations.
- The database application agent does not support concurrent SRDF or concurrent SRDF/Star configurations where R1 is a source to two or more concurrent targets.

Support of VMAX SRDF/Metro

In a VMAX replication environment, the database application agent for ProtectPoint provides limited support of the Symmetrix Remote Data Facility (SRDF) in an SRDF/Metro configuration.

The SRDF/Metro support requires a Request for Product Qualification as detailed in the online software compatibility guide, available at <http://compatibilityguide.emc.com:8080/CompGuideApp/>. SRDF/Metro is a high availability facility, rather than a disaster recovery facility as provided by other SRDF implementations.

In its basic configuration, SRDF/Metro consists of pairs of R1 and R2 devices that are connected by an SRDF link, as in any other SRDF configuration. However, in SRDF/Metro, both the R1 and R2 devices are write accessible by the host systems concurrently. The SRDF R2 device acquires the external identity (geometry, device WWN) of the R1 device. Each pair of devices appears to the host systems as a single virtual device across the two SRDF paired VMAX arrays.

Note

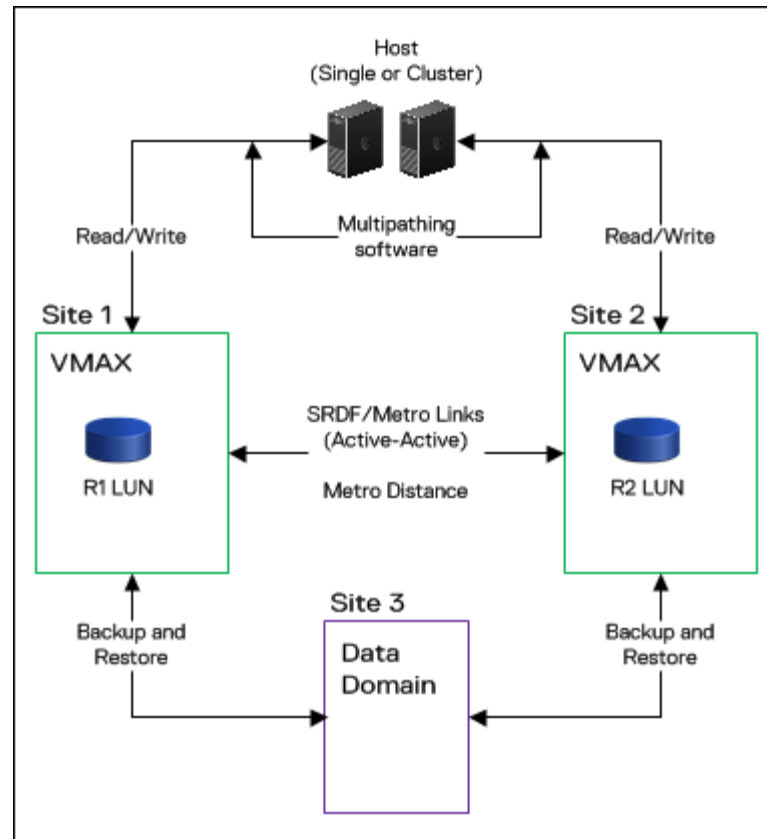
If the devices are not fully Metro paired, that is Metro did not reach the ActiveActive or ActiveBias state, the devices would not have the same external identity. The devices would present themselves as different devices to the host.

When both devices are accessible, the single host or cluster hosts can read and write to both the R1 and R2 devices. SRDF/Metro ensures that each copy remains current and consistent, and addresses any write conflicts that might occur between the paired SRDF devices.

The *SRDF/Metro Overview and Best Practices Technical Notes* provides an overview of SRDF/Metro and information about best practices. The VMAX documentation provides more information about VMAX replication and the SRDF/Metro functionality.

The following figure shows the only SRDF/Metro topology that the database application agent supports. Both VMAX arrays are attached to the same Data Domain system.

Figure 9 SRDF/Metro supported topology



SRDF/Metro requirements and configuration support for the ProtectPoint database application agent are as follows:

- Both VMAX arrays are connected to the same Data Domain system through a Fibre Channel (FC) network connection.
- ProtectPoint FAST.X backup and restore devices are configured for both VMAX arrays.
- Native VMAX LUNs are configured for both VMAX arrays. [Configuring ProtectPoint VMAX restores from local snapshots](#) on page 96 provides more details.
- Solutions Enabler 8.4 or later is installed on the application host systems.
- The database application agent supports all the SRDF/Metro witness configurations. The VMAX documentation, such as the latest *Solutions Enabler SRDF Family CLI User Guide*, provides details.
- The database application agent does not support concurrent or cascaded SRDF configurations.

To configure SRDF/Metro for ProtectPoint, the storage administrator must create the initial SnapVX link from both R1 and R2 devices to their FAST.X encapsulated backup devices.

The following additional considerations and requirements apply for backup and restore operations with the ProtectPoint database application agent:

- Database application agent backups:

- If the application data resides on devices that are part of an SRDF Metro group, the database application agent automatically discovers the R1-R2 pair for the devices, and then uses the R1 device for the rest of the operations.
- Database application agent point-in-time restores:
 - During a point-in-time restore through a FAST.X encapsulated restore device, the restore workflow always allocates the restore device from the VMAX array used for the backup (the R1 site at the backup time). A point-in-time restore from a different VMAX system is not supported.
 - A point-in-time restore uses the local SnapVX snapshot only when the local snapshot on the restore target VMAX is the same as the one that was used during the backup for the specified backup identifier. Otherwise, the point-in-time restore uses the copy on the Data Domain system.
- Database application agent rollback restores:
 - Before you perform a rollback restore, the SRDF link must be manually suspended.
 - The database application agent discovers the R1 at the time of the restore and rolls back to the R1, regardless of which site was used during the backup. If the device is no longer part of the SRDF/Metro, the rollback behavior is the same as for a non-SRDF/Metro configuration. In this case, the rollback is performed to the original source device that was backed up unless the rollback occurs to alternate LUNs as described in [Rollback restores of ProtectPoint for VMAX backups](#) on page 97.
 - A rollback restore uses the local SnapVX snapshot only when the local snapshot on the restore target VMAX is the same as the one that was used during the backup for the specified backup identifier. Otherwise, the rollback restore uses the copy on the Data Domain system.
 - After you perform a rollback restore, the SRDF link must be reestablished.

Usage limits on Data Domain resources

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

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

Note

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

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

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

- When a soft limit is exceeded, the Data Domain host generates an alert, and if a tenant-unit notification list is configured, sends an email to the addresses on the list. An application can continue to use more of the limited resource after a soft limit is exceeded.

- When a hard limit is exceeded, an application cannot use any more of the limited resource.

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

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

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

The database application agent user can run the `ddbmadmin` command to determine the streams limits of a storage unit. For example:

```
ddbmadmin -l -z <configuration_file_name>
```

```
Stream limits for storage unit rp731 on the primary Data Domain host
dd_host1.emc.com:
    active write streams: 11
    active read streams: 0
    soft limit write streams: none
    soft limit read streams: none
    soft limit combined streams: 40
    hard limit combined streams: 60
```

Database application agent operations with Data Domain usage limits on capacity

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

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

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

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

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

When the capacity hard limit of a vdisk pool is exceeded during a ProtectPoint backup, the backup fails.

Database application agent operations with Data Domain usage limits on streams

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

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

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

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

Database application agent ProtectPoint operations with Data Domain usage limits

For ProtectPoint operations, follow the general recommendations and requirements for resource usage by all the supported applications, such as the requirements for parallelism and accommodation of resources for backups and restores.

In a ProtectPoint workflow, the snapshot agent produces the snapshot static images, and both the snapshot agent and application produce the metadata records for snapshot catalog purposes.

Usage limits on Data Domain streams

The catalog records are saved and retrieved sequentially during a ProtectPoint operation. Only one stream is required for these records during an entire backup, restore, or deletion. When no streams are available, the operation fails.

Usage quota on Data Domain capacity

The storage resources that a ProtectPoint backup consumes are the vdisk static images and the metadata records produced by both the application and snapshot agent for the snapshot catalog.

Note

DD OS 5.7.2 or later supports the capacity hard limit on a DD vdisk pool. For earlier versions of DD OS, do not use the capacity quota setting on a vdisk pool.

The static images are stored in a vdisk pool, and their size is uncompressed and LUN-based. Calculate the vdisk pool capacity based on the source LUN size and the number of LUNs in a snapshot backup, rather than on the size of the database that is backed up. For ProtectPoint with RecoverPoint, you must also consider that the backup occurs for the whole consistency group. In addition to the application agent-initiated backups, RecoverPoint can invoke backups based on the replication policy setting of a consistency group.

RecoverPoint might require internal storage space for a backup and restore, for example, for the storage unit that is used for working files during the static image staging in creation or retrieval. Ensure that the configured resource limit and quota of the storage unit are sufficient for the snapshot backup. If in doubt, do not limit the

usage of this storage unit. The RecoverPoint management user guides provide specific configuration recommendations.

During the creation of a static image, if the hard capacity limit is reached for a vdisk pool in a ProtectPoint snapshot, the backup fails at the save stage. Error messages appear in the snapshot agent logs:

```
0 11/17/16 15:59:41 1267619584 14632 1479416331 ddbsmnd NSR info [msg
#404 dd_snapshot_impl.cpp 1339 PSDBG -1] Snapshot to Data Domain
failed: Snapshot to DD lfcp0031.lss.emc.com for DD WWN
6002188000000002D0057F65F1000007 failed: Error creating static image
for DD WWN 6002188000000002D0057F65F1000007. ([5005] no room left).

109581 11/17/16 15:45:04 1647015680 14051 1479415465 ddbsmnd NSR info
Snapshot to DD lfcp0031.lss.emc.com for DD WWN
6002188000000002D0057F65F1000007 failed: Error creating static image
for DD WWN 6002188000000002D0057F65F1000007. ([5009] I/O error).
```

The catalog records are stored on the storage unit that is specified by the `DEVICE_PATH` parameter. The size of the save sets is typically small. When the hard limit is reached, the backup fails at the snapshot save stage, either by the snapshot agent or application agent.

When the creation of the metadata record of the application agent fails, error messages appear in the operational or debug log of the application agent. For example, a typical error message in the DB2 application agent log is as follows:

```
(pid = 26762) (12/13/2016 11:39:05 AM) lnm_meta_close: Exiting with
error:
Closing a file failed [5057] ([26762] [140615742465824] Tue Dec 13
11:39:05 2016
    ddp_close_file() failed, Err: 5057-File handle is stale
).

(pid = 26762) (12/13/2016 11:39:05 AM) Unable to complete the saving
of the metadata of the backup. An error occurred in closing the index
session.
```

When the creation of the metadata record of the snapshot agent fails, error messages appear in the operational and debug log of the agents. For example:

- A typical error message in the ddbsmnd debug log is as follows:

```
0 10/28/16 08:58:30 4052920064 6962 1477670213 ddbsmnd NSR notice
10/28/16 08:58:30.081704 fsys_open failed for pathname /
nsm_30_141/pp-db2-30-142/27/2.0/data_ss/$db2_acs:$/
_ts10k_147767/1477670309.ss with error Opening the file '/
nsm_30_141/pp-db2-30-142/27/2.0/data_ss/$db2_acs:$/
_ts10k_147767/1477670309.ss' failed [5005] ([ 6962]
[140161720690432] Fri Oct 28 08:58:30 2016
ddp_open_file() failed for File: //nsm_30_141/pp-db2-30-142/27/2.0/
data_ss/$db2_acs:$/_ts10k_147767/1477670309.ss, Err: 5005-nfs
setattr failed (nfs: No space left on device)).
```

- A typical error message in the ddbsmnd trace log is as follows:

```
[msg #286 backup_subr.cpp 616 PSDBG 5] After calling nsr_start,
there was a non-retryable error
Error [msg #288 backup_subr.cpp 617 PSDBG -1] NULL[msg #289
backup_subr.cpp 624 PSDBG 5] After calling nsr_start, Can not
recover from error, exiting
[msg #290 backup_subr.cpp 800 PSDBG 5] BR::Leaving function
br_cvr_save_func
Error [msg #292 FileBackupService.cpp 1190 PSDBG -1] Failed to
```

```
create snapshot. Reason: nulError [msg #294 FileBackupService.cpp
1193 PSDBG -1] SAVE FAIL[msg #295 FileBackupService.cpp 4020
PSDBG 5] Invalid cover id, assuming there is no cover id
```

- A typical error message in the DB2 application agent debug log is as follows:

```
(pid = 6869) (10/28/2016 08:58:30 AM) INFO: Object = /Mount3/
NODE0000/LOGSTREAM0000, type=SAVE, status=ERROR
```

```
(pid = 6869) (10/28/2016 08:58:30 AM) ERROR: An error has
encountered, Object = /Mount3/NODE0000/LOGSTREAM0000,
returnCode=-1. Exit with error.
```

Road map for product operations

Use the following road map to install and configure the database application agent software, and perform backups and restores with the product on the supported database and application hosts.

Procedure

1. Review the *Data Domain Boost for Enterprise Applications and ProtectPoint Database Application Agent Version 4.0 Release Notes* for the most up-to-date information about the product.
2. Ensure that the required Data Domain systems are set up, configured, and enabled according to [Data Domain System Configuration for DD Boost Operations](#) on page 49.
3. For ProtectPoint operations, ensure that the required VMAX, XtremIO, and Data Domain system configurations are completed. The *ProtectPoint Version 4.0 Primary and Protection Storage Configuration Guide* provides details.
4. Install the database application agent software according to [Product Installation](#) on page 59.
5. Configure the product according to [Product Configuration](#) on page 77.
6. Protect the database system by using the required procedures:
 - Procedures for DB2 operations:
 - [DD Boost Operations on DB2 Systems](#) on page 141
 - [ProtectPoint Operations on DB2 Systems](#) on page 181
 - Procedures for Oracle operations:
 - [DD Boost Operations on Oracle Systems](#) on page 213
 - [ProtectPoint Operations on Oracle Systems](#) on page 229
 - Procedures for SAP HANA operations:
 - [DD Boost Operations on SAP HANA Systems](#) on page 251
 - Procedures for SAP with Oracle operations:
 - [DD Boost Operations on SAP with Oracle Systems](#) on page 271
 - [ProtectPoint Operations on SAP with Oracle Systems](#) on page 289
7. Troubleshoot issues with product operations by using the following information:
 - [General troubleshooting tips](#) on page 132
 - DB2 troubleshooting tips:
 - [DB2 troubleshooting tips for DD Boost operations](#) on page 174

- [DB2 troubleshooting tips for ProtectPoint operations](#) on page 210
- Oracle troubleshooting tips:
 - [Oracle troubleshooting tips for DD Boost operations](#) on page 227
 - [Oracle troubleshooting tips for ProtectPoint operations](#) on page 248
- SAP HANA troubleshooting tips:
 - [SAP HANA troubleshooting tips for DD Boost operations](#) on page 269
- SAP with Oracle troubleshooting tips:
 - [SAP with Oracle troubleshooting tips for DD Boost operations](#) on page 287
 - [SAP with Oracle troubleshooting tips for ProtectPoint operations](#) on page 310
- *Data Domain Boost for Enterprise Applications and ProtectPoint Database Application Agent Version 4.0 Release Notes*

CHAPTER 2

Data Domain System Configuration for DD Boost Operations

This chapter includes the following topics:

- [Licensing the Data Domain system](#).....50
- [Enabling DD Boost on a Data Domain system](#).....50
- [Changing the DD Boost access rights](#)..... 51
- [Enabling encryption over a WAN connection](#)..... 52
- [Enabling the DD Boost operations through a firewall](#).....52
- [Setting up the storage units](#).....53
- [Enabling the distributed segment processing](#)..... 53
- [Enabling the advanced load balancing and link failover](#).....54
- [Enabling the DD Boost-over-FC service](#).....55
- [Validating and troubleshooting the database and Data Domain system connection](#).....57

Licensing the Data Domain system

Note

The Data Domain administrator must configure the Data Domain system for DD Boost operations. This chapter provides examples of the basic configurations. The Data Domain documentation provides details on the Data Domain system configurations.

You need the proper Data Domain licenses, such as the Data Domain Boost or replication license, to use the database application agent software.

Contact your Data Domain representative for more information and to purchase licensed features.

The *Data Domain Operating System Administration Guide* provides details about all the licensed features and how to display and enable Data Domain licenses.

Enabling DD Boost on a Data Domain system

You can enable DD Boost on a Data Domain system through the `ddboost enable` command or from the Data Domain System Manager on the **Data Management > DD Boost** page as described in the *Data Domain Operating System Administration Guide*.

Note

DD Boost requires a separate license.

Use the Data Domain command line interface to complete the required administration tasks. The *Data Domain Operating System Command Reference Guide* provides details about the commands.

Procedure

1. On the Data Domain system, log in as an administrative user.
2. To verify that the file system is enabled and running, run the following command:

```
# fileSYS status
```

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

To enable the file system, run the following command:

```
# fileSYS enable
```

3. To verify that the DD Boost license is enabled, run the following command:

```
# license show
```

```
Feature licenses:
##  License Key           Feature
--  -
1   ABCD-EFGH-IJKL-MNOP   DDBOOST
--  -
```

If the DD Boost license is disabled, run the following command to add the DD Boost license by using the license key that Data Domain provided:

```
# license add <license_key>
```

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

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

Note

The username, password, and role must be set up on the Data Domain system as described in the *Data Domain Operating System Administration Guide*. The username and password are case-sensitive and must match the username and password provided in the procedure [Configuring the lockbox](#) on page 103.

To establish the username and password, run the following commands:

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

[Changing the DD Boost access rights](#) on page 51 provides information about how changing a username and access rights affects the operations on a Data Domain system.

5. To enable DD Boost, run the following command:

```
# ddbboost enable
```

```
DD Boost enabled
```

6. To verify that DD Boost is enabled, run the following command:

```
# ddbboost status
```

Changing the DD Boost access rights

By default, when the DD Boost service is first enabled on a Data Domain system, the service is accessible to all database servers. You can use the `ddbboost access` command to override this default and restrict the access to specific database servers.

For example, the Data Domain administrator can run the following commands to remove the default access permission for all servers and add new access permissions for two specific database servers, `dbserver1.datadomain.com` and `dbserver2.datadomain.com`. The *Data Domain Operating System Command Reference Guide* provides details about the commands.

```
# ddbboost disable
# ddbboost access del clients *
# ddbboost access add clients dbserver1.datadomain.com
dbserver2.datadomain.com
# ddbboost enable
```

These commands establish a set of access controls that enable DD Boost access only to the two database servers, `dbserver1.datadomain.com` and `dbserver2.datadomain.com`.

Consider the following guidelines when you change the DD Boost access rights:

- Ensure that no backup operations are running to the Data Domain system when you change any access rights. You can run the `ddbboost disable` command to prevent operations while access is being changed.
- Specify only a fully qualified domain name, IP address, or resolvable DNS name for the client when modifying the client access control list.
- After the access rights are changed, you can run the `ddbboost enable` command to enable DD Boost and the access rights will take effect.

You can run the `ddbboost clients show` command to verify which database hosts have the DD Boost access rights. If the command output is simply `*`, then all database servers have the access rights. For example:

```
# ddbboost clients show

DD Boost access allowed from the following clients
*
```

```
# ddbboost clients show

DD Boost access allowed from the following clients:
dbserver1.datadomain.com
dbserver2.datadomain.com
```

Enabling encryption over a WAN connection

The database application agent provides support for DD Boost clients to have in-flight data encryption with a Data Domain 5.5 or later operating system over a WAN connection.

To enable the in-flight data encryption over a WAN connection, configure the Data Domain system with either medium-strength or high-strength encryption and set the authentication mode to anonymous. Currently, the product only supports anonymous as the authentication mode. For example, run the following `ddbboost` command to set the required in-flight data encryption for the client systems:

```
ddbboost clients add <client_list> [encryption-strength {medium | high} authentication-mode anonymous]
```

The configuration is transparent to the application agent. The latest *Data Domain Boost Administration Guide* provides details.

Enabling the DD Boost operations through a firewall

The Data Domain system as initially configured does not operate through a firewall, neither for a database server connection to a Data Domain system nor for one Data Domain system connection to another. If you need the Data Domain system to operate through a firewall, contact your network support provider.

The following ports must be open in a firewall to enable DD Boost backups and optimized duplication:

- TCP 2049 (NFS)

- TCP 2051 (Replication)
- TCP 111 (NFS portmapper)
- TCP xxx (select a port for NFS mountd, where the default MOUNTD port is 2052)

Setting up the storage units

One or more storage units must be created on each Data Domain system that will be used with the database application agent. 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.

Note

Storage unit names are case-sensitive.

You must provide the storage unit name when you configure the operations with the database application agent. [Product Configuration](#) on page 77 provides more information.

You can create a storage unit through the `ddbboost storage-unit` command or from the Data Domain System Manager on the **Data Management > DD Boost** page as described in the *Data Domain Operating System Administration Guide*.

For example, you can run the following command on the Data Domain system for each storage unit that you want to create:

```
# ddbboost storage-unit create <storage_unit_name> user <username>
```

You can run the following command to list the status of the storage units:

```
# ddbboost storage-unit show
```

Name	Pre-Comp (GiB)	Status
SU_ABCDE03	5.8	RW
SU_ABCDE5	9.8	RW/Q

D	: Deleted
Q	: Quota Defined
RO	: Read Only
RW	: Read Write

You must create at least one storage unit on each Data Domain system that you will use with the database application agent. You can share a storage unit on a Data Domain system among multiple database hosts.

The storage on a Data Domain system can be provisioned through optional quota limits for a storage unit. Quota limits can be specified either when a storage unit is created or later through separate commands. The *Data Domain Operating System Command Reference Guide* provides details about the `ddbboost` command.

Enabling the distributed segment processing

Distributed segment processing is a DD Boost software feature that uses the DD Boost library on the database server and the Data Domain software on the DDR. The

database application agent loads the DD Boost library during backup and restore operations.

You must configure the distributed segment processing option on the Data Domain system. The option setting applies to all the database servers and all the software that uses DD Boost on this Data Domain system.

You can manage the distributed segment processing through the `ddboost option` command or from the Data Domain System Manager on the **Data Management > DD Boost** page as described in the *Data Domain Operating System Administration Guide*.

To confirm whether or not DD Boost has distributed segment processing enabled, you can run the command `ddboost option show`.

To configure the distributed segment processing option, you can run the following command:

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

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

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

Note

You cannot enable distributed segment processing on Solaris SPARC systems except T4 and T5.

Enabling the advanced load balancing and link failover

Note

This topic applies only if you use an Ethernet connection, not Fibre Channel, for backup and restore operations with the database application agent.

The advanced load balancing and link failover feature enables the combination of multiple Ethernet links into a group and the registration of only one interface on the Data Domain system with the database application agent.

The Data Domain documentation provides details about the features and benefits of advanced load balancing and link failover.

If an interface group is configured when the Data Domain system receives data from the DD Boost client, the data transfer is load balanced and distributed as separate jobs on the private network, providing greater throughput, especially for customers who use multiple 1 GbE connections.

You can manage advanced load balancing and link failover through the `ddboost ifgroup` command or from the Data Domain System Manager on the **Data Management > DD Boost** page as described in the *Data Domain Operating System Administration Guide*.

You can perform the following steps to create an interface group on the Data Domain system by adding existing interfaces to the group and registering the Data Domain

system with the database application agent. After the interface group is set up, you can add or delete interfaces from the group.

Procedure

1. To add the interfaces into the group, run the `ddboost ifgroup` command. The interfaces must have been created with the `net` command. For example:

```
# 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
```

This example assumes that no additional named interface groups have been created and uses the default interface group.

2. Select one interface on the Data Domain system to register with the database application agent. Create a failover aggregated interface and register that interface with the database application agent. The *Data Domain Operating System Administration Guide* describes how to create a virtual interface for link aggregation.

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

3. To enable the feature on the Data Domain system, run the following command:

```
# ifgroup enable
```

4. To verify the configuration, run the following command:

```
# 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

Enabling the DD Boost-over-FC service

DD OS release 5.3 and later provides support for the Data Domain Fibre Channel (DFC or FC) mechanism of communication between the DD Boost library and the Data Domain system.

Note

Support of the DD Boost-over-FC service with the database application agent software requires that a presales FC qualifier has been submitted and approved. Contact your Data Domain representative for more details.

The Data Domain documentation provides details about the features and benefits of the DD Boost-over-FC service.

The Data Domain system must have an HBA that the DD Boost-over-FC service supports. The *Data Domain Operating System Command Reference Guide* and *Data*

Domain Operating System Administration Guide provide information about using the `scsitarget` command for managing the SCSI target subsystem.

The Data Domain administrator can complete the following steps to configure the DD Boost-over-FC service.

Procedure

1. To enable the DD Boost-over-FC service, run the following command:

```
ddboost option set fc enabled
```

2. To optionally set the `dfc-server-name`, run the following command:

```
ddboost fc dfc-server-name set <server_name>
```

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

- lowercase letters (a–z)
- uppercase letters (A–Z)
- digits (0–9)
- underscore (_)
- dash (-)

Note

The dot or period character (.) is not valid within a `dfc-server-name`. You cannot use the fully qualified domain name of a Data Domain system as the `dfc-server-name`.

3. To create a DD Boost FC group, run the following command:

```
ddboost fc group create <group_name>
```

For example:

```
ddboost fc group create lab_group
```

4. To configure the device set of the DD Boost FC group, run the following command:

```
ddboost fc group modify <group_name> device-set count <count>
endpoint {all | none | <endpoint_list>}
```

For example:

```
ddboost fc group modify lab_group device-set count 8 endpoint
all
```

5. To add initiators to the DD Boost FC group, run the following command:

```
ddboost fc group add <group_name> initiator <initiator_spec>
```

For example:

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

6. Verify that the DFC devices are visible on the database server.
7. Ensure that the user who performs the backups and restores has the required permission to access the DFC devices.

[Configuring product operations over FC and IP networks](#) on page 88 provides details about configuring the database application agent to use the FC network connection.

DD Boost-over-FC path management

The ifgroup-based advanced load balancing and link failover mechanism described in a previous topic 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 documentation provides more details.

Validating and troubleshooting the database and Data Domain system connection

Depending on the type of network connection being used, you can run the appropriate command to validate the communication between the database server host and the Data Domain system:

- If you have a DD Boost-over-IP system, you can log in to the database server and run the `rpcinfo` command if the command is available on the system. For example:

```
# rpcinfo -p <Data_Domain_system_hostname>
```

The command output must include the ports listed in [Enabling the DD Boost operations through a firewall](#) on page 52. For example:

```
# rpcinfo -p <Data_Domain_system_hostname>
```

```

program vers proto  port  service
 100000    2    tcp    111   portmapper
 100000    2    udp    111   portmapper
 100024    1    udp    779   status
 100024    1    tcp    782   status
537220272  2    tcp    3006
 100005    1    tcp    2052  mountd
 100005    1    udp    2052  mountd
 100005    2    tcp    2052  mountd
 100005    2    udp    2052  mountd
 100005    3    tcp    2052  mountd
 100005    3    udp    2052  mountd
 100003    3    tcp    2049  nfs
 100003    3    udp    2049  nfs
285824256  1    udp    709
537329792  1    tcp    3007
537220001  2    tcp    2051
537220001  3    tcp    2051
537220439  1    tcp    695
537220017  1    tcp    727

```

- If you have a DD Boost-over-FC system, you can log in to the database server and run the appropriate command to verify that the DFC devices are visible on the client.

The *Data Domain Operating System Command Reference Guide* provides details about the supported commands.

You can use the `ddbadmin` command to verify the username, password, and valid access permissions for the Data Domain system. [Configuring the lockbox](#) on page 103 provides more details about how to use the command.

The Knowledgebase Article number 201919, titled *How to troubleshoot DataDomain DDBoost connectivity and performance*, provides information about how to use the `ddpconnchk` tool to troubleshoot specific DD Boost issues. The article is available on the Support website.

CHAPTER 3

Product Installation

This chapter includes the following topics:

- [Road map to install or update the software](#).....60
- [AIX installation](#)..... 62
- [HP-UX installation](#)..... 65
- [Linux installation](#).....67
- [Solaris installation](#)..... 69
- [Microsoft Windows installation](#)..... 71
- [Software components](#)..... 74

Road map to install or update the software

You must perform the required steps to install or update the database application agent software on the database or application server host.

Before you begin

You must install the same version of the database application agent on each node in a high-availability environment.

During an upgrade of the software, ensure that no backups or restores are running on either the client or each node in the high-availability environment. In a DB2 archived log backup configuration, ensure that the old loaded vendor library is cleaned up, as described in [DB2 backups of transaction logs](#) on page 143.

The database application agent supports coexistence with the following software:

- ProtectPoint file system agent version 4.0 or later.
- Any other backup product used to protect data that the database application agent does not protect.

The database application agent does not support coexistence of ProtectPoint with the NetWorker client, but supports coexistence of DD Boost with the NetWorker client with some restrictions. The *Data Domain Boost for Enterprise Applications and ProtectPoint Database Application Agent Version 4.0 Release Notes* provides more details.

Procedure

1. Ensure that the database or application server host contains the supported database or application server software, installed and functioning in a supported environment.

The online software compatibility guide, available at <http://compatibilityguide.emc.com:8080/CompGuideApp/>, describes the supported software and operating system versions.

2. Ensure that you have reviewed the information in the *Data Domain Boost for Enterprise Applications and ProtectPoint Database Application Agent Version 4.0 Release Notes*.
3. Ensure that the Data Domain system has been prepared according to [Data Domain System Configuration for DD Boost Operations](#) on page 49.
4. If you will perform ProtectPoint operations for a database on a VMAX system, ensure that the VMAX and Data Domain systems have been configured according to instructions in the *ProtectPoint Version 4.0 Primary and Protection Storage Configuration Guide*.
5. If you will perform ProtectPoint with RecoverPoint operations for a database on an XtremIO system, ensure that the XtremIO and Data Domain systems and the RecoverPoint components have been configured according to instructions in the *ProtectPoint Version 4.0 Primary and Protection Storage Configuration Guide*.
6. Ensure that you have operating system root or administrator privileges on the database or application server host.
7. Obtain the required software licenses.

Contact your sales representative for more details about the required licenses for the environment.

8. If you are updating from Data Domain Boost (DD Boost) for Oracle Recovery Manager (RMAN), obtain the required DDBEA license. [Migrating an Oracle configuration from DD Boost for RMAN 1.x or later](#) on page 220 provides details on how to migrate the Oracle configuration to the database application agent 4.0.
9. If you are updating from the database application agent 1.0 (formerly known as DDBDA 1.0) or the database application agent 2.x or 3.x, uninstall the database application agent according to the instructions in the appropriate guide:
 - *Data Domain Boost for Databases and Applications 1.0 Administration Guide*
 - Version 2.0, 2.5, 3.0, or 3.5 of *Data Domain Boost for Enterprise Applications and ProtectPoint Database Application Agent Installation and Administration Guide*

Do not use any upgrade option of the installer on a Linux or UNIX platform, for example, the `rpm -U` command on Linux.

NOTICE

After an update of the database application agent from a pre-4.0 release to release 4.0 on Linux or UNIX, the root user must run the `ddbadmin -U` command if the lockbox is in a nondefault location. [Configuring the lockbox](#) on page 103 provides details about lockbox requirements.

As an alternative on Windows, instead of uninstalling the previous version of the database application agent, you can run a direct update procedure with the Windows installation wizard.

10. Download the database application agent 4.0 software package from the Support website, and then extract the installation package from the file.

You must uncompress the downloaded file twice. First, uncompress the file by using `WinZip` on Windows (recommended) or an `unzip` utility on UNIX/Linux that supports encryption. Then uncompress the resulting file on the application host by using `WinZip` again on Windows or the `gunzip` utility on UNIX/Linux.

The following examples show the steps to download and extract the software package.

Example: Preparing for the database application agent installation on Windows

You can complete the following steps on Windows to download and extract the database application agent software package.

- a. Download the database application agent software package to the Windows host.
- b. Use `WinZip` to uncompress the file.
- c. Use `WinZip` again to uncompress the resulting zip file.

Example: Preparing for the database application agent installation on AIX

You can complete the following steps on AIX to download and extract the database application agent software package.

- a. Download the `dbappagent40_aixpower.tar.gz` file to the AIX application host.

- b. Uncompress and extract the database application agent package:

```
gunzip dbappagent40_aixpower.tar.gz
tar -xvpBf dbappagent40_aixpower.tar
```

11. Install the downloaded the database application agent 4.0 software by following the installation instructions in this chapter.
12. Enable ProtectPoint operations on UNIX or Linux by running the following command to start the snapshot agent:

```
/opt/dpsapps/dbappagent/bin/ddbsm start
```

AIX installation

You must complete the required procedures to install and uninstall the database application agent software on AIX. The following topics provide detailed instructions.

Installing the software on AIX

You can install the database application agent software on AIX by running the `installp` command line interface (CLI) program or the AIX System Management Interface Tool (SMIT), which is a graphical user interface (GUI) program. In a supported cluster, you must install the software on each node that will perform backups and recovery.

Procedure

1. Complete the preparation tasks in [Road map to install or update the software](#) on page 60. Ensure that you log in as the root user and you are in the correct directory, which contains the downloaded software installation files.

Note

If you do not start the installation from the correct directory, the installation might fail.

2. Run either the `installp` CLI program or the SMIT GUI program:
 - To run the CLI program, type the following command:

```
installp -a -d /dir_pathname EMCdbappagent.rte
```

where `/dir_pathname` is the complete pathname of the directory that contains the software installation files.

To verify that the installation succeeded, type the following command:

```
lslpp -L all | grep -i emcdbappagent
```

```
EMCdbappagent.rte 4.0.0.0 C F EMC database app agent
```

If the `lslpp` command output includes `EMCdbappagent.rte 4.0.0.0`, then the installation succeeded.

- To run the SMIT GUI program, perform the following steps:
 - a. Type the following command:

```
smitty install_latest
```

- b. In the **Entry Field**, type the complete pathname of the directory that contains the software installation software files.
- c. Select the option **SOFTWARE to install**.
- d. Type **yes** in response to the following prompts:

```
Accept new license agreements?
Preview new license agreements?
```

- e. To display the list of the software packages, select **F4=List**.
- f. To install the software, select **EMCdbappagent.rte**.
- g. Select **Install and Update Software**.
- h. To begin the installation, press **Enter**.

The installation on AIX stores the different types of software files in the directories shown in the following table.

Table 5 Software installation directories on AIX

Types of installed files or directories	Installation directory
Executable files	/opt/dpsapps/dbappagent/bin
Configuration file templates	/opt/dpsapps/dbappagent/config
Debug log files directory	/opt/dpsapps/dbappagent/logs, linked to /var/opt/ddbda/logs
Library files	/opt/dpsapps/dbappagent/lib/lib64

The installation creates a number of symbolic links as described in [Software links created during installation](#) on page 75.

3. If you will perform operations over a Fibre Channel (FC) connection, you can use either the DFC driver that is packaged with the database application agent or the AIX SCSI generic device driver. Use of the SCSI generic device driver does not require the installation of any drivers.

To check for the type of driver that is installed, you can run the `lsdev` command and review the command output:

- For the DFC driver that is packaged with the database application agent, the command output includes the device names as `DDdfc*` and the type as `Data Domain DDdfc Release 1.0.0.4`. For example:

```
lsdev
```

```
DDdfc    Available    Data Domain DDdfc Release 1.0.0.4
DDdfc1   Available    Data Domain DDdfc Release 1.0.0.4
DDdfc2   Available    Data Domain DDdfc Release 1.0.0.4
DDdfc3   Available    Data Domain DDdfc Release 1.0.0.4
```

- The command output for the SCSI generic device driver includes the device names as `hdisk*` and the type as `Other FC SCSI Disk Drive`. For example:

```
lsdev

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

Note

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

If you want to use the DFC driver that is packaged with the database application agent but the driver is not installed, install the driver as follows:

- Extract the driver package, `DDdfc.rte.1.0.0.4.bff`, from the AIX software package:

```
gunzip < dbappagent40_aixpower.tar.gz | tar xvf -
```

- Install the driver:

```
installp -d /dir_pathname/DDdfc.rte.1.0.0.4.bff all
```

- To confirm that the driver is installed, run the `lsdev` command and review the command output, as previously described.

To configure and integrate the FC connection with Data Domain, follow the instructions in the *Data Domain Fibre Channel Configuration and Integration with Data Domain Boost for Enterprise Applications Database Application Agent Technical Notes*. The document is available on the Support website at <https://support.emc.com>.

- Verify the installed version of the product software by running one of the following commands, where `file_name` is the complete pathname of the `ddbadmin` program file:

```
what file_name
lslpp -I all | grep -i emcdbappagent
```

- To ensure that the DBA can perform backup and restore operations as a non-root user, follow the configuration instructions in this Knowledgebase article:

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

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

- Configure the database application agent software by following the instructions in [Product Configuration](#) on page 77.

Uninstalling the software on AIX

You can uninstall the database application agent software on AIX by running the `installp` command or the SMIT GUI program. In a supported cluster, you must perform the uninstall procedure on each node that contains the software.

Procedure

1. Ensure that no database or application backups are running.
2. To uninstall the software, use one of the following methods as the root user:

Note

You do not need to shut down a database to uninstall the software.

- Use the CLI by typing the following command:

```
installp -u EMCdbappagent.rte
```

- Use the SMIT GUI program:
 - a. Type the following `smitty` command:

```
smitty remove
```

- b. To display a list of the installed software packages, select **F4=List**.
- c. Select the package to uninstall:

EMCdbappagent.rte
- d. Set the **PREVIEW Only** option to **No**.
- e. To uninstall the software, press **Enter**.
- f. Exit the SMIT GUI program.

The uninstall procedure does not remove certain files and directories that contain logs and lockbox files. You must manually remove these items after saving a copy, if required.

HP-UX installation

You must complete the required procedures to install and uninstall the database application agent software on HP-UX. The following topics provide detailed instructions.

Installing the software on HP-UX

You can install the database application agent software on HP-UX by using the `swinstall` utility to run the command line interface (CLI) or the graphical user interface (GUI) program. In a supported cluster, you must install the software on each node that will perform backups and recovery.

Procedure

1. Complete the preparation tasks in [Road map to install or update the software](#) on page 60. Ensure that you log in as the root user and you are in the correct directory, which contains the downloaded software installation files.

Note

If you do not start the installation from the correct directory, the installation might fail.

2. To run either the CLI or GUI program, type the `swinstall` command:

- To run the `swinstall` CLI program, type the following command:

```
swinstall -x mount_all_filesystems=false -s /dir_pathname/
EMCdbappagent.pkg EMCdbappagent
```

where `/dir_pathname` is the complete pathname of the directory that contains the software installation files.

- To run the `swinstall` GUI program, type the following command:

```
swinstall -x mount_all_filesystems=false -i -s /
dir_pathname/EMCdbappagent.pkg EMCdbappagent
```

where `/dir_pathname` is the complete pathname of the directory that contains the software installation files. Perform the following steps in the GUI program:

- From the **Actions** menu, select **Install (analysis)**.

When the analysis is complete, a `Ready with Warnings` message appears. The message is normal.

- To continue the installation, click **OK**.

The installation on HP-UX stores the different types of software files in the directories shown in the following table.

Table 6 Software installation directories on HP-UX

Types of installed files or directories	Installation directory
Executable files	/opt/dpsapps/dbappagent/bin
Configuration file templates	/opt/dpsapps/dbappagent/config
Debug log files directory	/opt/dpsapps/dbappagent/logs, linked to /var/opt/dbbda/logs
Library files	/opt/dpsapps/dbappagent/lib/hpux64

The installation creates a number of symbolic links as described in [Software links created during installation](#) on page 75.

The installation on HP-UX stores informational messages including installation errors in the `/var/adm/sw/swagent.log` file. If an error occurs during the installation, check this file to obtain details about the error.

3. Verify the installed version of the product software by running one of the following commands, where *file_name* is the complete pathname of the `ddbadmin` program file:

```
what file_name
swlist | grep db
```

4. Configure the database application agent software by following the instructions in [Product Configuration](#) on page 77.

Uninstalling the software on HP-UX

You can uninstall the database application agent software on HP-UX by running the `swremove` command or GUI program. In a supported cluster, you must perform the uninstall procedure on each node that contains the software.

Procedure

1. Ensure that no database or application backups are running.
2. To uninstall the software, use one of the following methods as the root user:

Note

You do not need to shut down a database to uninstall the software.

- Use the CLI by typing the following command:

```
swremove EMCdbappagent
```

- Use the `swremove` GUI program:
 - a. Type the following `swremove` command:

```
swremove -i EMCdbappagent
```

- b. Select **Actions > Remove (analysis)**.
- c. To complete the uninstall, click **OK** when the system analysis is complete.
- d. To confirm the uninstall, click **Yes**.

The uninstall procedure does not remove certain files and directories that contain logs and lockbox files. You must manually remove these items after saving a copy, if required.

Linux installation

You must complete the required procedures to install and uninstall the database application agent software on Linux. The following topics provide detailed instructions.

Installing the software on Linux

You can install the database application agent software on Linux by running the `rpm` command. In a supported cluster, you must install the software on each node that will perform backups and recovery.

Procedure

1. Complete the preparation tasks in [Road map to install or update the software](#) on page 60. Ensure that you log in as the root user and you are in the correct directory, which contains the downloaded software installation files.

Note

If you do not start the installation from the correct directory, the installation might fail.

2. On a RHEL 6 Linux platform, ensure that you have downloaded and installed the `compat-libstdc++-33` package.
3. To install the software, type the required `rpm` command:

- On Linux x64:

```
rpm -ivh emcdbappagent-4.0.0.0-1.x86_64.rpm
```

- On Linux Power PC 64-bit big-endian:

```
rpm -ivh emcdbappagent-4.0.0.0-1.ppc64.rpm
```

- On Linux Power PC 64-bit little-endian:

```
rpm -ivh emcdbappagent-4.0.0.0-1.ppc64le.rpm
```

4. To verify that the installation was successful, type the `rpm -aq` command:

```
rpm -aq | grep -i emc
```

The command output must include the following line:

```
emcdbappagent-4.0.0.0-1
```

The installation on Linux stores the different types of software files in the directories shown in the following table.

Table 7 Software installation directories on Linux

Types of installed files or directories	Installation directory
Executable files	/opt/dpsapps/dbappagent/bin
Configuration file templates	/opt/dpsapps/dbappagent/config
Debug log files directory	/opt/dpsapps/dbappagent/logs, linked to /var/opt/dbda/logs

Table 7 Software installation directories on Linux (continued)

Types of installed files or directories	Installation directory
Library files	/opt/dpsapps/dbappagent/lib/lib64 (Linux x64, Linux Power PC little-endian) /opt/dpsapps/dbappagent/lib (Linux Power PC big-endian)

The installation creates a number of symbolic links as described in [Software links created during installation](#) on page 75.

5. Verify the installed version of the product software by running one of the following commands, where *file_name* is the complete pathname of the `ddbadmin` program file:

```
strings file_name | grep "@(#)"
rpm -aq | grep -i emc
```

6. Configure the database application agent software by following the instructions in [Product Configuration](#) on page 77.

Uninstalling the software on Linux

You can uninstall the database application agent software on Linux by running the `rpm` command. In a supported cluster, you must perform the uninstall procedure on each node that contains the software.

Procedure

1. Ensure that no database or application backups are running.
2. To uninstall the software, type the following command as the root user:

Note

You do not need to shut down a database to uninstall the software.

```
rpm -e emcdbappagent-4.0.0.0-1
```

The uninstall procedure does not remove certain files and directories that contain logs and lockbox files. You must manually remove these items after saving a copy, if required.

Solaris installation

You must complete the required procedures to install and uninstall the database application agent software on Solaris. The following topics provide detailed instructions.

Installing the software on Solaris

You can install the database application agent software on Solaris by running the `pkgadd` command. In a supported cluster, you must install the software on each node that will perform backups and recovery.

Procedure

1. Complete the preparation tasks in [Road map to install or update the software](#) on page 60. Ensure that you log in as the root user and you are in the correct directory, which contains the downloaded software installation files.

Note

If you do not start the installation from the correct directory, the installation might fail.

If the Solaris system has Solaris zones (containers) and the software is to run on a sparse root zone, install the software on the global zone and on each required sparse root zone.

2. Verify that the `basedir` variable setting is `basedir=default` in the `/var/sadm/install/admin/default` file.
3. To install the software, type the following `pkgadd` command:

```
pkgadd -d /dir_pathname EMCdbappagent
```

where `/dir_pathname` is the complete pathname of the directory that contains the software package.

4. Type `y` when prompted whether or not to continue the installation.

The installation on Solaris stores the different types of software files in the directories shown in the following table.

Table 8 Software installation directories on Solaris

Types of installed files or directories	Installation directory
Executable files	<code>/opt/dpsapps/dbappagent/bin</code>
Configuration file templates	<code>/opt/dpsapps/dbappagent/config</code>
Debug log files directory	<code>/opt/dpsapps/dbappagent/logs</code> , linked to <code>/var/opt/dbda/logs</code>
Library files	<code>/opt/dpsapps/dbappagent/lib/amd64</code> (Solaris AMD) <code>/opt/dpsapps/dbappagent/lib/sparcv9</code> (Solaris SPARC)

The installation creates a number of symbolic links as described in [Software links created during installation](#) on page 75.

5. Verify the installed version of the product software by running one of the following commands, where `file_name` is the complete pathname of the `ddbadmin` program file:

Note

The `what` command is available on Solaris 11 only.

```
what file_name
pkginfo -l EMCdbappagent
```

6. Configure the database application agent software by following the instructions in [Product Configuration](#) on page 77.

NOTICE

On Solaris x64, you might need to run the `ddbadmin -U` command with the `LOCKBOX_IMPORT` option if both of the following conditions are true:

- You have updated to the latest version of the database application agent.
- The system contains an existing lockbox that was created with the pre-4.0 database application agent.

[Updating the lockbox](#) on page 114 provides details.

Uninstalling the software on Solaris

You can uninstall the database application agent software on Solaris by running the `pkgrm` command. In a supported cluster, you must perform the uninstall procedure on each node that contains the software.

Procedure

1. Ensure that no database or application backups are running.
2. To uninstall the software, type the following command as the root user:

Note

You do not need to shut down a database to uninstall the software.

To uninstall the software on Solaris zones, first uninstall the software on the global zone and then uninstall the software on each required sparse root zone.

```
pkgrm EMCdbappagent
```

3. Type `y` when prompted.

The uninstall procedure does not remove certain files and directories that contain logs and lockbox files. You must manually remove these items after saving a copy, if required.

Microsoft Windows installation

You must complete the required procedures to install and uninstall the database application agent software on Windows. You can also modify, repair, and remove an existing installation on Windows. The following topics provide detailed instructions.

Installing the software on Windows

You can install the database application agent software on Windows by running the software installer, `emcdbappagent.exe`. In a supported cluster, you must install the software on each node that will perform backups and restores.

Procedure

1. Complete the preparation tasks in [Road map to install or update the software](#) on page 60. Ensure that you log in as an administrator and you are in the correct directory, which contains the downloaded software installation files.

Note

If you do not start the installation from the correct directory, the installation might fail.

2. Run the Windows installer, `emcdbappagent.exe`, and then follow the instructions provided by the installation wizard.

If you are updating from the database application agent 1.0, formerly known as DDBDA 1.0, to the database application agent 4.0 on Windows, you can select to run a direct update procedure through the installation wizard.

3. To exit the installer, click **Finish**.

The installation on Windows stores all the executable files, configuration file templates, and library files in the directory `C:\Program Files\DPSAPPS\DBAPPAGENT\bin\`. The installation also creates the debug log files directory, `C:\Program Files\DPSAPPS\DBAPPAGENT\logs\`.

The installation creates a number of symbolic links as described in [Software links created during installation](#) on page 75.

4. Verify that the system Path environment variable includes the software installation directory. For example:
 - a. From the **Start** menu, select **Computer**.
 - b. From the context menu, select **System properties**.
 - c. Select **Advanced system settings**.
 - d. On the **Advanced** tab, click **Environment Variables**.
 - e. Under **System Variables**, verify the Path variable. The Path variable must include the software installation directory, `C:\Program Files\DPSAPPS\DBAPPAGENT\bin`.

Note

The directory pathname can include spaces, but there cannot be spaces before or after the pathname.

5. To verify the installed version of the product software, use one of the following methods:
 - Check the system properties for the `ddbadmin` program:
 - a. Right-click the file `C:\Program Files\DPSAPPS\DBAPPAGENT\bin\ddbadmin`.

- b. Select **Properties**.
 - c. On the **Details** tab, verify the product version.
 - Check the Control Panel for the `ddbmadmin` program:
 - a. From the **Start** menu, select **Control Panel**.
 - b. Select **Programs > Programs and Features**.
 - c. In the **Uninstall or change a program** window, verify the product version.
6. Configure the database application agent software by following the instructions in [Product Configuration](#) on page 77.

Uninstalling the software on Windows

You can uninstall the database application agent software on Windows by using the Windows installer or Windows Control Panel. In a supported cluster, you must perform the uninstall procedure on each node that contains the software.

Procedure

1. Log in as the Windows system administrator on the software host.
2. Ensure that no database or application backups are running.
3. To uninstall the software, use one of the following methods:

Note

You do not need to shut down a database to uninstall the software. Before you complete the uninstall, save a copy of any configuration files, if required.

- Windows installer method:
 - a. Run the Windows installer, `emcdbappagent.exe`.
 - b. On the **Change, repair, or remove installation** page, select **Remove**, and then click **Next**.
 - c. To uninstall the software, in the **Remove the Program** dialog box, click **Remove**.
- Windows Control Panel method:
 - a. In the **Control Panel** window, select **Add or Remove Programs** or **Programs and Features**, depending on the Microsoft Windows version.
 - b. In the **Add or Remove Programs** window, select **database app agent for DBEA and ProtectPoint**, and then click **Remove**.

The uninstall procedure does not remove certain files and directories that contain logs and lockbox files. You must manually remove these items after saving a copy, if required.

Software components

The following table lists the major software components installed on the database or application host during an installation of the database application agent software.

Table 9 Product software components

Component name	Description
Components used with all applications:	
ddbmadmin	Program that is used for lockbox operations with the database application agent.
ddbsm ddbsm_shutdown ddbsmdd	Snapshot management files that are used for ProtectPoint operations with the database application agent.
libddbprotectpoint_rp.xx libcurl.xx libddbrestclient.xx	Snapshot control library files that are used for ProtectPoint with RecoverPoint operations.
libddbprotectpoint_v3.xx	Snapshot control library file that is used for ProtectPoint with VMAX operations.
libddvdisk	vdisk SDK library file that is used for ProtectPoint operations.
liblocktbl.so libccme*.xx libcryptocme*.xx libCSP*.so	Library files that are used for lockbox operations on UNIX and Linux x64.
libccme*.so libcryptocme*.xx libEnshroud-BSAFEMES.so libLockbox.so libssp.so.0 libcrypto.so libssl.so	Library files that are used for lockbox operations on Linux Power PC.
locktbl.dd ccme*.dll cryptocme*.xxx CSP*.dll	Library files that are used for lockbox operations on Windows.

Table 9 Product software components (continued)

Component name	Description
libDDBoost.xx	DD Boost library file that is used by the database application agent.
Components used with DB2 only:	
db2_ddbda.cfg	Configuration file template for DB2 operations.
ddbmdb2adutil	Executable that is used for DB2 database backup and log backup retrieval operations.
libddboostdb2.xx	Library that is used for DB2 operations.
Components used with Oracle only:	
ddsbtcn.exe	Executable that is used for Oracle operations on Windows.
oracle_ddbda.cfg	Configuration file template for Oracle operations.
libddboostora.xx	Library that is used for Oracle operations.
Components used with SAP HANA only:	
hdbbackint	Executable that is used for SAP HANA operations.
sap_hana_ddbda.utl	Configuration file template for SAP HANA operations.
Components used with SAP with Oracle only:	
backint	Executable that is used for SAP with Oracle (BR*Tools) operations.
ddsbtcnsap.exe	Executable that is used for SAP with Oracle operations with RMAN on Windows.
libddboostsapora.xx	Library that is used for SAP with Oracle operations with RMAN.
sap_oracle_ddbda.utl	Configuration file template for SAP with Oracle operations.

Software links created during installation

When you install the database application software, the installation process creates a number of symbolic links for the software binaries and libraries.

On UNIX or Linux, the software installation creates symbolic links to the locations of the previous release libraries and binaries. For example, the installation creates the following symbolic links on Linux:

- /opt/ddbda/bin/hdbbackint linked to /opt/dpsapps/dbappagent/bin/hdbbackint
- /opt/ddbda/bin/backint linked to /opt/dpsapps/dbappagent/bin/backint

- `/opt/ddbda/bin/ddbadmin` **linked to** `/opt/dpsapps/dbappagent/bin/ddbadmin`
 - `/usr/lib/ddbda/lib64/libddboostdb2.so` **linked to** `/opt/dpsapps/dbappagent/lib/lib64/libddboostdb2.so`
 - `/usr/lib/ddbda/lib64/libddboostora.so` **linked to** `/opt/dpsapps/dbappagent/lib/lib64/libddboostora.so`
 - `/usr/lib/ddbda/lib64/libddboostsapora.so` **linked to** `/opt/dpsapps/dbappagent/lib/lib64/libddboostsapora.so`
-

Note

The UNIX library directories from the previous release are as follows:

- On AIX: `/usr/lib/ddbda/lib64`
 - On HP-UX: `/usr/lib/ddbda/hpux64`
 - On Solaris AMD: `/usr/lib/ddbda/amd64`
 - On Solaris SPARC: `/usr/lib/ddbda/sparcv9`
-

On Windows, the software installation creates the following library links:

- **Hard link:**
`C:\Program Files\EMC DD Boost\DA\bin\libddboostdb2.dll` **linked to** `C:\Program Files\DPSAPPS\DBAPPAGENT\bin\libddboostdb2.dll`
- **Symbolic links:**
`C:\Program Files\EMC DD Boost\DA\bin\libddboostora.dll` **linked to** `C:\Program Files\DPSAPPS\DBAPPAGENT\bin\libddboostora.dll`
`C:\Program Files\EMC DD Boost\DA\bin\libddboostsapora.dll` **linked to** `C:\Program Files\DPSAPPS\DBAPPAGENT\bin\libddboostsapora.dll`

CHAPTER 4

Product Configuration

This chapter includes the following topics:

- [Road map for configuration](#)..... 78
- [Setting up the configuration file](#)..... 78
- [Configuring product operations over FC and IP networks](#)..... 88
- [Configuring the optimization of ProtectPoint backups for third-party multipathing software](#)..... 89
- [Configuring restores of replicated backups](#)..... 90
- [Configuring ProtectPoint VMAX restores directly from Data Domain](#)..... 96
- [Configuring ProtectPoint VMAX restores from local snapshots](#)..... 96
- [Configuring rollback restores of ProtectPoint backups](#)..... 97
- [Configuring usage limits on Data Domain resources](#)..... 101
- [Configuring the lockbox](#)..... 103
- [Configuring the display and deletion of save set information](#)..... 117
- [Configuring the use of Data Domain Cloud Tier for data movement to the cloud](#)..... 126
- [General troubleshooting tips](#)..... 132
- [ProtectPoint specific troubleshooting tips](#)..... 137

Road map for configuration

Note

The configuration procedures described in this chapter apply to all the supported database servers and all the supported workflows, including the DD Boost, ProtectPoint, and ProtectPoint with RecoverPoint workflows. Review the configuration information in the subsequent chapters for additional procedures that apply to specific database servers.

Use the following road map to configure the database application agent software for backups and restores on the supported database servers.

Procedure

1. Set up the configuration file to be used for the product operations according to [Setting up the configuration file](#) on page 78.
2. Configure the product to use the required network connections according to [Configuring product operations over FC and IP networks](#) on page 88.
3. If required, configure restores of replicated backups according to [Configuring restores of replicated backups](#) on page 90.
4. If required, configure restores of ProtectPoint for VMAX backups directly from a Data Domain system, without the use of a VMAX system, according to [Configuring ProtectPoint VMAX restores directly from Data Domain](#) on page 96.
5. If required, configure rollback restores of ProtectPoint backups according to [Configuring rollback restores of ProtectPoint backups](#) on page 97.
6. Configure any required usage limits on Data Domain resources according to [Configuring usage limits on Data Domain resources](#) on page 101.
7. Configure the lockbox that is used by the product according to [Configuring the lockbox](#) on page 103.
8. If required, configure the display or deletion of save set information according to [Configuring the display and deletion of save set information](#) on page 117.
9. Complete the required application-specific configurations according to the appropriate configuration instructions. Each of the subsequent chapters contains a configuration topic for a specific application and operation type.

Setting up the configuration file

You must set up a configuration file to be used for backups and restores with the database application agent. You must customize a configuration file template that the software installation provides by setting specific parameters in the file.

The software installation provides the following templates for the configuration file:

- `db2_ddbda.cfg`—Template for the DB2 configuration file
- `oracle_ddbda.cfg`—Template for the Oracle configuration file
- `sap_hana_ddbda.utl`—Template for the SAP HANA configuration file
- `sap_oracle_ddbda.utl`—Template for the SAP with Oracle configuration file

The configuration file templates are installed in the following directory:

- On UNIX and Linux: `/opt/dpsapps/dbappagent/config/`
- On Windows: `C:\Program Files\DPSAPPS\DBAPPAGENT\config\`

Make a copy of the required configuration file template, for example, in the original directory or an alternate location, and modify the parameter settings in the file as required. Follow the guidelines in [Syntax rules for the configuration file](#) on page 79.

Note

The uninstall of the product software removes the original configuration file templates.

[Common parameters](#) on page 80 describes the common parameters that you set in the configuration file for backups and restores of all the supported databases and applications, including both DD Boost and ProtectPoint operations.

The following common parameters are mandatory for all operations with the database application agent:

- `DDBOOST_USER`
- `DEVICE_HOST`
- `DEVICE_PATH`

[Common ProtectPoint parameters for VMAX](#) on page 83 describes the common parameters that you set in the configuration file for ProtectPoint backups and restores of DB2, Oracle, and SAP with Oracle database data that resides on a VMAX system.

[Common ProtectPoint with RecoverPoint parameters for XtremIO](#) on page 85 describes the common parameters that you set in the configuration file for ProtectPoint with RecoverPoint backups and restores of DB2, Oracle, and SAP with Oracle database data that resides on an XtremIO system.

The subsequent topics describe additional parameters that you can set in the configuration file for specific operations. Later chapters describe additional database-specific parameters.

Syntax rules for the configuration file

The configuration file includes the following sections:

- **General section**—In this section, the case-sensitive section heading `[GENERAL]` is followed by parameter settings that provide information about the application, the lockbox pathname, and the log files.
- **Primary system section**—In this section, the case-sensitive section heading `[PRIMARY_SYSTEM]` is followed by parameter settings that provide information about the primary Data Domain system.
- **Secondary system section**—In this section, the case-sensitive section heading `[SECONDARY_SYSTEM]` is followed by parameter settings that provide information about the replication of data from the primary Data Domain system to the secondary Data Domain system.

This section is required only for Data Domain replication. The database application agent supports Data Domain MTree replication of data from one Data Domain system to another, but does not provide a mechanism to manage the replication. Use the Mtree replication controls in the Data Domain operating system.

- **RecoverPoint cluster section**—In this section, the case-sensitive section heading `[RP_CLUSTER_1]` is followed by parameter settings that provide information about the primary RecoverPoint cluster.

This section is required only when a RecoverPoint cluster is used with the database application agent.

The configuration file must conform to the following syntax rules:

- Each parameter setting must be in one of the following formats:

```
NAME=value
NAME=value1;value2;value3
```

where:

- *NAME* is the parameter name.
- *value*, *value1*, *value2*, *value3* are the assigned parameter values.
- Parameter names and values are case-sensitive, unless specified otherwise.
- If a parameter value contains a quotation mark, then the value must be enclosed within outer quotes that are different from the inner quote:
 - Use double quotes to enclose a parameter value that contains a single quote. For example: `DEVICE_PATH= "/new's"`
 - Use single quotes to enclose a parameter value that contains a double quote.
- A mandatory parameter must always be set for the specified operation. If an optional parameter is not set, the operation uses the default value of the parameter, if a default value exists.
- Use white spaces as preferred. The database application agent ignores all the white spaces.
- When a line starts with the # symbol, any text on the line is a comment. The database application agent ignores all the comments.

Note

For each parameter that is not required and is not set to a value, ensure that the parameter line starts with the # symbol.

Common parameters

The following table describes the common parameters that the database application agent uses for both DD Boost and ProtectPoint backups and restores of all the supported databases and applications. For each parameter, the table lists the section heading of the configuration file section that contains the parameter.

Table 10 Common parameters

<p>Parameter: CLIENT</p> <p>Section: [GENERAL]</p> <p>Specifies the application hostname or the hostname of the client that is being protected and has backups that are stored on the Data Domain system.</p> <p>Mandatory in a high-availability environment or for a restore to an alternative host. The application-specific chapters provide details on this parameter setting in a high-availability environment.</p> <p>Recommended in all other environments.</p>
--

Table 10 Common parameters (continued)

Note

The `CLIENT` parameter setting for a restore must match the `CLIENT` parameter setting used during the backup.

Valid values:

- Hostname of the local physical host on which the backup or restore runs (default).
- Client hostname.

Parameter: DDBOOST_COMPRESSED_RESTORE**Section: [GENERAL]**

Specifies whether to perform a compressed restore that uses the DD Boost workflow. A compressed restore uses less bandwidth by restoring the backed-up data in a compressed form from the Data Domain system to the application host.

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

Optional.

Valid values:

- FALSE (default).
- TRUE.

Parameter: LOCKBOX_PATH**Section: [GENERAL]**

Specifies the complete directory pathname of the lockbox on the database or application host.

Optional.

Note

The lockbox must be properly configured according to [Configuring the lockbox](#) on page 103.

Valid values:

- Default directory pathname of the lockbox:
 - On UNIX or Linux: `/opt/dpsapps/common/lockbox`
 - On Windows: `C:\Program Files\DPSAPPS\common\lockbox`
- Valid complete directory pathname of the lockbox. For example:
`LOCKBOX_PATH=/opt/lockbox`

Parameter: DDBOOST_USER**Section: [PRIMARY_SYSTEM]**

Specifies the username of the DD Boost user configured on the primary Data Domain system, when this parameter is set in the primary system section of the configuration file. The primary system section has the `[PRIMARY_SYSTEM]` heading.

Mandatory.

Table 10 Common parameters (continued)**NOTICE**

You must set the initial value of the parameter in the configuration file before the `ddbmadmin` command is used for any lockbox procedures. After this initial setting, you can modify the parameter value and rerun the `ddbmadmin` command as described in [Configuring the lockbox](#) on page 103.

Valid values:

- Undefined (default).
- Valid username of the DD Boost user on the primary Data Domain system. For example:

```
DDBOOST_USER=user1
```

Parameter: DEVICE_HOST**Section: [PRIMARY_SYSTEM]**

Specifies the hostname of the primary Data Domain system where the backup is stored, when this parameter is set in the primary system section of the configuration file. The primary system section has the `[PRIMARY_SYSTEM]` heading.

Mandatory.

NOTICE

You must set the initial value of the parameter in the configuration file before the `ddbmadmin` command is used for any lockbox procedures. After this initial setting, you can modify the parameter value and rerun the `ddbmadmin` command as described in [Configuring the lockbox](#) on page 103.

Set the `DEVICE_HOST` parameter to the same value in all the configuration files on a particular client host. All the backups for a given client should be stored under the same `DEVICE_HOST` if you upgrade from the database application agent to NetWorker software in the future.

Valid values:

- Undefined (default).
- Valid hostname as the fully qualified domain name of the primary Data Domain system. For example:

```
DEVICE_HOST=dd.host.com
```

Parameter: DEVICE_PATH**Section: [PRIMARY_SYSTEM]**

Specifies the name of the storage unit or a top-level directory within the storage unit on the primary Data Domain system, when this parameter is set in the primary system section of the configuration file. The primary system section has the `[PRIMARY_SYSTEM]` heading.

Mandatory for a backup or restore on the primary Data Domain system.

Table 10 Common parameters (continued)**NOTICE**

You must set the initial value of the parameter in the configuration file before the `ddbmadmin` command is used for any lockbox procedures. After this initial setting, you can modify the parameter value and rerun the `ddbmadmin` command as described in [Configuring the lockbox](#) on page 103.

Set the `DEVICE_PATH` parameter to the same value in all the configuration files on a particular client host. All the backups for a given client should use the same `DEVICE_PATH` setting if you upgrade from the database application agent to NetWorker software in the future.

Valid values:

- Undefined (default).
- Valid directory name of the storage unit on the primary Data Domain system, without the `/data/coll` prefix. The specified name is case-sensitive.

For example, if the pathname is `/data/coll/su1`, then the valid storage unit name for this parameter is `/su1`:

```
DEVICE_PATH=/su1
```

Common ProtectPoint parameters for VMAX

The following table describes the common parameters that the database application agent uses only for ProtectPoint backups and restores of DB2, Oracle, and SAP with Oracle data that resides on a VMAX system. For each parameter, the table lists the section heading of the configuration file section that contains the parameter.

The subsequent topic describes the common parameters that the database application agent uses for ProtectPoint with RecoverPoint operations with an XtremIO system.

Note

Unless specified otherwise, the following common ProtectPoint parameters are case-insensitive and optional.

Table 11 Common ProtectPoint parameters for VMAX**Parameter: DDVDISK_USER****Section: [PRIMARY_SYSTEM]**

Specifies the vdisk username on the primary Data Domain system.

Mandatory only if the Data Domain user that connects to the DD vdisk services is different from the DD Boost user specified in `DDBOOST_USER`.

Valid values:

- DD Boost username specified in the `DDBOOST_USER` parameter setting (default). For example:

```
DDBOOST_USER=user1
```

- Valid DD vdisk username provided during the creation of the vdisk device pool.

Table 11 Common ProtectPoint parameters for VMAX (continued)

<p>Parameter: DEVICE_POOL</p> <p>Section: [PRIMARY_SYSTEM]</p> <p>Specifies the name of the DD vdisk device pool that provides the backup or restore LUNs.</p> <p>The database application agent uses this parameter for LUN validation purposes. If this parameter is set, then the database application agent confirms that all the backup or restore LUNs involved in the operation are in the specified device pool.</p> <p>Optional.</p> <p>NOTICE</p> <p>If this parameter is set in the configuration file, then you must register the username and password of the DD vdisk device pool with the lockbox. Configuring the lockbox on page 103 provides details.</p> <hr/> <p>Valid values:</p> <ul style="list-style-type: none"> • Undefined (default). • Valid name of a DD vdisk device pool.
<p>Parameter: RESTORE_DEVICE_GROUP</p> <p>Section: [PRIMARY_SYSTEM]</p> <p>Specifies the DD vdisk device group in the vdisk device pool that contains the restore LUNs to use for the restore of a ProtectPoint for VMAX backup directly from Data Domain. The restore directly from Data Domain does not involve a VMAX system.</p> <p><code>RESTORE_DEVICE_POOL</code> in the primary system section of the configuration file (section with the [PRIMARY_SYSTEM] heading) specifies the device pool on a local (primary) Data Domain system. <code>RESTORE_DEVICE_POOL</code> in the secondary system section of the configuration file (section with the [SECONDARY_SYSTEM] heading) specifies the device pool on a remote (secondary) Data Domain system.</p> <p>Mandatory for a restore directly from Data Domain.</p> <p>Configuring ProtectPoint VMAX restores directly from Data Domain on page 96 provides details.</p> <p>Valid values:</p> <ul style="list-style-type: none"> • Undefined (default). • Valid name of a DD vdisk device group.
<p>Parameter: RESTORE_DEVICE_POOL</p> <p>Section: [PRIMARY_SYSTEM]</p> <p>Specifies the name of the DD vdisk device pool to use for the restore of a ProtectPoint for VMAX backup directly from Data Domain. The specified device pool must contain the restore LUNs that are provided on the restore host.</p> <p>Mandatory for a restore directly from Data Domain.</p> <p>Configuring ProtectPoint VMAX restores directly from Data Domain on page 96 provides details.</p> <p>Valid values:</p> <ul style="list-style-type: none"> • Undefined (default). • Valid name of a DD vdisk device pool.

Table 11 Common ProtectPoint parameters for VMAX (continued)

<p>Parameter: RESTORE_TYPE_ORDER</p> <p>Section: [GENERAL]</p> <p>Specifies the type of ProtectPoint restore to perform.</p> <p>The database application agent performs a point-in-time restore by default. This restore mounts the static images through restore LUNs to the recovery host and copies the files to the requested location.</p> <p>You can also specify a rollback restore, which is a LUN-level restore. Configuring rollback restores of ProtectPoint backups on page 97 provides more details.</p> <p>Optional.</p> <p>Valid values:</p> <ul style="list-style-type: none"> • PIT (default). • rollback.
<p>Parameter: SYMM_SNAP_REMOTE</p> <p>Section: [GENERAL]</p> <p>Specifies whether the ProtectPoint backup is an SRDF based backup as described in VMAX replication on page 37.</p> <p>Optional for a ProtectPoint backup from a VMAX system.</p> <p>Valid values:</p> <ul style="list-style-type: none"> • FALSE (default). • TRUE.
<p>Parameter: VMAX_FASTX_RESTORE_SG</p> <p>Section: [PRIMARY_SYSTEM]</p> <p>Specifies the name of the VMAX storage group to use during a ProtectPoint restore to a selected FAST.X or VMAX native restore device on VMAX. By default, the NsrSnapSG storage group is used for a ProtectPoint restore to a VMAX system.</p> <p>Optional for a ProtectPoint restore to a primary VMAX system.</p> <p>Valid values:</p> <ul style="list-style-type: none"> • NsrSnapSG (default). • Valid name of a VMAX storage group.

Common ProtectPoint with RecoverPoint parameters for XtremIO

The following table describes the common parameters that the database application agent uses only for ProtectPoint with RecoverPoint backups and restores of DB2, Oracle, and SAP with Oracle data that resides on an XtremIO system. For each parameter, the table lists the section heading of the configuration file section that contains the parameter.

The previous topic describes the common parameters that the database application agent uses for ProtectPoint operations with a VMAX system.

Note

Unless specified otherwise, the following ProtectPoint with RecoverPoint parameters are case-insensitive and optional.

Table 12 Common ProtectPoint with RecoverPoint parameters for XtremIO

<p>Parameter: DDVDISK_USER</p> <p>Section: [PRIMARY_SYSTEM]</p> <p>Specifies the vdisk username on the primary Data Domain system.</p> <p>Mandatory only if the Data Domain user that connects to the DD vdisk services is different from the DD Boost user specified in <code>DDBOOST_USER</code>.</p> <p>Valid values:</p> <ul style="list-style-type: none"> • DD Boost username specified in the <code>DDBOOST_USER</code> parameter setting (default). For example: <code>DDBOOST_USER=user1</code> • Valid DD vdisk username provided during the creation of the vdisk device pool.
<p>Parameter: DEVICE_POOL</p> <p>Section: [PRIMARY_SYSTEM]</p> <p>Specifies the name of a DD vdisk device pool to use for a ProtectPoint with RecoverPoint restore.</p> <p>If this parameter is set, then the database application agent locates the static images in the specified device pool instead of the device pool that is registered to the RecoverPoint cluster.</p> <p>Optional for a restore.</p> <hr/> <p>Note</p> <p>This parameter is not used for a backup.</p> <hr/> <p>Valid values:</p> <ul style="list-style-type: none"> • Name of the DD vdisk device pool that is registered to the RecoverPoint cluster (default). • Valid name of a DD vdisk device pool.
<p>Parameter: RESTORE_DEVICE_GROUP</p> <p>Section: [PRIMARY_SYSTEM]</p> <p>Specifies the DD vdisk device group in the vdisk device pool that contains the restore LUNs to use for a ProtectPoint with RecoverPoint restore.</p> <p>Mandatory except for rollback restore.</p> <hr/> <p>Note</p> <p>This parameter is not used for a rollback restore.</p> <hr/> <p>Valid values:</p> <ul style="list-style-type: none"> • Undefined (default).

Table 12 Common ProtectPoint with RecoverPoint parameters for XtremIO (continued)

- Valid name of a DD vdisk device group.

Parameter: RESTORE_DEVICE_POOL**Section: [PRIMARY_SYSTEM]**

Specifies the name of the DD vdisk device pool to use for a ProtectPoint with RecoverPoint restore. The specified device pool must contain the restore LUNs.

Mandatory except for rollback restore.

Note

This parameter is not used for a rollback restore.

Valid values:

- Undefined (default).
- Valid name of a DD vdisk device pool.

Parameter: RESTORE_TYPE_ORDER**Section: [GENERAL]**

Specifies the type of ProtectPoint with RecoverPoint restore to perform.

The database application agent performs a point-in-time restore by default. This restore mounts the static images through DD restore LUNs to the recovery host and copies the files to the requested location.

You can also specify a rollback restore. [Configuring rollback restores of ProtectPoint backups](#) on page 97 provides more details.

Optional.

Valid values:

- PIT (default).
- rollback.

Parameter: RP_MGMT_HOST**Section: [RP_CLUSTER_1]**

Specifies the hostname or IP address of the RecoverPoint management host. The database application agent does not support multiple RecoverPoint management hosts.

Mandatory.

NOTICE

You must set the initial value of the parameter in the configuration file before the `ddbmadmin` command is used for any lockbox procedures. After this initial setting, you can modify the parameter value and rerun the `ddbmadmin` command as described in [Configuring the lockbox](#) on page 103.

Valid values:

- Undefined (default).

Table 12 Common ProtectPoint with RecoverPoint parameters for XtremIO (continued)

<ul style="list-style-type: none"> Valid RecoverPoint management hostname or IP address.
<p>Parameter: RP_USER</p> <p>Section: [RP_CLUSTER_1]</p> <p>Specifies the RecoverPoint username that connects to the RecoverPoint cluster, which protects the XtremIO production volumes.</p> <p>Mandatory.</p> <p>NOTICE</p> <p>You must set the initial value of the parameter in the configuration file before the <code>ddbmadmin</code> command is used for any lockbox procedures. After this initial setting, you can modify the parameter value and rerun the <code>ddbmadmin</code> command as described in Configuring the lockbox on page 103.</p> <p>Valid values:</p> <ul style="list-style-type: none"> Undefined (default). Valid RecoverPoint username.

Configuring product operations over FC and IP networks

You can perform operations with the database application agent over either a Fibre Channel (FC) or Ethernet (IP) network connection between the database or application host and the Data Domain system.

Operations over an IP network are enabled by default.

To enable the database application agent software to use DD Boost over FC, ensure that you meet the following requirements:

- An FC connection is enabled according to [Enabling the DD Boost-over-FC service](#) on page 55.
- The configuration file contains the parameter settings described in the following table for the primary Data Domain system. [Setting up the configuration file](#) on page 78 provides details about setting parameters in the configuration file. [Configuring restores of replicated backups](#) on page 90 describes additional parameter settings for restores from a secondary Data Domain system over an FC network.
- The operating system user who performs the backup or restore has the correct device permissions as described in the following articles:
 - Fibre Channel Devices with Products using DD Boost in Linux/UNIX Environment* (Document ID dd95007)
 - Fibre Channel Devices with Products using DD Boost in Windows Environment* (Document ID dd95005)

Use the document ID to search for these articles on the Support website at <https://support.emc.com>.

- When a Data Domain host is rebooted, you might need to perform a rescan of the operating system devices on the application host to ensure that the DFC devices are recognized.

For each parameter, the following table lists the section heading of the configuration file section that contains the parameter.

Table 13 Parameters for operations over FC networks

<p>Parameter: DDBOOST_FC</p> <p>Section: [PRIMARY_SYSTEM]</p> <p>Specifies whether a backup or restore on the primary Data Domain system uses a Fibre Channel (FC) or IP network connection.</p> <hr/> <p>Note</p> <p>The primary Data Domain system must be configured to support an FC connection if this parameter is set to TRUE.</p> <hr/> <p>Valid values:</p> <ul style="list-style-type: none"> • FALSE (default) = Backup or restore uses an IP network connection. • TRUE = Backup or restore uses an FC network connection.
<p>Parameter: DEVICE_FC_SERVICE</p> <p>Section: [PRIMARY_SYSTEM]</p> <p>Specifies the name of the FC service configured on the primary Data Domain system to be used for a backup or restore.</p> <p>Mandatory when the <code>DDBOOST_FC</code> parameter is set to TRUE.</p> <hr/> <p>Note</p> <p>The <code>DEVICE_HOST</code> parameter must be set during FC operations. Common parameters on page 80 provides details on the parameter.</p> <hr/> <p>Valid values:</p> <ul style="list-style-type: none"> • Undefined (default). • Name of the DD Boost server on the primary Data Domain system, without a DFC- prefix.

Configuring the optimization of ProtectPoint backups for third-party multipathing software

When you use third-party multipathing software, such as PowerPath, on the database application agent host, ensure that the `/etc/lvm/lvm.conf` file on the host contains the recommended filter setting.

The filter setting in the `/etc/lvm/lvm.conf` file optimizes the ProtectPoint backups by improving the performance of LVM commands, such as the `lvs` command, during the backups.

For example, the following filter setting in the `lvm.conf` file prevents the listing of any duplicate physical volumes (PVs):

```
filter = [ "a|/dev/sda1.*|", "a|/dev/mapper/.*/.*|", "a|/dev/emcpower.*|",
"x|.*/.*|" ]
```

Without this filter setting in the file, the ProtectPoint backup is delayed as the LVM commands produce listings of the duplicate devices.

Configuring restores of replicated backups

The database application agent supports the restore of replicated backups from a secondary Data Domain system when the primary Data Domain system is unavailable or when you specifically want to restore from a replica.

To enable the database application agent to automatically restore replicated backups from a secondary Data Domain system, ensure that the configuration file contains the appropriate parameter settings from the following table. [Setting up the configuration file](#) on page 78 provides details about setting parameters in the configuration file.

Note

Unless specified otherwise, the parameters in the following table apply to the restores of replicated ProtectPoint VMAX backups and replicated ProtectPoint with RecoverPoint backups. The restore of a replicated ProtectPoint with RecoverPoint backup is a point-in-time restore and does not support a rollback restore. For a restore to an XtremIO system, the restore device group must be the same on both the primary Data Domain system and the secondary Data Domain system.

For restores of replicated ProtectPoint backups of data from a VMAX system, review the additional considerations in [Configuring restores of replicated ProtectPoint backups](#) on page 94.

For restores from a secondary system over a DD Boost-over-FC network, ensure that the operating system user who performs the restores has the correct device permissions as described [Configuring product operations over FC and IP networks](#) on page 88.

For each parameter, the following table lists the section heading of the configuration file section that contains the parameter.

Table 14 Parameters for restores from a secondary Data Domain system

Parameter: DDBOOST_USER

Section: [SECONDARY_SYSTEM]

Specifies the username of the DD Boost user configured on a secondary Data Domain system, when this parameter is set in the secondary system section of the configuration file. The secondary system section has the [SECONDARY_SYSTEM] heading.

Mandatory when `DEVICE_HOST` is set for a restore from the secondary Data Domain system.

Table 14 Parameters for restores from a secondary Data Domain system (continued)**NOTICE**

You must set the initial value of the parameter in the configuration file before the `ddbmadmin` command is used for any lockbox procedures. After this initial setting, you can modify the parameter value and rerun the `ddbmadmin` command as described in [Configuring the lockbox](#) on page 103.

Valid values:

- Undefined (default).
- Valid username of the DD Boost user on the secondary Data Domain system. For example:

```
DDBOOST_USER=user2
```

Parameter: DDVDISK_USER**Section: [SECONDARY_SYSTEM]**

Specifies the DD vdisk username that was specified during creation of the replication vdisk device pool on the secondary Data Domain system, when this parameter is set in the secondary system section of the configuration file. The secondary system section has the [SECONDARY_SYSTEM] heading.

Mandatory only when both of the following conditions are true:

- `DEVICE_HOST` is set for a restore from the secondary Data Domain system.
- The DD vdisk username is different from the DD Boost username specified in `DDBOOST_USER`.

NOTICE

You must set the initial value of the parameter in the configuration file before the `ddbmadmin` command is used for any lockbox procedures. After this initial setting, you can modify the parameter value and rerun the `ddbmadmin` command as described in [Configuring the lockbox](#) on page 103.

Valid values:

- DD Boost username specified in the `DDBOOST_USER` parameter setting (default).
- Valid DD vdisk username on the secondary Data Domain system.

Parameter: DEVICE_HOST**Section: [SECONDARY_SYSTEM]**

Specifies the hostname of a secondary Data Domain system, when this parameter is set in the secondary system section of the configuration file. The secondary system section has the [SECONDARY_SYSTEM] heading.

Mandatory for a restore from the secondary Data Domain system.

NOTICE

You must set the initial value of the parameter in the configuration file before the `ddbmadmin` command is used for any lockbox procedures. After this initial setting, you can modify the parameter value and rerun the `ddbmadmin` command as described in [Configuring the lockbox](#) on page 103.

Valid values:

- Undefined (default).
- Valid hostname as the fully qualified domain name of the secondary Data Domain system. For example:

Table 14 Parameters for restores from a secondary Data Domain system (continued)

<p>DEVICE_HOST=dd.repl.com</p>
<p>Parameter: DEVICE_PATH</p> <p>Section: [SECONDARY_SYSTEM]</p> <p>Specifies the name of the storage unit or a top-level directory within the storage unit on a secondary Data Domain system, when this parameter is set in the secondary system section of the configuration file. The secondary system section has the [SECONDARY_SYSTEM] heading.</p> <p>Mandatory for a restore from the secondary system when the DEVICE_HOST parameter is set in the secondary system section of the configuration file.</p> <p>NOTICE</p> <p>You must set the initial value of the parameter in the configuration file before the <code>ddbadmin</code> command is used for any lockbox procedures. After this initial setting, you can modify the parameter value and rerun the <code>ddbadmin</code> command as described in Configuring the lockbox on page 103.</p> <p>Valid values:</p> <ul style="list-style-type: none"> Undefined (default). Valid directory name of the storage unit on the secondary Data Domain system, without the <code>/data/col1</code> prefix. The specified name is case-sensitive. <p>For example, if the pathname on the secondary Data Domain system is <code>/data/col1/su2</code>, then the valid storage unit name for this parameter is <code>/su2</code>:</p> <p>DEVICE_PATH=/su2</p>
<p>Parameter: DEVICE_POOL</p> <p>Section: [SECONDARY_SYSTEM]</p> <p>Specifies the name of the MTree replication DD vdisk device pool on the secondary Data Domain system, when this parameter is set in the secondary system section of the configuration file. The secondary system section has the [SECONDARY_SYSTEM] heading.</p> <p>For a restore to a VMAX system, this name is not the name of the local pool that is created on the secondary Data Domain system.</p> <p>For a restore to an XtremIO system, this name is the name of the read-only pool that is created automatically by replication on the secondary Data Domain system.</p> <p>Mandatory only when both of the following conditions are true:</p> <ul style="list-style-type: none"> DEVICE_HOST is set for a restore from the secondary Data Domain system. The replication pool name is different from the source pool name. <p>NOTICE</p> <p>If this parameter is set in the configuration file, then you must register the username and password of the secondary vdisk pool with the lockbox. Configuring the lockbox on page 103 provides details.</p> <p>Valid values:</p> <ul style="list-style-type: none"> Undefined (default).

Table 14 Parameters for restores from a secondary Data Domain system (continued)

- Valid DD vdisk pool name on the secondary Data Domain system.

Parameter: DDBOOST_FC**Section: [SECONDARY_SYSTEM]**

Specifies whether a restore uses an FC or IP network connection from a secondary Data Domain system, where backups were replicated from the primary Data Domain system. This parameter is set in the secondary system section of the configuration file, which has the [SECONDARY_SYSTEM] section heading.

Note

The secondary Data Domain system must be configured to support an FC connection if this parameter is set to TRUE.

Valid values:

- FALSE (default) = Restore uses an IP network connection.
- TRUE = Restore uses an FC network connection.

Parameter: DEVICE_FC_SERVICE**Section: [SECONDARY_SYSTEM]**

Specifies the name of the FC service configured on a secondary Data Domain system, where backups are replicated from the primary Data Domain system. This parameter is set in the secondary system section of the configuration file, which has the [SECONDARY_SYSTEM] section heading.

Mandatory when the DDBOOST_FC parameter is set to TRUE for the secondary system.

Note

The DEVICE_HOST parameter must be set during FC operations.

Valid values:

- Undefined (default).
- Name of the DD Boost server on the secondary Data Domain system, without a DFC- prefix.

Parameter: RESTORE_DEVICE_POOL**Section: [SECONDARY_SYSTEM]**

This parameter is used for two types of restores:

- Restores of ProtectPoint for VMAX backups directly from the secondary Data Domain system to the restore host, without the use of a VMAX system.
- Restores of ProtectPoint with RecoverPoint backups from the secondary Data Domain system to the restore host.

Specifies the name of the DD vdisk device pool that provides the restore LUNs on the secondary Data Domain system, when this parameter is set in the secondary system section of the configuration file. The secondary system section has the [SECONDARY_SYSTEM] heading.

Mandatory only when both of the following conditions are true:

- DEVICE_HOST is set for a restore from the secondary Data Domain system.
- The restore LUNs are configured in a DD vdisk device pool that is different from the pool specified by RESTORE_DEVICE_POOL for the primary system.

Table 14 Parameters for restores from a secondary Data Domain system (continued)**NOTICE**

If this parameter is set in the configuration file, then you must register the username and password of the secondary vdisk pool with the lockbox. [Configuring the lockbox](#) on page 103 provides details.

Valid values:

- Undefined (default).
- Valid DD vdisk pool name on the secondary Data Domain system.

Parameter: VMAX_FASTX_RESTORE_SG**Section: [SECONDARY_SYSTEM]**

Specifies the name of the VMAX storage group to use during a ProtectPoint restore to a selected FAST.X restore device on VMAX. By default, the NsrSnapSG storage group is used for a ProtectPoint restore to a VMAX system.

Optional for a ProtectPoint restore to a secondary VMAX system.

Valid values:

- NsrSnapSG (default).
- Valid name of a VMAX storage group.

Configuring restores of replicated ProtectPoint backups

Additional considerations apply for the restore of replicated ProtectPoint backups.

The database application agent stores the device pool name, device group name, and static image name in the backup catalog during a ProtectPoint backup. The database application agent uses these values to locate the replicated backup image for restore.

The backups and catalog files reside in both of the following locations, and the Data Domain administrator must replicate both:

- DD Boost storage unit
- DD vdisk device pool

Due to replication lag, a restore of a replicated ProtectPoint backup from a secondary Data Domain system might fail in the following cases:

- Catalog entries are replicated, but static images are not fully replicated—When the database application agent tries to instantiate a static image in this case, the operation fails because the static images are not fully replicated. Data Domain MTree specifies that a partial static image cannot be used to instantiate a device.
- Static images are replicated, but catalog entries are not fully replicated—Because the catalog entries are not fully replicated, the restore fails or falls back to an old backup. The resulting behavior depends on the type of database.

The Data Domain administrator must provision the restore LUNs on the secondary Data Domain system. A local pool must be created for the restore LUNs according to the Data Domain documentation:

- To restore a ProtectPoint for VMAX backup and use VMAX FAST.X LUNs as the restore LUNs, the restore LUNs must be encapsulated and visible through the VMAX array that is connected to the restore host. This array can be a different array from where the application data originally resided.

Note

A rollback restore from a replica must be performed to the original source LUNs in the original VMAX system.

- To restore a ProtectPoint for VMAX backup directly from Data Domain to the restore host (without using a VMAX system), the restore LUNs are DD vdisk devices that are visible to the restore host.
- To restore a ProtectPoint with RecoverPoint backup, the restore LUNs are DD vdisk devices that are visible to the restore host.

You can set the `VMAX_FASTX_RESTORE_SG` parameter in the `[SECONDARY_SYSTEM]` section of the configuration file to specify the VMAX storage group to use during a restore of a replicated ProtectPoint backup to a selected restore device. The restore uses the `NsrSnapSG` storage group by default.

As a best practice, the original DD Boost storage unit name and DD vdisk device pool name should be used for the storage unit and device pool created for replication on the secondary Data Domain system.

If the Data Domain administrator uses different names for the storage unit and device pool on the secondary system, then the `[SECONDARY_SYSTEM]` section of the configuration file must include the following settings for a restore from the replica:

- `DEVICE_HOST`—Specifies the secondary Data Domain hostname.
 - `DDBOOST_USER`—Specifies the secondary DD Boost username.
-

Note

The password for this DD Boost username must be set in the lockbox according to the procedure in [Configuring the lockbox](#) on page 103.

- `DDVDISK_USER`—Specifies the secondary DD vdisk username.
-

Note

The password for this DD vdisk username must be set in the lockbox according to the procedure in [Configuring the lockbox](#) on page 103.

- `DEVICE_PATH`—Specifies the secondary DD Boost storage unit name, if different from the original name.
- `DEVICE_POOL`—Specifies the secondary DD vdisk device pool name, if different from the original.
- `RESTORE_DEVICE_POOL`—Specifies the DD vdisk device pool that provides the restore LUNs on the secondary Data Domain system.

For example, the following parameters are set in the `[SECONDARY_SYSTEM]` section of the configuration file and the lockbox updated accordingly:

```
DEVICE_HOST=dev_host.lss.emc.com
DDBOOST_USER=boost_rep
DDVDISK_USER=vdisk_rep
DEVICE_PATH=/IT_data_rep
DEVICE_POOL=IT_data_pool_rep
```

During a restore, the database application agent tries to connect to the primary Data Domain system first and then falls back to the secondary Data Domain system if the primary system connection fails.

Configuring ProtectPoint VMAX restores directly from Data Domain

You can optionally configure a restore of a ProtectPoint for VMAX backup directly from a Data Domain system to the restore host, which does not involve a VMAX system. The backup is restored to the restore host by using restore LUNs that are provisioned directly from Data Domain. This type of restore is a point-in-time restore, not a rollback restore.

The restore of a ProtectPoint for VMAX backup directly from a Data Domain system is especially useful in a local/remote configuration scenario. For example, you back up data from a VMAX system to a local (primary) Data Domain system and then replicate the backup to a remote (secondary) Data Domain system. You can perform a restore of the backup from the secondary Data Domain system, without using a VMAX system.

To configure the restore from a local (primary) Data Domain system, set the `RESTORE_DEVICE_POOL` and `RESTORE_DEVICE_GROUP` parameters in the `[PRIMARY_SYSTEM]` section of the configuration file. These parameters specify to use the restore LUNs in the DD vdisk device pool and device-group, instead of `FAST.X` restore LUNs on a VMAX system.

To configure the restore from a replicated (secondary) Data Domain system, set `RESTORE_DEVICE_POOL` in the `[SECONDARY_SYSTEM]` section of the configuration file, as required for the secondary Data Domain system. In this case, the restore device pool on the secondary Data Domain system can be different from the restore device pool on the primary system. However, the restore device group must be the same on both the primary and secondary Data Domain systems.

Configuring ProtectPoint VMAX restores from local snapshots

In addition to storing the backups on the Data Domain system, the database application agent keeps the last SnapVX snapshot on the VMAX system as its backup. To achieve a faster restore, you can restore from this last SnapVX snapshot, which is also known as a local snapshot.

The restore can be a point-in-time restore or a LUN-level rollback restore. You can perform the rollback restore to either the original source LUNs or alternate target LUNs in the same VMAX system.

The advantage of a restore from the local snapshot is that you can restore the data more quickly.

When you use the database application agent to perform a ProtectPoint backup on VMAX, the name of the local SnapVX snapshot on the VMAX system is `PROTECTPOINT_SNAP_<timestamp>`. The database application agent keeps the last SnapVX snapshot on the VMAX system.

When you use the database application agent to perform a restore, the agent restores from the local snapshot if it is in a valid state. Otherwise, the database application agent restores from the backup that is located on the Data Domain system.

To enable restores from the last SnapVX snapshot, you must follow the recommended data layout. For Oracle and SAP with Oracle, the recommended data layout uses

separate LUNs for the datafiles and archived redo logs. Also, DD Boost is used to back up the SAP BR*Tools.

Unless you follow the recommended data layout, the last SnapVX snapshot is not equivalent to the last backup that is created with the database utility.

Note

Do not manually create a snapshot that has the same name as a snapshot created by the database application agent.

To prepare for a point-in-time restore with the database application agent from the last local SnapVX snapshot on VMAX, you must configure the native VMAX LUN as the database application agent's restore LUN. By default, the database application agent selects the restore LUNs from the VMAX storage group NsrSnapSG unless the `VMAX_FASTX_RESTORE_SG` parameter is set to a different storage group name.

The VMAX storage group NsrSnapSG or the VMAX storage group specified by `VMAX_FASTX_RESTORE_SG` must contain both types of database application agent restore LUNs:

- Native VMAX LUNs
- FAST.X LUNs

To prepare for a redirected rollback restore from the local SnapVX snapshot on VMAX, ensure that you also meet the configuration requirements in the following topic.

The *ProtectPoint Version 4.0 Primary and Protection Storage Configuration Guide* provides details about how to perform the required configurations.

Configuring rollback restores of ProtectPoint backups

The database application agent does not support partitioned disks in a snapshot operation, such as a snapshot backup or rollback restore. In a rollback restore, on the target devices, any extra file systems and volume management that reside on partitioned disks and are not involved in the restore must be manually cleaned up before the restore. Otherwise, the rollback restore might fail.

You must set `RESTORE_TYPE_ORDER=rollback` in the configuration file to specify the rollback restore of a ProtectPoint backup. A rollback restore is a destructive restore because the rollback overwrites the entire contents of a snapshot unit, such as a volume group or disk or a RecoverPoint consistency group.

Rollback restores of ProtectPoint for VMAX backups

A rollback restore to a VMAX system is a LUN-level restore. You can perform a rollback restore of a ProtectPoint for VMAX backup to either the original source LUNs or alternate target LUNs on the same VMAX system. The rollback restore performs a restore of the whole volume group or the whole LUN when a volume manager is not used:

- You can perform a regular rollback restore to restore the backup to the original source LUNs on the backup host.
- You can perform a redirected rollback restore to relocate a database to an alternate host, configured on alternate target LUNs on the same VMAX array. In the redirected rollback restore, perform the restore of a full database backup only.

Note

The database application agent does not support a redirected rollback restore to alternate LUNs on the original backup host. During a rollback restore to the original backup host, the snapshot backup is restored to the original source LUNs. The backup file systems must exist on the source LUNs before the restore; re-create the file systems if required.

The database application agent 3.5 introduced the support of a redirected rollback restore of a ProtectPoint for VMAX backup in a DB2 pureScale environment. This operation restores to different target LUNs on the same VMAX system, and these LUNs are provisioned to an alternate DB2 pureScale environment. [DB2 pureScale requirements for ProtectPoint operations](#) on page 199 provides details.

The database application agent 4.0 introduced support of a redirected rollback restore of a ProtectPoint for VMAX backup on the same VMAX array for all the supported DB2, Oracle, and SAP Oracle systems.

For a regular rollback restore to the original source LUNs on the backup host, the file system with the same mount point as used in the backup must exist and be mounted on the host. If Logical Volume Manager (LVM) is used, then the volume group name must be the same.

For a redirected rollback restore to a different set of LUNs on an alternate host, ensure that you meet the following requirements:

- You have met the common requirements for a redirected restore to a different host. For example, the username, user ID (UID), group name, and group ID (GID) of the target database/instance owner match the original values captured during the backup.
- The file system with the same mount point as used in the backup must exist and be mounted on the destination host.
- The number of devices on which the file system resides on the destination host must be equal to the original number of devices in the backup.
- The size of the target LUN must be equal to or greater than the size of the original LUN.
- When multiple LUNs are included in the rollback restore, the destination LUN size must be greater than or equal to the static image size.
- For an Oracle rollback restore, the Oracle-Managed Files (OMF) feature is disabled for the Oracle database on the alternate host because the Oracle rollback restore in this release does not support the renaming of the restored files.
- If file system management is used, such as LVM or Veritas Volume Manager:
 - If a file system or volume manager exists on the backed-up devices, the file system or volume manager version on the recovery host might need to be the same as or higher than the version on the backed-up devices. The file system and volume manager documentation provides details.
 - The names of the volume group, logical volume, and physical device on the target devices do not need to match the original names, provided that no conflicts exist in the logical volume and volume group names.
 - The number of file systems and logical volumes of the target volume group do not need to match the numbers in the original volume group configuration.

Note

Any extra file systems, volume groups, and logical volumes on the recovery host must be listed in the `psrollback.res` file, so that these items are skipped during the safety checks. The `psrollback.res` file is described in the subsequent topic about safety checks during rollback restores.

- On the target devices, any volume group (and its logical volumes and file systems) or file systems that are not involved in the restore must be manually cleaned up before the restore. The extra file system must be unmounted, and the extra volume group must be removed. Otherwise, the rollback restore might fail.
- The database application agent software must be installed and configured properly on the host that performs the rollback restore. The `CLIENT` parameter must be set to the original value, as recorded in the backup.

Rollback restores of ProtectPoint with RecoverPoint backups that use RecoverPoint pre-5.0

With RecoverPoint pre-5.0, a rollback restore to an XtremIO system is a RecoverPoint consistency group-level restore, which restores all the LUNs in a consistency group. The rollback restore of a ProtectPoint with RecoverPoint backup is performed to the source XtremIO LUNs in the consistency group.

Note

With any version of RecoverPoint, a DB2 rollback restore to an XtremIO system is always a RecoverPoint consistency group-level restore.

With RecoverPoint pre-5.0, a ProtectPoint with RecoverPoint backup and rollback restore occur at the consistency group level, regardless of which objects are included in the backup command. As a best practice for the ProtectPoint with RecoverPoint rollback restore, when you perform the backup or rollback restore, do not exclude the logs or any database files that are part of the RecoverPoint consistency group being backed up or restored.

If any LUNs in the backed-up consistency group contain objects that were not included in the backup command, ensure that you manually unmount those LUNs before the rollback restore and then manually mount the LUNs back after the restore.

Rollback restores of ProtectPoint with RecoverPoint backups that use RecoverPoint 5.0 or later

With RecoverPoint 5.0 or later, a rollback restore of a ProtectPoint Oracle or SAP Oracle backup to an XtremIO system is a LUN-level restore, which can restore a partial RecoverPoint consistency group. The rollback restore performs a restore of the whole volume group or the whole LUN when a volume manager is not used. The whole volume group or whole LUN that is restored is known as the restore unit.

Note

A DB2 rollback restore to an XtremIO system is always a RecoverPoint consistency group-level restore.

With RecoverPoint 5.0 or later, the rollback restore of a ProtectPoint Oracle or SAP Oracle backup can restore the following objects:

- A database when the logs are part of the same consistency group but on a different restore unit.
- A pluggable database when the root and other pluggable databases and online logs are part of the same consistency group but on a different restore unit.

- Tablespaces when the rest of the database and the online logs are part of the same consistency group but on a different restore unit.

Safety checks during rollback restores of ProtectPoint backups

During a rollback restore, the database application agent performs safety checks by default. The safety checks ensure that there are no files, directories, partitions, or volumes (data targets) on the rollback target LUN other than those restored with ProtectPoint. If there are additional such data targets on the target LUN that are not included in the restore session, the database application agent fails the rollback restore as a safety precaution to prevent the overwriting of data.

For a rollback restore of a RecoverPoint consistency group with RecoverPoint pre-5.0, the safety checks also prevent the additional data targets on all the XtremIO LUNs of the target consistency group from being overwritten.

To override the safety checks, you can use the `psrollback.res` file. In the file, you must list all the files and directories to be excluded from the rollback safety checks.

For example, `lv011` is the logical volume at backup time, and `lv011` and `lv012` are logical volumes on the destination host. You must include `lv012` in the `psrollback.res` file to enable the rollback restore to proceed. You can also list the device name to ensure that the safety check skips all the file systems that reside on the device. To prevent `lv012` from being overwritten during the rollback restore, do not list `lv012` or the device name in the file.

NOTICE

Use the `psrollback.res` file with extreme caution to prevent possible data corruption. If you use this file to override the safety checks, the rollback restore might overwrite some database files that were not included in the restore session, such as Oracle online redo logs, which could result in data loss.

On Linux or Solaris SPARC, if a disk is configured with partitions, you can perform a rollback restore only if you list the entire disk in the `psrollback.res` file. The rollback restore then overwrites the entire disk. For example, if `/fs1` and `/fs2` are configured with partitions `/dev/sdc1` and `/dev/sdc2` respectively, then you must enable the rollback restore of `/fs1` by listing the entire disk `/dev/sdc` in `psrollback.res`. The rollback restore overwrites the entire disk `/dev/sdc`, so `/fs2` is also restored.

If a logical volume manager (LVM) controls the file system of an application host, then you must list in the `psrollback.res` file all the physical disks that belong to the LVM volume group. For example, if a volume group contains the disks `/dev/sdc` and `/dev/sdd`, and `/fs1` is the mount point of the file system, then the `psrollback.res` file must include the following lines:

```
/fs1/lost+found
/fs1/test
/dev/sdc
/dev/sdd
```

The `psrollback.res` file location is as follows:

- On UNIX systems: `/opt/dpsapps/dbappagent/config/psrollback.res`
- On Windows systems: `C:\Program Files\DPSAPPS\DBAPPAGENT\config\psrollback.res`

Refer to the later ProtectPoint chapters for any application-specific restrictions on rollback restore operations.

Example 1 Overriding safety checks during a rollback restore

If you are restoring `/fs1/data1.df` and `/fs1/data2.df` but there are other files in the `/fs1` directory, such as the files `lost+found` and `test`, you can exclude these other files from the safety checks during a rollback restore if you do not need these files. To exclude the files, list the file pathnames in the `psrollback.res` file:

```
more /opt/dpsapps/dbappagent/config/psrollback.res
```

```
/fs1/lost+found
/fs1/test
```

Configuring usage limits on Data Domain resources

You must complete the required procedures on the Data Domain host to configure the capacity or streams usage limits for the database application agent. The following topics provide the configuration details.

Refer to the configuration sections in subsequent chapters for additional guidelines and best practices related to the capacity or streams usage limits on application-specific systems.

Configuring usage quota on Data Domain capacity

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

Procedure

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

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

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

NOTICE

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

Configuring usage limits on Data Domain streams

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

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

Procedure

1. Determine which application agent hosts will use the storage unit.
2. Determine how many backup and restore streams to allow for the storage unit.
3. Create the storage unit. You can set the streams limit as part of the `ddboost storage-unit create` command or (after the storage unit is created) with the `ddboost storage-unit modify` command. The Data Domain documentation provides more details.

Note

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

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

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

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

NOTICE

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

Configuring the lockbox

A lockbox is an encrypted file that the database application agent uses to store and protect confidential information from unauthorized access. The lockbox stores the Data Domain system information, including credentials for the DD Boost user.

Starting with release 3.5, the database application agent, Microsoft application agent, and ProtectPoint file system agent use the same lockbox in the common lockbox location.

Lockbox requirements

The common lockbox file used by the database application agent is named `agents.clb`. The default directory location of the lockbox file is as follows:

- On UNIX or Linux: `/opt/dpsapps/common/lockbox`
- On Windows: `C:\Program Files\DPSAPPS\common\lockbox`

Note

When the database application agent is updated from a previous release, the pre-3.5 lockbox file is moved to this common lockbox location and renamed if another agent has not yet created the common lockbox. The pre-3.5 lockbox is retained in the following directory only if another agent has already created the common lockbox:

- On UNIX or Linux: `/var/opt/ddbda/lockbox`
- On Windows: `C:\Program Files\EMC DD Boost\DA\config\lockbox`

The root or administrative user can specify a nondefault directory for the lockbox file during the lockbox creation. For example, the lockbox can be stored in a shared directory in a cluster environment.

In the same directory as the lockbox file, the product also maintains additional files needed for correct lockbox operations.

NOTICE

The root or administrative user must have read and write permissions to the lockbox, and all the database users must have at least the read permission to the lockbox.

On UNIX, the root user can assign the lockbox group ownership to a different group, such as a DBA group, which enables the group users to perform specific lockbox operations. This feature of assigning the lockbox group ownership is not supported on Windows.

The default file permissions on a lockbox file can be changed to restrict the lockbox access to a specified group of users. The lockbox files in the directory must all have the same permissions.

When the database agent is installed on the same application host as the ProtectPoint file system agent, a separate lockbox location is required for each agent.

Configuring the lockbox with the ddbmadmin command

On Windows, the administrative user must run the `ddbmadmin` command to perform all the lockbox operations. The lockbox group ownership cannot be changed on Windows.

On UNIX, the root user can run the `ddbmadmin` command to perform all the lockbox operations. If the UNIX root user assigns the lockbox group ownership to a DBA group, the group users can also run `ddbmadmin` to perform the following lockbox operations:

- Register a Data Domain system with the lockbox.
- Unregister a Data Domain system.
- Update the lockbox configuration.

On UNIX, only the root user can perform the following lockbox operations:

- Create the lockbox on the database host to be used for backups or restores.
- Change the lockbox group ownership.
- Grant lockbox access to a specific host.
- Revoke lockbox access from a specific host.

Note

Only one host at a time can access a shared lockbox with the `ddbmadmin` command.

The following `ddbmadmin` commands perform the lockbox operations:

```
ddbmadmin -L [-a LOCKBOX_PATH=<lockbox_dir_pathname>] [-a LOCKBOX_OWNER_GID=
<group_ID_of_lockbox_owner>] [-D 9]

ddbmadmin -P -z <configuration_file> [-D 9]

ddbmadmin -X -z <configuration_file> [-a CONFIRM={yes|no}] [-D 9]

ddbmadmin -G [-a LOCKBOX_PATH=<lockbox_dir_pathname>] [-a
LOCKBOX_REMOTE_HOST=<hostname_to_add>] [-a VIRTUAL_HOST={yes|no}] [-D 9]

ddbmadmin -R [-a LOCKBOX_PATH=<lockbox_dir_pathname>] [-a
LOCKBOX_REMOTE_HOST=<hostname_to_delete>] [-D 9]

ddbmadmin -U [-a LOCKBOX_PATH=<lockbox_dir_pathname>] [-D 9]

ddbmadmin -U -a LOCKBOX_IMPORT=TRUE -a LOCKBOX_PATH=<nondefault_lockbox_directory>
```

In these `ddbmadmin` commands, *<configuration_file>* is the complete pathname of the configuration file, such as `/opt/dpsapps/dbappagent/config/ddbda.cfg`.

NOTICE

You must use a complete pathname in each case when the `ddbmadmin` command requires a file name.

The following table describes the `ddbmadmin` command options for lockbox operations.

Table 15 Options of the `ddbmadmin` command for lockbox operations

Option	Description
<code>-a</code>	This option enables the <code>ddbmadmin</code> command to run in noninteractive mode. Running the ddbmadmin command in noninteractive mode on page 105 provides details.
<code>-L</code>	This option creates the lockbox and changes the UNIX group ownership if requested. Creating the lockbox and changing the UNIX group ownership on page 107 provides details.
<code>-P -z</code> <code><configuration_file></code>	This option registers a primary and optional secondary Data Domain system with the lockbox. Adding Data Domain systems to the lockbox on page 108 provides details.
<code>-X -z</code> <code><configuration_file></code>	This option unregisters a primary and optional secondary Data Domain system. Removing Data Domain systems from the lockbox on page 113 provides details.
<code>-G</code>	This option grants lockbox access to a specific host. Configuring the lockbox in a high-availability environment on page 115 provides details.
<code>-R</code>	This option revokes lockbox access from a specific host. Configuring the lockbox in a high-availability environment on page 115 provides details.
<code>-U</code>	This option updates the lockbox configuration. Updating the lockbox on page 114 provides details.
<code>-D 9</code>	This option generates debugging information during a lockbox operation. The option is used to troubleshoot lockbox issues.

Running the `ddbmadmin` command in noninteractive mode

You can run the `ddbmadmin` command in a noninteractive mode by specifying the `-a` option with any of the other options except the `-P` option. The `ddbmadmin -P` command does not support the noninteractive mode because you must manually specify a password to register a Data Domain system with the lockbox.

You can run the `ddbmadmin` command in an interactive mode by omitting the `-a` option. In the interactive mode, the command prompts for any required inputs as described in the following topics.

The following table provides examples of the noninteractive `ddbmadmin` commands.

Table 16 Examples of noninteractive `ddbmadmin` commands

Creating the lockbox and changing the UNIX group ownership if requested:

```
ddbmadmin -L -a LOCKBOX_PATH=<lockbox_dir_pathname> -a
LOCKBOX_OWNER_GID=<group_ID_of_lockbox_owner>
```

Table 16 Examples of noninteractive `ddbmadm` commands (continued)

<p>Example:</p> <pre>ddbmadm -L -a LOCKBOX_PATH=/opt/lockbox -a LOCKBOX_OWNER_GID=501</pre> <p>Output:</p> <pre>Lockbox has been successfully created in the directory '/opt/lockbox' with group ownership 501.</pre>
<p>Unregistering a primary and optional secondary Data Domain system:</p> <pre>ddbmadm -X -z <configuration_file> -a CONFIRM={yes no}</pre> <p>CONFIRM=yes confirms that the unregistration must continue.</p> <p>Example:</p> <pre>ddbmadm -X -z /opt/dpsapps/dbappagent/config/ddbda.cfg -a CONFIRM=yes</pre> <p>Output:</p> <pre>Lockbox directory is '/tmp/lb'. Device host 'magni' for DD Boost user 'dduser1' has been unregistered from the lockbox.</pre>
<p>Granting lockbox access to a specific host:</p> <pre>ddbmadm -G -a LOCKBOX_PATH=<lockbox_dir_pathname> -a LOCKBOX_REMOTE_HOST=<hostname_to_add> -a VIRTUAL_HOST={yes no}</pre> <p>VIRTUAL_HOST=no indicates that the host is not a virtual host.</p> <p>Example:</p> <pre>ddbmadm -G -a LOCKBOX_PATH=C:\lockbox -a LOCKBOX_REMOTE_HOST=host2.xyz.com -a VIRTUAL_HOST=no</pre> <p>Output:</p> <pre>Host 'host2.xyz.com' has been granted access to the lockbox in the directory 'C:\lockbox'. Ensure that the administrator on host 'host2.xyz.com' runs the ddbmadm -U command to enable backup and recovery operations on host 'host2.xyz.com'.</pre>
<p>Revoking lockbox access from a specific host:</p> <pre>ddbmadm -R -a LOCKBOX_PATH=<lockbox_dir_pathname> -a LOCKBOX_REMOTE_HOST=<hostname_to_delete></pre>

Table 16 Examples of noninteractive `ddbadmin` commands (continued)

Example:

```
ddbadmin -R -a LOCKBOX_PATH=C:\lockbox -a LOCKBOX_REMOTE_HOST=host2.xyz.com
```

Output:

```
Revoked access from the host 'host2.xyz.com' to the lockbox in the directory 'C:\lockbox'.
```

Updating the lockbox configuration:

```
ddbadmin -U -a LOCKBOX_PATH=<lockbox_dir_pathname>
```

Example:

```
ddbadmin -U -a LOCKBOX_PATH=C:\lockbox
```

Output:

```
Lockbox 'agents.clb' in the directory 'C:\lockbox' has been updated.
```

Creating the lockbox and changing the UNIX group ownership

Only the root or administrative user can run the `ddbadmin -L` command. This command creates the lockbox if it does not exist, and updates the UNIX group ownership of the lockbox if requested. If a lockbox already exists, this command only updates the UNIX group ownership if requested.

On UNIX, this command sets the directory permissions to 775 for the new or existing lockbox, and sets the lockbox file permissions to 664.

The `ddbadmin -L` command prompts for a lockbox directory pathname. If no pathname is specified, the lockbox is created in the default directory. On UNIX, the command also prompts for a group ID for the lockbox owner. If 0 or no ID is specified, the root user group maintains the lockbox ownership.

The following UNIX example creates the lockbox files in the `/opt/lockbox` directory, and sets the group owner to `dba_grp1`, which has the group ID 501:

```
# ddbadmin -L
```

```
Provide the full path for the lockbox, or press Enter to accept the
default directory (<default lockbox_directory>): /opt/lockbox
Provide a group ID for lockbox ownership, or type 0 to accept the
'root user' group as the lockbox owner: 501
Lockbox has been successfully created in the directory '/opt/
lockbox' with group ownership 501.
```

```
# ls -l (in the /opt/lockbox directory)
```

```
-rw-rw-r-- 1 root root 3.6K Nov 16 17:38 agents.clb
-rw-rw-r-- 1 root root 6 Nov 16 17:38 agents.clb.FCD
```

```
-rw-rw-r-- 1 root root 3.6K Nov 16 17:38 agents.clb.bak
-rw-rw-r-- 1 root root 5 Nov 16 17:38 agents.clb.bak.FCD
```

Adding Data Domain systems to the lockbox

To enable backups to and restores from a Data Domain system, you must run the `ddbmadmin -P -z <configuration_file>` command to register the Data Domain system to the host. This command creates the required lockbox (if it does not exist) or updates an existing lockbox.

Note

The lockbox is created only if the root or administrative user runs the command. On UNIX, if the root user assigns the lockbox group ownership to a DBA group, the group users can run the command to register the Data Domain systems.

The following topics provide details about adding Data Domain systems to the lockbox for DD Boost operations and for ProtectPoint operations with VMAX and XtremIO systems.

Adding Data Domain systems to the lockbox for DD Boost operations

Before you run the `ddbmadmin -P -z <configuration_file>` command, the configuration file must exist and contain the mandatory parameter settings. For example, the following parameters are set in the [PRIMARY_SYSTEM] section of the configuration file.

```
DDBOOST_USER=dduser1
DEVICE_HOST=magni
DEVICE_PATH=/hermes-ddboost
```

You can optionally register a secondary Data Domain system to be used for restores when the primary system is unavailable. In this case, the [SECONDARY_SYSTEM] section of the configuration file must also contain the parameter settings for the secondary system. For example:

```
DDBOOST_USER=dduser2
DEVICE_HOST=telly
DEVICE_PATH=/windows-poseidon-boost
```

To register the secondary system, the primary system parameters must be in the configuration file.

Optionally, you can set `LOCKBOX_PATH` in the configuration file to a nondefault lockbox location. [Setting up the configuration file](#) on page 78 provides details about setting parameters in the configuration file.

NOTICE

If you edit the configuration file and modify any of these parameter settings after you registered them with the lockbox, you must rerun the `ddbmadmin -P -z <configuration_file>` command to update the lockbox entries. Otherwise, backups and restores can fail.

Based on the configuration file settings, the `ddbmadmin -P -z <configuration_file>` command prompts for the required passwords for the primary and secondary systems. The command verifies the passwords by logging in to the systems, encrypts the passwords, and stores the encrypted passwords in the lockbox.

For example, the configuration file contains the following parameter settings:

```
[GENERAL]
LOCKBOX_PATH=/tmp/lb

[PRIMARY_SYSTEM]
DDBOOST_USER=dduser1
DEVICE_HOST=magni
DEVICE_PATH=/hermes-ddboost

[SECONDARY_SYSTEM]
DDBOOST_USER=dduser2
DEVICE_HOST=telly
DEVICE_PATH=/windows-poseidon-boost
```

The `ddbmadmin -P -z <configuration_file>` command displays the following prompts and information:

```
# ddbmadmin -P -z /opt/dpsapps/dbappagent/config/ddbda.cfg

Performing the registration of the device host 'magni' for DD Boost
user 'dduser1'.
Enter password:
Confirm password:
Logging in to the device host 'magni' with DD Boost credentials.
Logging in to the device host 'magni' with DD Boost credentials was
successful.

Continue with the registration of the secondary device host 'telly'
for DD Boost user 'dduser2'? [y/n]: y
Performing the registration of the device host 'telly' for DD Boost
user 'dduser2'.
Enter password:
Confirm password:
Logging in to the device host 'telly' with DD Boost credentials.
Logging in to the device host 'telly' with DD Boost credentials was
successful.

Lockbox directory is '/tmp/lb'.
Device host 'magni' for DD Boost user 'dduser1' has been registered
in the lockbox.
Device host 'telly' for DD Boost user 'dduser2' has been registered
in the lockbox.
```

Adding Data Domain systems to the lockbox for ProtectPoint operations with VMAX

In addition to the DD Boost parameters described in the preceding topic, the configuration file for ProtectPoint operations with a VMAX system also requires the following parameters:

- In `[PRIMARY_SYSTEM]` section of the configuration file: `DDVDISK_USER` and `DEVICE_POOL` (`DDVDISK_USER` is required only if it is different than `DDBOOST_USER`)
- In `[SECONDARY_SYSTEM]` section of the configuration file: `DDVDISK_USER` and `DEVICE_POOL` (`DDVDISK_USER` is required only if it is different than `DDBOOST_USER`)

If the DD Boost and DD vdisk usernames are the same for a Data Domain system, then the `ddbadmin -P -z <configuration_file>` command automatically sets the DD vdisk password for that system to the DD Boost password.

If the device pool value is set, then the command tries to log in to the Data Domain system by using the DD vdisk credentials before storing the information in the lockbox.

The `ddbadmin -P -z` command creates a separate lockbox entry for each of the following cases:

- DDVDISK_USER is set but DEVICE_POOL is not set in the [PRIMARY_SYSTEM] section of the configuration file
- DDVDISK_USER and DEVICE_POOL are both set in the [PRIMARY_SYSTEM] section of the configuration file

NOTICE

Ensure that you run the `ddbadmin -P -z` command for the parameter setting combination that will be used during operations with the database application agent.

If you edit the configuration file and modify any of these parameter settings after you registered them with the lockbox, you must rerun the `ddbadmin -P -z <configuration_file>` command to update the lockbox entries. Otherwise, backups and restores can fail.

For example, the configuration file for ProtectPoint operations with a VMAX system contains the following parameter settings:

```
[GENERAL]
LOCKBOX_PATH=/tmp/lb

[PRIMARY_SYSTEM]
DDBOOST_USER=dduser1
DEVICE_HOST=magni
DEVICE_PATH=/hermes-ddboost
DDVDISK_USER=dduser1
DEVICE_POOL=pool

[SECONDARY_SYSTEM]
DDBOOST_USER=dduser2
DEVICE_HOST=telly
DEVICE_PATH=/windows-poseidon-boost
DDVDISK_USER=ddvdiskuser2
DEVICE_POOL=pool2
```

In this case, the `ddbadmin -P -z` command does not prompt for the DD vdisk password for the primary system because the DD vdisk and DD Boost usernames are the same:

```
# ddbadmin -P -z /opt/dpsapps/dbappagent/config/ddbda.cfg

Performing the registration of the device host 'magni' for DD Boost
user 'dduser1'.
Enter password:
Confirm password:
Logging in to the device host 'magni' with DD Boost credentials.
Logging in to the device host 'magni' with DD Boost credentials was
successful.
```

```

Continue with the registration of the secondary device host 'telly'
for DD Boost user 'dduser2'? [y/n]: y
Performing the registration of the device host 'telly' for DD Boost
user 'dduser2'.
Enter password:
Confirm password:
Logging in to the device host 'telly' with DD Boost credentials.
Logging in to the device host 'telly' with DD Boost credentials was
successful.

Performing the registration of the device host 'magni' for DD VDISK
user 'dduser1'.
Using the credentials from the DD Boost user registration for
device host 'magni' because the DD VDISK username is the same as
the DD Boost username.
Logging in to the device host 'magni' with DD VDISK credentials.
Logging in to the device host 'magni' with DD VDISK credentials was
successful.

Continue with the registration of the secondary device host 'telly'
for DD VDISK user 'ddvdiskuser2'? [y/n]: y
Performing the registration of the device host 'telly' for DD VDISK
user 'ddvdiskuser2'.
Enter password:
Confirm password:
Logging in to the device host 'telly' with DD VDISK credentials.
Logging in to the device host 'telly' with DD VDISK credentials was
successful.

Lockbox directory is '/tmp/lb'.
Device host 'magni' for DD Boost user 'dduser1' has been registered
in the lockbox.
Device host 'telly' for DD Boost user 'dduser2' has been registered
in the lockbox.

Device host 'magni' for DD VDISK user 'dduser1' has been registered
in the lockbox.
Device host 'telly' for DD VDISK user 'ddvdiskuser2' has been
registered in the lockbox.

```

Adding Data Domain systems and RPA to the lockbox for ProtectPoint with RecoverPoint operations

In addition to the DD Boost parameters described in a preceding topic, the configuration file for ProtectPoint with RecoverPoint operations with an XtremIO system also requires the following parameters. These parameters which must be set in the [RP_CLUSTER_1] section of the configuration file:

- RP_MGMT_HOST
- RP_USER

You must run the `ddbadmin -P -z` command to create a lockbox entry for the RecoverPoint user. The command prompts for the user password to store in the lockbox.

Note

You can register more than one user for the same RecoverPoint management host in the same lockbox by using different configuration files. For example, different applications might use different RecoverPoint clusters and each cluster might use the same RecoverPoint management host but a different user.

When you run the `ddbadmin -P -z` command, the command tries to log in to the RecoverPoint appliance to validate the RecoverPoint user credentials before storing the information in the lockbox.

NOTICE

Ensure that you run the `ddbadmin -P -z` command for the parameter setting combination that will be used during operations with the database application agent. If you edit the configuration file and modify any of these parameter settings after you registered them with the lockbox, you must rerun the `ddbadmin -P -z <configuration_file>` command to update the lockbox entries. Otherwise, backups and restores can fail.

For example, the configuration file for ProtectPoint with RecoverPoint operations with an XtremIO system contains the following parameter settings:

```
[GENERAL]
LOCKBOX_PATH=/tmp/lb

[PRIMARY_SYSTEM]
DDBOOST_USER=dduser1
DEVICE_HOST=magni
DEVICE_PATH=/hermes-ddboost

[RP_CLUSTER_1]
RP_USER=rpadmin
RP_MGMT_HOST=ledmrp08.lss.emc.com
```

The `ddbadmin -P -z` command reads the `RP_USER` parameter for the specified RecoverPoint management host from the configuration file and prompts for the user password:

```
# ddbadmin -P -z /opt/dpsapps/dbappagent/config/ddbda.cfg

Performing the registration of the device host 'magni' for DD Boost
user 'dduser1'.
Enter password:
Confirm password:
Logging in to the device host 'magni' with DD Boost credentials.
Logging in to the device host 'magni' with DD Boost credentials was
successful.

Lockbox directory is '/tmp/lb'.
Device host 'magni' for DD Boost user 'dduser1' has been registered
in the lockbox.

Performing the registration of the RecoverPoint management host
'ledmrp08.lss.emc.com' for RecoverPoint user 'rpadmin'.
Enter password:
Confirm password:
Logging in to the RecoverPoint management host
'ledmrp08.lss.emc.com'...
Logging in to the RecoverPoint management host '
ledmrp08.lss.emc.com' was successful.

The configuration file contains the following for the above
registration command:
RP_USER=rpadmin
RP_MGMT_HOST=ledmrp08.lss.emc.com
```


Removing Data Domain systems from the lockbox

You can run the `ddbmadmin -X -z <configuration_file>` command to unregister Data Domain systems. The command uses the configuration file settings to determine the required information about the Data Domain systems, and prompts for consent to unregister each system.

Note

On UNIX, if the root user assigned the lockbox group ownership to a DBA group, the group users can run the command to unregister Data Domain systems.

During unregistration, the `ddbmadmin -X -z` command deletes the information in the lockbox about the registered systems and then deletes the corresponding parameters in the configuration file. The command also reads and deletes the following ProtectPoint parameters if set in the configuration file:

- In [PRIMARY_SYSTEM] section of the configuration file: `DDVDISK_USER` and `DEVICE_POOL` for the primary Data Domain system.
- In [SECONDARY_SYSTEM] section of the configuration file: `DDVDISK_USER` and `DEVICE_POOL` for the secondary Data Domain system.
- In [RP_CLUSTER_1] section of the configuration file: `RP_USER` for a RecoverPoint cluster and XtremIO system.

The `DEVICE_HOST` parameters in the [PRIMARY_SYSTEM] and [SECONDARY_SYSTEM] sections of the configuration file are not deleted during the unregistration process.

A primary system can be unregistered only after the secondary system has been successfully unregistered. This requirement applies to both the DD Boost and DD vdisk users.

In the following example, the configuration file settings prior to the unregistration are as follows:

```
[GENERAL]
LOCKBOX_PATH=/tmp/lb

[PRIMARY_SYSTEM]
DDBOOST_USER=dduser1
DEVICE_HOST=magni
DEVICE_PATH=/hermes-ddboost
DDVDISK_USER=ddvdiskuser1
DEVICE_POOL=pool

[SECONDARY_SYSTEM]
DDBOOST_USER=dduser2
DEVICE_HOST=telly
DEVICE_PATH=/windows-poseidon-boost
DDVDISK_USER=ddvdiskuser2
DEVICE_POOL=pool2
```

The following command performs the unregistration. The primary system is not unregistered for the DD vdisk user because the consent is not provided:

```
# ddbmadmin -X -z /opt/dpsapps/dbappagent/config/ddbda.cfg

The lockbox directory to be used is '/tmp/lb'.
```

```
The device pathname and DD Boost username will be deleted from
'/opt/dpsapps/dbappagent/config/ddbda.cfg' after unregistration.
Continue with the unregistration of the secondary device host
'telly' for DD Boost user 'dduser2'? [y/n]: y
The lockbox directory to be used is '/tmp/lb'.
```

```
The device pathname and DD Boost username will be deleted from
'/opt/dpsapps/dbappagent/config/ddbda.cfg' after unregistration.
Continue with the unregistration of the device host 'magni' for DD
Boost user 'dduser1'? [y/n]: y
The lockbox directory to be used is '/tmp/lb'.
```

```
The device pool and DD VDISK username will be deleted from '/opt/
dpsapps/dbappagent/config/ddbda.cfg' after unregistration.
Continue with the unregistration of the secondary device host
'telly' for DD VDISK user 'ddvdiskuser2'? [y/n]: y
The lockbox directory to be used is '/tmp/lb'.
```

```
The device pool and DD VDISK username will be deleted from '/opt/
dpsapps/dbappagent/config/ddbda.cfg' after unregistration.
Continue with the unregistration of the device host 'magni' for DD
VDISK user 'ddvdiskuser1'? [y/n]: n
Confirmation for the unregistration of the device host 'magni' was
not provided.
Unregistration of the device host 'magni' for DD VDISK user
'ddvdiskuser1' will not be performed.
```

```
Lockbox directory is '/tmp/lb'.
Device host 'magni' for DD Boost user 'dduser1' has been
unregistered from the lockbox.
Device host 'telly' for DD Boost user 'dduser2' has been
unregistered from the lockbox.
Device host 'telly' for DD VDISK user 'ddvdiskuser2' has been
unregistered in the lockbox.
```

The unregistration removes some of the parameter settings in the configuration file. The configuration file settings after the unregistration are as follows:

```
[GENERAL]
LOCKBOX_PATH=/tmp/lb

[PRIMARY_SYSTEM]
DEVICE_HOST=magni
DDVDISK_USER=ddvdiskuser1
DEVICE_POOL=pool

[SECONDARY_SYSTEM]
DEVICE_HOST=telly
```

Updating the lockbox

You can run the `ddbmadmin -U` command to update the lockbox configuration. This operation ensures that the lockbox is continuously accessible to the host. For example, you can run the command to update the lockbox before a backup or restore operation.

Note

On UNIX, if the root user assigned the lockbox group ownership to a DBA group, the group users can run the command to update the lockbox.

NOTICE

To ensure that the lockbox is continuously accessible, ensure that the `ddbmadmin -U` command is run every time that a system is restarted. If you do not run this command after every system restart, then the lockbox might become inaccessible after a major system update. [Major system update can produce an error about lockbox stable value threshold](#) on page 136 provides details.

You must run the `ddbmadmin -U` command with the `-a LOCKBOX_IMPORT=TRUE` option if both the following conditions are true:

- You have updated from a pre-4.0 version of the database application agent on a Solaris x64 system.
- The pre-4.0 lockbox is stored in a nondefault directory, which is any directory other than `/opt/dpsapps/common/lockbox` or `/var/opt/ddbda/lockbox`.

When you use the lockbox import option, you must also use the `-a LOCKBOX_PATH=<nondefault_lockbox_directory>` option to specify the nondefault pathname of the pre-4.0 lockbox:

```
ddbmadmin -U -a LOCKBOX_IMPORT=TRUE -a
LOCKBOX_PATH=<nondefault_lockbox_directory>
```

Note

You only need to run this command once after you have performed the software update of the database application agent.

Configuring the lockbox in a high-availability environment

In a high-availability environment, you can create a lockbox on the local disk of each node. In this case, each lockbox must contain the same information. Alternately, you can configure a lockbox in a shared location and grant the lockbox access to all the nodes in the environment.

Note

When the lockbox is located in an NFS-shared location, the NFS share must grant access to the root or administrative user. For example, the NFS share is exported with the `no root squash` option.

You must complete the following steps to configure a lockbox in a shared location.

Procedure

1. To register a Data Domain system, select one host (node), and then run the `ddbmadmin -P -z <configuration_file>` command.
2. To grant lockbox access to another host, such as `host2`, run the `ddbmadmin -G` command on the host from step 1.

Note

You must run the `ddbmadmin -G` command separately for each host (node) when the environment contains multiple hosts.

The `ddbmadmin -G` command prompts for the hostname of the host to be granted the lockbox access. Provide the required hostname:

- If **host2** is a UNIX or Linux system, provide the output of the `uname -n` command on **host2**.
- If **host2** is a Windows system:
 - Provide the fully qualified domain name.
 - In a Microsoft cluster for Oracle, also provide the Oracle cluster service hostname.

The following examples show the `ddbmadmin -G` command on the different platforms:

- On UNIX or Linux:

```
# ddbmadmin -G

Provide full pathname for the lockbox, or press Enter to
accept the default directory (/opt/dpsapps/common/
lockbox): /opt/lockbox
Hostname to grant access to lockbox: host2.xyz.com
Host 'host2.xyz.com' has been granted access to the lockbox
in the directory '/opt/lockbox'.
Ensure that the root user on host 'host2.xyz.com' runs the
ddbmadmin -U command to enable backup and recovery
operations on host 'host2.xyz.com'.
```

- On Windows:

```
C:\Program Files\DPSAPPS\DBAPPAGENT\bin> ddbmadmin -G

Provide full pathname for the lockbox, or press Enter to
accept the default directory (C:\Program Files\DPSAPPS
\common\lockbox): C:\lockbox
Hostname to grant access to lockbox: host2.xyz.com
Is 'host2.xyz.com' a virtual hostname that is part of a
cluster? [y/n]: n
Host 'host2.xyz.com' has been granted access to the lockbox
in the directory 'C:\lockbox'.
Ensure that the root user on host 'host2.xyz.com' runs the
ddbmadmin -U command to enable backup and recovery
operations on host 'host2.xyz.com'.
```

In a Microsoft cluster for Oracle setup, you must also grant access to the Oracle cluster service hostname and confirm that it is a virtual hostname, as shown in the following example:

```
C:\Program Files\DPSAPPS\DBAPPAGENT\bin> ddbmadmin -G

Provide full pathname for the lockbox, or press Enter to
accept the default directory (C:\Program Files\DPSAPPS
\common\lockbox): C:\lockbox
Hostname to grant access to lockbox: oraclecluster.xyz.com
Is 'oraclecluster.xyz.com' a virtual hostname that is part
of a cluster? [y/n]: y
Host 'oraclecluster.xyz.com' has been granted access to the
lockbox in the directory 'C:\lockbox'.
Ensure that the root user on host 'host2.xyz.com' runs the
ddbmadmin -U command to enable backup and recovery
operations on host 'oraclecluster.xyz.com'.
```

3. For each additional host that needs to be added to the lockbox, repeat step 2. For each host to which you grant lockbox access, run the `ddbmadmin -G` command separately.

NOTICE

Keep a record of the hosts that are granted access to the lockbox. You cannot use the `ddbmadmin` command to obtain a list of all the hosts that have lockbox access.

In an environment where the lockbox is shared among multiple hosts, the user from each host that performs any operations with the database application agent must have operating system read access to the lockbox files. The root or administrative user on each host must have read and write access to the files.

You can run the `ddbmadmin -R` command to revoke the lockbox access from a host. You must run the command on a host other than the host from which you revoke the lockbox access. For example, you cannot run the command on `host1` to revoke the lockbox access from `host1`.

When a new node is added to a high-availability environment or replaces an existing node, grant the lockbox access to the new node according to step 2.

Configuring the display and deletion of save set information

When you perform a large number of backups with the database application agent, the available space on the Data Domain system can become greatly reduced. To delete old backups and free up space on the system, use the database native backup management tools if available.

Some database tools, such as the SAP Oracle tools, do not provide the backup management functionality to delete backups. In this case, you can use the `ddbmadmin` program for space management on the Data Domain system.

Note

The information in this section applies to all the supported DD Boost backups, ProtectPoint backups, and ProtectPoint with RecoverPoint backups. The `ddbmadmin` program deletes the backup information on the Data Domain system but does not delete any information on the database server.

For DB2 and Oracle backup deletions, use the database native backup management tools:

- For DB2 backup deletion:
 - For a ProtectPoint backup, use the `db2acsutil` command with the `delete` option.
 - For a DD Boost backup, follow the DB2 automatic recovery object deletion policy.
- For Oracle backup deletion, use the `RMAN delete` command.

A backup with the database application agent consists of backup save sets, where a save set is a collection of one or more save files created during the backup session. A save file is an operating system file or block of data, the simplest object that you can back up or restore. A backup creates one or more save files within a save set. The `ddbmadmin` program can perform deletions at the save set level only, deleting all the save files in a save set.

You can use the `ddbmadmin` program to perform any of the following operations:

- Display all the clients for a specified device path on the Data Domain system.
- Display information about the backup save sets.
- Display information about the save files.
- Delete the save sets created during a specified time interval.
- Upgrade the SAP Oracle backup index from the database application agent 1.0 to release 4.0.

You must run the `ddbmadmin` command at the command line with the required options. Certain command options are mandatory for different operations. The `-z <configuration_file>` option is mandatory for all operations, and specifies the configuration file used with the database application agent.

Note

A separate configuration file is required for each Data Domain system.

All the information that the `ddbmadmin` command prints to standard output is added to the operational log file, `ddbmadmin.messages.log`.

The following topics provide details on how to use the `ddbmadmin` command and options for the supported operations.

Using the `ddbmadmin` command to display and delete save sets

A DBA user can run the `ddbmadmin` command with the appropriate options to perform the following operations:

- Display all the clients for a specified device path on the Data Domain system:

```
ddbmadmin -i -z <configuration_file> [-D 9]
```

[Using the `ddbmadmin` command to display clients for a device path](#) on page 121 provides details.

- Display information for the save sets within a save time range:

```
ddbmadmin -s [-t] [-b <start_time>] -e <end_time> -n  
<application> -z <configuration_file> [-D 9]
```

[Using the `ddbmadmin` command to display save set information](#) on page 121 provides details.

- Display information for the save files within a save time range:

```
ddbmadmin -f [-b <start_time>] -e <end_time> -n <application> -  
z <configuration_file> [-D 9]
```

[Using the `ddbmadmin` command to display save file information](#) on page 122 provides details.

- Delete the save sets within a save time range:

```
ddbmadmin -d [-t] [-b <start_time>] -e <end_time> -  
n <application> -z <configuration_file> [-D 9] [-c]
```

[Using the `ddbmadmin` command to delete save sets](#) on page 124 provides details.

- Upgrade the SAP Oracle backup index from the database application agent 1.0 to release 4.0:

```
ddbadmin -u -n <application> -z <configuration_file>
```

[Using the ddbadmin command to upgrade the backup index](#) on page 125 provides details.

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

Table 17 The `ddbadmin` command options for save set display and deletion

Option	Description
<code>-b <start_time></code>	Optional. Specifies the start of the time range, in a date and time format. Date and time format used with the ddbadmin command options on page 120 provides details on the supported date and time formats for <code><start_time></code> . Without this option, the earliest backup time is used by default for the start of the time range.
<code>-c</code>	Optional. Specifies to run the operation in a noninteractive mode. Without this option, the operation is interactive by default.
<code>-d</code>	Specifies to perform a deletion of one or more backup save sets created during the specified time range.
<code>-D 9</code>	Generates debugging information during the operation. The option is used to troubleshoot operational issues.
<code>-e <end_time></code>	Mandatory with other options except the <code>-i</code> and <code>-u</code> options. Specifies the end of the time range, in a date and time format. Date and time format used with the ddbadmin command options on page 120 provides details on the supported date and time formats for <code><end_time></code> . The option setting <code>-e now</code> specifies the current time.
<code>-f</code>	Specifies to display information about the backup save files created during the specified time range.
<code>-i</code>	Specifies to display all the clients for the device path of the <code>DEVICE_PATH</code> parameter in the configuration file.
<code>-n <application></code>	Mandatory with all other options except the <code>-i</code> option. Specifies the application name to use for the deletion, display, or upgrade operation. A valid <code>application</code> value is <code>db2</code> , <code>oracle</code> , <code>saphana</code> , or <code>saporacle</code> .
<code>-s</code>	Specifies to display information about the backup save sets created during the specified time range.
<code>-t</code>	Specifies to display the location of the save sets in either the Data Domain system (active tier) or the cloud tier. Configuring the use of Data Domain Cloud Tier for data movement to the cloud on page 126 provides details about operations with the Data Domain cloud tier.

Table 17 The ddbmadmin command options for save set display and deletion (continued)

Option	Description
-u	Specifies to upgrade an SAP Oracle backup index from the database application agent 1.0 to the release 4.0 index format. The backup namespace used in the index is changed from “backup” in the database application agent 1.0 to “saporacle” in release 4.0.
-z <configuration_file>	<p>Mandatory. Specifies the complete pathname of the configuration file as described in Setting up the configuration file on page 78.</p> <hr/> <p>Note</p> <p>The <code>CLIENT</code> parameter is mandatory in the configuration file when you run the <code>ddbmadmin</code> command from a different system than the one where the backup was performed.</p> <p>A separate configuration file is required for each different set of mandatory parameters.</p>

Date and time format used with the ddbmadmin command options

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

- **Time of day**—A time of day is in the form `hh[:mm[:ss]]` (or `hhmm`) [*meridian*] [*zone*]. If you do not specify a meridian (am or pm), a 24-hour clock is used. You can specify a time of day as just `hh` followed by a *meridian*. If you do not specify a timezone (for example, GMT), then the current timezone is used, as determined by the second parameter, *now*.
- **Date**—A date is a specific month and day, and possibly a year. The acceptable formats are `mm/dd[/yy]` and `month_name dd[, yy]`. If omitted, the year defaults to the current year. If you specify a year as a number in the range 70 and 99, 1900 is added. If a year is in the range 00 and 30, 2000 is added. The treatment of other years less than 100 is undefined. If a number is not followed by a day or relative time unit, the number is interpreted as a year if a *time_of_day*, *month_name*, and *dd* have already been specified; otherwise, it will be treated as a *time_of_day*.
- **Day**—A day of the week can be specified. The current day is used if appropriate. A day can be preceded by a number, indicating which instance of that day is preferred; the default is 1. Negative numbers indicate times past. Some symbolic numbers are accepted: last, next, and the ordinals first through twelfth (second is ambiguous, and is not accepted as an ordinal number). The symbolic number next is equivalent to 2; thus, next monday does not refer to the coming Monday, but refers to the one a week later.
- **relative time**—Specifications relative to the current time can be used. The format is [*number*] *unit*. Acceptable units are decade, year, quarter, month, fortnight, week, day, hour, minute, and second.

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

The actual date is formed as follows. First, any absolute date and/or time is processed and converted. Using that time as the base, day-of-week specifications are added. Last, relative specifications are used. If a date or day is specified, and no absolute or

relative time is given, midnight is used. Finally, a correction is applied so that the correct hour of the day is produced after allowing for daylight savings time differences.

Using the ddbmadmin command to display clients for a device path

You can run the following `ddbmadmin` command to display all the clients for a specified device path on the Data Domain system:

```
ddbmadmin -i -z <configuration_file> [-D 9]
```

[Table 17](#) on page 119 provides details about the command options. The configuration file must contain the required parameters including `DEVICE_PATH`.

Typically, you run this operation to prepare for save set deletion with the `ddbmadmin` command. The following example shows the `ddbmadmin -i` command output.

Example 2 Displaying the clients for a device path

You run the following `ddbmadmin` command to display the clients for a device path specified by the `DEVICE_PATH` parameter in the configuration file:

```
ddbmadmin -i -z /vnxspace1/ddbda20/initBOS.utl
```

The command output displays one row for each client:

```
List of clients for device path '/ddbda-linux':
client = hades.backup
client = artemis.backup
client = eros.backup
client = ate.backup
client = erebos
client = chronos
client = hera.backup
client = athena
client = eros.lss.emc.com
```

Using the ddbmadmin command to display save set information

You can run the following `ddbmadmin` command to display information about the backup save sets:

```
ddbmadmin -s [-t] [-b <start_time>] -e <end_time> -n <application> -z <configuration_file> [-D 9]
```

[Table 17](#) on page 119 provides details about the command options.

Typically, you run this operation to prepare for save set deletion with the `ddbmadmin` command. You can compare the save set information with details in the database backup catalog and determine which backup save sets to delete.

The `ddbmadmin -s` command output includes one row for each save set found, where the row contains the client name, save set date, save set size, and save set name. The rows are sorted in descending order by the save set date and time. The following keywords appear in the command output:

- `client`—Hostname of the client that created the save set.
- `date and time`—Date and time when the save set was created.
- `size`—Size of backup data in the save set.
- `name`—Name of the save set.

The following example shows the `ddbmadmin -s` command output.

Example 3 Displaying the save set information

You run the following `ddbmadmin` command to display the save sets from SAP Oracle backups:

```
ddbmadmin -s -b 'Mar 1 12:00:00 2015' -e 'Apr 8 12:52:29 2015' -n
saporacle -z /vnxspace1/dbappagent40/initBOS.utl
```

The command output displays one row for each save set:

```
Save set information:
client = eros.lss.emc.com, date and time = 03/23/15 11:24:07
(1427124247), size = 20982944, name = backint:BOS
client = eros.lss.emc.com, date and time = 03/23/15 11:24:09
(1427124249), size = 20982944, name = backint:BOS
client = eros.lss.emc.com, date and time = 03/23/15 11:24:11
(1427124251), size = 20982948, name = backint:BOS
client = eros.lss.emc.com, date and time = 03/23/15 11:24:13
(1427124253), size = 20982944, name = backint:BOS
client = eros.lss.emc.com, date and time = 03/23/15 11:30:21
(1427124621), size = 9024, name = backint:BOS:PS:
client = eros.lss.emc.com, date and time = 03/23/15 11:30:23
(1427124623), size = 10496, name = backint:BOS:PS:
client = eros.lss.emc.com, date and time = 03/23/15 11:30:25
(1427124625), size = 12816, name = backint:BOS:PS:
client = eros.lss.emc.com, date and time = 03/23/15 11:30:27
(1427124627), size = 10412, name = backint:BOS:PS:
client = eros.lss.emc.com, date and time = 03/23/15 11:30:29
(1427124629), size = 9032, name = backint:BOS:PS:
client = eros.lss.emc.com, date and time = 03/23/15 11:30:43
(1427124643), size = 72768, name = backint:BOS
Total number of save sets = 10.
```

Using the `ddbmadmin` command to display save file information

You can run the following `ddbmadmin` command to display information about the backup save files:

```
ddbmadmin -f [-b <start_time>] -e <end_time> -n <application> -
z <configuration_file> [-D 9]
```

Table 17 on page 119 provides details about the command options. The configuration file must contain the required parameters.

Typically, you run this operation to prepare for save set deletion with the `ddbmadmin` command. You can compare the save file information with details in the database backup catalog and determine which backup save sets to delete.

The `ddbmadmin -f` command output includes one row for each save file found, where the row contains the save file name and the backup date and time. The following example shows the `ddbmadmin -f` command output.

Example 4 Displaying the save file information

You run the following `ddbmadmin` command to display the save files from SAP Oracle backups:

```
ddbmadmin -f -b 'Mar 1 12:00:00 2015' -e 'Apr 8 12:52:29 2015' -n
saporacle -z /vnxspace1/dbappagent40/initBOS.utl
```

The command output displays one row for each save file:

```
/, date = 1427124643 Mon Mar 23 11:30:43 2015.
/data_disk2//, date = 1427124643 Mon Mar 23 11:30:43 2015.
/vnxspace1//, date = 1427124643 Mon Mar 23 11:30:43 2015.
/home//, date = 1427124643 Mon Mar 23 11:30:43 2015.
/home/, date = 1427124643 Mon Mar 23 11:30:43 2015.
/vnxspace1/, date = 1427124643 Mon Mar 23 11:30:43 2015.
/data_disk2/, date = 1427124643 Mon Mar 23 11:30:43 2015.
/home/oracle/, date = 1427124643 Mon Mar 23 11:30:43 2015.
/home/oracle/app/, date = 1427124643 Mon Mar 23 11:30:43 2015.
/vnxspace1/dbappagent40/, date = 1427124643 Mon Mar 23 11:30:43 2015.
/data_disk2/oradata-bos/, date = 1427124643 Mon Mar 23 11:30:43 2015.
/home/oracle/app/oracle/, date = 1427124643 Mon Mar 23 11:30:43 2015.
/home/oracle/app/oracle/product/, date = 1427124643 Mon Mar 23
11:30:43 2015.
/data_disk2/oradata-bos/sapreorg/, date = 1427124643 Mon Mar 23
11:30:43 2015.
/data_disk2/oradata-bos/sapbackup/, date = 1427124643 Mon Mar 23
11:30:43 2015.
/home/oracle/app/oracle/product/11.2.0/, date = 1427124643 Mon Mar 23
11:30:43 2015.
/home/oracle/app/oracle/product/11.2.0/dbhome_2/, date = 1427124643
Mon Mar 23 11:30:43 2015.
/home/oracle/app/oracle/product/11.2.0/dbhome_2/dbs/, date =
1427124643 Mon Mar 23 11:30:43 2015.
/data_disk2/oradata-bos/sapreorg/strucBOS.log, date = 1427124643 Mon
Mar 23 11:30:43 2015.
/data_disk2/oradata-bos/sapreorg/spaceBOS.log, date = 1427124643 Mon
Mar 23 11:30:43 2015.
/vnxspace1/dbappagent40/initBOS.utl, date = 1427124643 Mon Mar 23
11:30:43 2015.
/home/oracle/app/oracle/product/11.2.0/dbhome_2/dbs/initBOS.ora, date
= 1427124643 Mon Mar 23 11:30:43 2015.
:
:
```

Using the ddbmadmin command to delete save sets

You can run the following `ddbmadmin` command to delete the backup save sets created during a specified time range:

```
ddbmadmin -d [-t] [-b <start_time>] -e <end_time> -n <application> -z <configuration_file> [-D 9] [-c]
```

[Table 17](#) on page 119 provides details about the command options. The configuration file must contain the required parameters.

For example, you can run the `ddbmadmin` command to delete all the save sets older than one year. The command deletes the save set data in the `.ss` files and the index metadata in the `.rec` files associated with the save sets.

NOTICE

Use the `ddbmadmin` command with caution because the command deletes save sets without considering the dependencies between the save sets in a backup. The `ddbmadmin` command uses the backup save time for deletions, and deletes all the backups that are stored under the same device path in the deletion time range. After you delete certain save sets from a backup, you might be unable to restore data from the backup. Therefore, improper save set deletion with the `ddbmadmin` command can lead to failed restores and data loss.

The deletion of save sets cannot be undone and must not be interrupted once started. For applications that support backup deletion, such as DB2 with the automatic deletion of recovery objects, use the supported application interface to delete the obsolete backups properly.

Before the `ddbmadmin` program starts to delete save sets, the program prompts you to confirm the deletion. If the program finds at least one uncommitted save set due to a running backup or a crashed backup, the program displays warning messages that refer to incomplete save sets. You can either terminate or continue the deletion.

The following examples show the command output from an interactive deletion (without the `-c` option) and the results when you terminate or continue the deletion after an uncommitted save set is detected.

Example 5 Terminating a deletion after the detection of an incomplete save set

You run the following `ddbmadmin` command to delete save sets from SAP Oracle backups:

```
ddbmadmin -d -b '03/23/2015 11:24:09' -e '03/23/2015 11:24:11' -n saporacle -z /vnxspace1/dbappagent40/initBOS.utl
```

The command output shows the list of save sets. The `n` response to the prompt causes the program to terminate the deletion:

Example 5 Terminating a deletion after the detection of an incomplete save set (continued)

```

Save set information:

client = eros.lss.emc.com, date and time = 03/23/15 11:24:09
(1427124249), size = 20982944, name = backint:BOS
client = eros.lss.emc.com, date and time = 03/23/15 11:24:11
(1427124251), size = 20982948, name = backint:BOS
Total number of save sets = 2.
The '/var/opt/dbdba/logs/ddbmadmin.messages.log' file contains a
list of the save sets to be deleted.
Continue with the deletion of the found save sets [y/n]: n

No save sets were deleted.

```

Example 6 Completing a deletion after the detection of an incomplete save set

You run the following `ddbmadmin` command to delete save sets from SAP Oracle backups:

```

ddbmadmin -d -b '03/23/2015 11:24:09' -e '03/23/2015 11:24:11' -
n saporacle -z /vnxspace1/dbappagent40/initBOS.utl

```

The command output shows the list of save sets. The `y` response to the prompt causes the program to continue and complete the deletion:

```

Save set information:

client = eros.lss.emc.com, date and time = 03/23/15 11:24:09
(1427124249), size = 20982944, name = backint:BOS
client = eros.lss.emc.com, date and time = 03/23/15 11:24:11
(1427124251), size = 20982948, name = backint:BOS
Total number of save sets = 2.
The '/var/opt/dbdba/logs/ddbmadmin.messages.log' file contains a
list of the save sets to be deleted.
Continue with the deletion of the found save sets [y/n]: y
List of deleted save set save times:

save time = 1427124249
save time = 1427124251

A total of 2 save sets were deleted successfully.

```

Using the `ddbmadmin` command to upgrade the backup index

Typically, a `backint` or `ddbmadmin` operation with the database application agent 4.0 automatically migrates the SAP Oracle backups that were created with the database application agent 1.0 to release 4.0. The database application agent 1.0 stores the SAP Oracle backups under the "backup" namespace in the application agent's catalog, whereas the database application agent 4.0 stores the backups under the "saporacle" namespace.

In the case where the automatic migration fails, run the following `ddbmadmin` command to upgrade the SAP Oracle backup manually:

```
ddbmadmin -u -n <application> -z <configuration_file>
```

Table 17 on page 119 provides details about the command options. The configuration file must contain the required parameters.

Note

The `ddbmadmin` command only converts the "backup" namespace name that was used by the database application agent 1.0 for SAP Oracle backups. If the namespace name is not converted, a restore by the database application agent 4.0 cannot find the required backup in the "backup" namespace.

The following example shows the `ddbmadmin -u` command output.

Example 7 Upgrading the SAP Oracle backup index

You run the following `ddbmadmin` command to upgrade the backup index for SAP Oracle backups from the database application agent 1.0 namespace to the database application agent 4.0 namespace:

```
ddbmadmin -u -n saporacle -z /dbtools/oracle/aix/64bit/product/11.1.0.6/dbs/initSAP.utl
```

```
The backup upgrade can take a long time and it should not be
interrupted.
The '/var/opt/ddbda/logs/ddbmadmin.messages.log' file contains a
list of items upgraded.
Do you want to continue with the upgrade [y/n]: y

Backup was upgraded successfully.
```

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

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

Data Domain (DD) Cloud Tier is a native feature of DD OS 6.0 and later for data movement from the active tier to low-cost, high-capacity object storage in the public, private, or hybrid cloud for long-term retention. The database application agent 3.5 introduced support of the DD Cloud Tier for movement of DD Boost backup data to the cloud, which frees up space on the Data Domain system (active tier).

Note

The database application agent does not support the DD Cloud Tier for movement of ProtectPoint backup data to the cloud.

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

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

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

A backup with the database application agent consists of backup save sets, where a save set is a collection of one or more save files created during the backup session. A save file is an operating system file or block of data, the simplest object that you can back up or restore. A backup creates one or more save files within a save set. The database application agent moves and recalls the backup data at the save set level only, moving all the save files in a save set.

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

Setting up the DD Cloud Tier policy for data movement to the cloud

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

DD Cloud Tier provides two types of policy, the application-based policy and the age-based policy. The database application agent supports only the application-based policy, which is managed by the application that creates the backup files on the Data Domain system. This policy moves the backup file content to the cloud according to the application's specifications.

NOTICE

Do not apply an age-based policy to a storage unit that is used by the database application agent. An age-based policy moves all the file content (including metadata) from a storage unit to the cloud according to the file age, as when all the files older than T days are moved. Such data movement by an age-based policy can cause the failure of metadata queries for the database application agent.

The DBA must contact the Data Domain administrator to create the application-based policy, also known as a data movement profile, for the MTree or storage unit that the database application agent uses for the DD Boost backups. The Data Domain documentation provides details about the DD Cloud Tier configuration procedures.

For policy management from a Linux system command line, you can use REST APIs through `curl` commands from the command line. You can run the commands to perform specific operations for the DD Cloud Tier policies or you can embed the commands in a script. Refer to following subtopics for more details and examples of how to use the REST APIs.

As an alternative with DD OS 6.1 or later, you can run the Data Domain command `data-movement policy` to configure the application-based policy.

Using the data-movement command with DD OS 6.1 or later

DD OS 6.1 or later enables you to configure the application-based policy through the following Data Domain command from the command line. This command sets the application-based policy for the specified Mtrees:

```
data-movement policy set app-managed {enabled | disabled} to-tier
cloud cloud-unit <unit-name> mtrees <mtree-list>
```

For example, the following command sets the application-based policy for the Mtree /data/coll/app-agent40:

```
data-movement policy set app-managed enabled to-tier cloud cloud-
unit Cloud mtrees /data/coll/app-agent40
```

You can run the following command to display the policy configuration result for verification purposes:

```
data-movement policy show
```

Mtree	Target (Tier/Unit Name)	Policy	Value
/data/coll/app-agent40	Cloud/Cloud	app-managed	enabled

Using REST APIs through curl commands

For example, you can run the following types of `curl` commands from a Linux command line to use REST APIs to perform operations for the DD Cloud Tier policies:

Note

In the following commands, you must replace any variables that start with the characters `$dd`, such as `$dduser` and `$ddhost`, with the appropriate values from your system setup. For example, the environment variables have the following settings. The `dduser` variable must be set to the DD administrator user:

```
ddhost="datadomain1.company.com"
dduser="ddadmin_user1"
ddmtree="/data/coll/sul"
ddcloudunit="ecs1"
ddpass="ddadmin_user1_password"
```

You must also replace `$TOKEN` in the commands with the result from the authentication command.

- The following example command performs an authentication:

```
curl --silent -k -i -X POST -H "Content-Type: application/json"
-d '{"username":"'${dduser}', "password":"'${ddpass}'"}'
https://$ddhost:3009/rest/v1.0/auth
```

- The following example command lists the policy for an MTree:

```
curl -k -D -i -X GET https://$ddhost:3009/rest/v1.0/dd-
systems/0/data-movement-policies -H "content-type:application/
xml" -H "X-DD-AUTH-TOKEN: $TOKEN" -d
"<data_movement_create><data_movement_policy_type>app-managed</
```



```
data_movement_policy_type><mtree_name>$ddmtree</mtree_name><cloud_unit_name>$ddcloudunit</cloud_unit_name></data_movement_create>"
```

- The following example command creates a policy on an Mtree:

```
curl -k -D -i -X POST https://$ddhost:3009/rest/v1.0/dd-systems/0/data-movement-policies -H "content-type:application/xml" -H "X-DD-AUTH-TOKEN: $TOKEN" -d "<data_movement_create><data_movement_policy_type>app-managed</data_movement_policy_type><mtree_name>$ddmtree</mtree_name><cloud_unit_name>$ddcloudunit</cloud_unit_name></data_movement_create>"
```

- The following example command deletes a policy from an MTree:

```
curl -k -D -i -X DELETE https://$ddhost:3009/rest/v1.0/dd-systems/0/data-movement-policies -H "content-type:application/xml" -H "X-DD-AUTH-TOKEN: $TOKEN" -d "<data_movement_create><data_movement_policy_type>app-managed</data_movement_policy_type><mtree_name>$ddmtree</mtree_name><cloud_unit_name>$ddcloudunit</cloud_unit_name></data_movement_create>"
```

Example Linux script to create a policy

For example, you can run the following type of script from a Linux system command line to create a DD Cloud Tier policy:

Note

In the script, you must replace any variables that start with the characters \$dd, such as \$dduser and \$ddhost, with the appropriate values from the particular system setup. If the #ddpass= line in the script is commented out as follows, the script prompts for the password of the DD user that is specified in the dduser setting. The dduser variable must be set to the DD administrator user.

If you copy and paste the following script into an editor, ensure that all the characters are transferred correctly, especially the characters at the ends of the lines. Some copy and paste operations might cause the end-of-line characters to be omitted.

```
#!/bin/bash
ddhost="datadomain1.company.com"
dduser="ddadmin_user1"
ddmtree="/data/Coll/sul"
ddcloudunit="ecs1"
#ddpass=ddadmin_user1_password

if [ -z $ddpass ] ; then
read -s -p "Enter DD Administrator Password: " ddpass
fi
TOKEN=`curl --silent -k -i -X POST -H "Content-Type: application/json" -d '{"username":"'"$dduser"'", "password":"'"$ddpass"'"}' https://$ddhost:3009/rest/v1.0/auth | gawk '/X-DD-AUTH-TOKEN: ./ { print $2; }'`
curl -k -D -i -X POST https://$ddhost:3009/rest/v1.0/dd-systems/0/data-movement-policies -H "content-type:application/xml" -H "X-DD-AUTH-TOKEN: $TOKEN" -d "<data_movement_create><data_movement_policy_type>app-managed</data_movement_policy_type><mtree_name>$ddmtree</mtree_name><cloud_unit_name>$ddcloudunit</cloud_unit_name></data_movement_create>"
```

Performing the data movement to the cloud

After you have set up the DD Cloud Tier policies, you can run the `ddbmadmin -m` command for the manual movement of backup data to the cloud.

As described in the previous topic, you set up the policy management script to specify the schedule for movement of backup data to the cloud.

A DBA can run the `ddbmadmin -m` command to mark the save set data files and a copy of the metadata for movement to the cloud. The command options specify the backup save sets that were created during a time range:

```
ddbmadmin -m [-t] [-b <start_time>] -e <end_time> -n <application> -z <configuration_file> [-D 9] [-c]
```

Note

For optimal Data Domain performance, it is recommended that you keep a backup on the Data Domain system (active tier) for at least 14 days before you move the backup to the cloud.

When the data movement policy is run, the data and metadata copy are then moved to the cloud. The original metadata remains on the Data Domain system.

The `-m` option specifies to mark the specified backup save sets and a copy of the corresponding metadata for movement to the cloud. [Table 17](#) provides details on the other command options. The configuration file must contain the required parameters.

A separate configuration file is required for each Data Domain system. All the information that the `ddbmadmin` command prints to standard output is added to the operational log file, `ddbmadmin.messages.log`.

For example, the following command marks for movement all the Oracle backup save sets within the time range from one month ago to the current time:

```
ddbmadmin -m -b '1 month ago' -e now -n oracle -z /config/oracle.cfg
```

When you run the `ddbmadmin -m` command, the program displays the list of save sets to be moved and prompts you to confirm the list for data movement. Then the program marks the save set files for movement. The files are moved to the cloud at the time determined by the data movement policy.

Performing the data recall from the cloud

After backup save sets have been moved to the cloud through the DD Cloud Tier policies, the database application agent by default recalls the save sets back to the Data Domain system. As an alternative, you can run the `ddbmadmin -r` command to manually recall the save sets. You can also disable the automatic recalls by the database application agent.

By default, the database application agent automatically recalls save sets from the cloud to the Data Domain system as needed to complete a restore operation.

You can manually recall the save sets from the cloud by running the `ddbmadmin -r` command.

Note

When data is recalled from the cloud, the data is actually removed from the cloud and moved back to the Data Domain system. You must ensure that the data is moved back to the cloud as required.

You can optionally set the `DDBOOST_AUTO_RECALL_DATA` parameter in the configuration file to specify the preferred method of data recall from the cloud:

- To enable the automatic recall method, set `DDBOOST_AUTO_RECALL_DATA=TRUE`. `TRUE` is the default value of the parameter. The database application agent automatically recalls data as needed for restore operations.
- To enable the manual recall method, set `DDBOOST_AUTO_RECALL_DATA=FALSE`. The database application agent can recall data only after you run the `ddbmadmin -r` command to initiate the data recall.

Note

When you use the ECS with DD OS 6.1 or later, the restores from the cloud are always seamless. The database application agent automatically restores the backup data directly from ECS, regardless of the `DDBOOST_AUTO_RECALL_DATA` setting. The Data Domain documentation provides more details and recommended practices on restores directly from ECS.

The manual recall method is recommended for restore scenarios where the recalls from the cloud is slow and might cause potential issues with server timeout.

When you enable the manual recall method, a restore operation fails when the database application agent tries to read a data file that has been moved to the cloud. The restore failure displays the following type of error message:

```
Unable to recover data with save time '1477335338'. Recall the file
from Data Domain Cloud Tier, and restart the recovery.
```

Based on the error message, determine the list of backup save sets within a specified time range to recall from the cloud. A DBA can run the `ddbmadmin` command with the `-r` option to specify the recall of the backup save sets:

```
ddbmadmin -r [-t] [-b <start_time>] -e <end_time> -n <application> -
z <configuration_file> [-D 9] [-c]
```

The `-r` option specifies to recall the specified backup save sets from the cloud. The `-t` option specifies to display the location of the save sets in either the Data Domain system (active tier) or the cloud tier. [Table 17](#) provides details on the other command options. The configuration file must contain the required parameters.

Note

A separate configuration file is required for each Data Domain system. All the information that the `ddbmadmin` command prints to standard output is added to the operational log file, `ddbmadmin.messages.log`.

For example, the following command specifies to recall the Oracle backup save sets within the time range from one month ago to the current time:

```
ddbmadmin -r -b '1 month ago' -e now -n oracle -z /config/oracle.cfg
```

When you restart the restore operation, the database application agent performs the specified data recall from the cloud and then completes the restore operation.

General troubleshooting tips

Review the following information about troubleshooting general issues that you might encounter in operations with the database application agent.

Debug log settings

The following table describes parameters that you can set in the configuration file to specify the debug log settings for the database application agent. You typically use these parameters when you work with Technical Support to troubleshoot issues with the product. Do not use these parameters for regular product operations.

[Setting up the configuration file](#) on page 78 describes how to set parameters in the configuration file.

[Debug log files](#) on page 133 describes the naming conventions for the debug log files.

For each parameter, the following table lists the section heading of the configuration file section that contains the parameter.

Table 18 Parameters for debugging

<p>Parameter: DEBUG_LEVEL</p> <p>Section: [GENERAL]</p> <p>Specifies whether the software writes debug messages to the debug log file, located in the directory specified by the <code>DIAGNOSTIC_DEST</code> parameter.</p> <hr/> <p>Note</p> <p>Use this parameter for debugging purposes with assistance from Technical Support only.</p> <hr/> <p>Valid values:</p> <ul style="list-style-type: none"> 0 (default) = The software does not generate debug messages. 9 = The software writes debug messages to the debug log file with a .log file name extension.
<p>Parameter: DIAGNOSTIC_DEST</p> <p>Section: [GENERAL]</p> <p>Specifies the directory location of the debug logs generated when the <code>DEBUG_LEVEL</code> parameter is set to 9.</p> <hr/> <p>Note</p> <p>The operational logs generated during normal product operations are not affected by this parameter setting.</p> <hr/> <p>Valid values:</p>

Table 18 Parameters for debugging (continued)

- Default directory of the debug log files:
 - /opt/dpsapps/dbappagent/logs (UNIX or Linux)
 - C:\Program Files\DPSAPPS\DBAPPAGENT\logs (Windows)
- Valid pathname of the directory of the debug log files.

Note

The default directory is used for SAP HANA or SAP Oracle if the specified nondefault directory does not exist or does not have write permissions for the OS user that runs the `hdbbackint` or `backint` program.

Parameter: DPRINTF**Section: [GENERAL]**

Specifies whether the software writes additional debug messages to the debug log file.

Note

For DB2 on Windows, it is recommended that you do not set this parameter to TRUE, especially for a DB2 multisession restore.

Valid values:

- FALSE (default) = The software does not generate additional debug messages.
- TRUE = The software writes additional debug messages to the debug log file.

Debug log files

The database application agent programs generate debugging information in debug log files with specific names.

Debug logs created by the ddbmadmin operations

- **Debug log created with `-D 9` option:**
`ddbmadmin_<yyyy>_<mm>_<dd>.<timestamp>.<pid>.log`
 For example: `ddbmadmin_2014_02_24.10_39_11.18678.log`
- **Regular log:** `ddbmadmin.messages.log`

Debug logs created for DB2 operations

- **Operational and error message log:** `ddboost_db2.messages.log`
- **Initial default log:** `libddboostdb2_default.log`
- **Session debug log:** `libddboostdb2_DB2_<timestamp>.<process/thread_id>.log`
- **XBSA error message log:** `xbsa.messages`

Debug logs created for Oracle operations

- **Operational and error message log:** `ddbda_oracle.messages.log`
- **Session debug log:**
`libddboostora_Oracle_<date>.<time>.<process_id>.log`

Debug logs created for SAP HANA operations

- **Operational log:** `hdbbackint<SID>.log`
For example: `hdbbackintEMC.log`
- **Database, archived log, and catalog file backup logs:**
 - `hdbbackint<SID>.debug.<pid>.log`
For example: `hdbbackintEMC.debug.11403.log`
 - `LGTOSAPs.debug.<SID>.<pid>.log`
For example: `LGTOSAPs.debug.EMC.11419.log`
- **Database recovery log:** `LGTOSAPr.debug.<SID>.<pid>.log`
For example: `LGTOSAPr.debug.EMC.14158.log`
- **Error log:** `hdbbackintHANA_<hostname>.op.<pid>.log`
For example: `hdbbackintHANA_fs1.op.14158.log`

Debug logs created for SAP with Oracle operations

- **Database and catalog file backup logs (backint backups):**
 - `backint<SID>.debug.<pid>`
For example: `backintSAP.debug.984`
 - `LGTOSAPs.debug.<SID>.<pid>`
For example: `LGTOSAPs.debug.SAP.1002`
- **Archived log backup logs (backint backups):**
 - `arch_backint<SID>.debug.<pid>`
For example: `arch_backintSAP.debug.4018`
 - `LGTOSAPs.debug.<SID>.<pid>`
For example: `LGTOSAPs.debug.SAP.4434`
- **RMAN operational and error message log:**
`ddboost_saporacle_rman.messages.log`
- **RMAN session debug log:**
`libddboostsapora_Oracle_<date>.<time>.<process_id>.log`

Backup or restore fails due to an inaccessible lockbox

An operation with the database application agent might fail with the following error message:

```
Unable to retrieve the primary device user password from the lockbox
```

If this happens and the lockbox is not shared with any other host, then you must run the `ddbmadmin -P -z <configuration_file>` command to register the Data Domain systems in the lockbox.

Command `ddbmadmin -P` encounters a conflict with an installed application

If another installed application uses the same libraries as the database application agent in the `/opt/dpsapps/dbappagent/lib/lib64` directory, then the `ddbmadmin -P -z <configuration_file>` command might encounter a conflict with the application.

Ensure that the software path and library path are set correctly before you run any `ddbmadmin` command.

For example, run the following commands before you run any `ddbmadmin` command on Linux:

```
# export PATH=/opt/dpsapps/dbappagent/lib/lib64:$PATH
# export LD_LIBRARY_PATH=/opt/dpsapps/dbappagent/lib/lib64:$LD_LIBRARY_PATH
```

Note

The `LD_LIBRARY_PATH` environment variable applies to Linux and Solaris. Set `LIBPATH` for AIX, and `SHLIB_PATH` for HP-UX. On Windows, ensure that the library path for the database application agent appears at the front in the `Path` environment variable setting, similar to the environment variable settings on UNIX and Linux.

Lockbox creation might fail on an NFS/CIFS share

The lockbox creation might fail with the following error on an NFS/CIFS share if the permissions on the share are insufficient:

```
The Lockbox file could not be opened.
```

Confirm the permissions on the NFS/CIFS share, and assign sufficient permissions as required for the lockbox creation.

Lockbox creation procedure when UAC is enabled on Windows

With UAC enabled on a Windows system, a user who logs in as a member of the Windows Administrator group is unable to create the lockbox in the default location by running the `ddbmadmin` command. For example, the lockbox creation with the `ddbmadmin -P -z configuration_file` command fails with the following error message:

```
Cannot create the directory 'C:\Program Files\EMC DD Boost\DA\config\lockbox'.
```

The Windows user can run the `ddbmadmin` command in the **Command Prompt** window as an administrator:

1. Click **Start**.
2. Right-click **Command Prompt**.
3. Select **Run as administrator**.

4. Run the required `ddbmadmin` command in the open **Command Prompt** window.

Major system update can produce an error about lockbox stable value threshold

When a host first accesses a stand-alone or shared lockbox, certain System Stable Values (SSVs) are stored in the lockbox for the host. The database application agent requires a specific number of the SSVs to be matched for the host for each subsequent lockbox access.

When a major update of the host system causes multiple SSVs to change, the required number of SSVs will not match when the host tries to access the lockbox during a backup or restore operation. In this case, the host's attempt to access the lockbox produces the following error:

```
The Lockbox stable value threshold was not met because the system fingerprint has changed. To reset the system fingerprint, open the Lockbox using the passphrase.
```

If you encounter this error, you must complete the required operation to enable the lockbox access for the host:

- In a stand-alone system, re-create the lockbox for the host and perform the registration again.
- In a high-availability system with a shared lockbox:
 1. Revoke the lockbox access of the host by running the `ddbmadmin -R` command from another host.
 2. Grant the lockbox access to the host by running the `ddbmadmin -G` command from another host.

Restore fails if the CLIENT parameter setting does not match the backup client name

A restore by the database application agent fails if the restore uses the `CLIENT` parameter setting but the parameter value does not match the hostname recorded in the backup by the database application agent.

For example, a DB2 restore to a new host uses the `CLIENT` parameter setting `saps2d01.vsp.com` to specify the original host that was backed up. However, the `CLIENT` parameter was not set for the backup, and the short hostname `saps2d01` was recorded in the backup by default. In this case, the restore fails because the `CLIENT` parameter setting for the restore is not identical to the hostname recorded in the backup.

The restore fails with the following type of error message:

```
Opening the directory 'directory_pathname' failed ([5004] nothing matched) .
```

Ensure that the `CLIENT` parameter setting for the restore matches the hostname recorded in the backup on the Data Domain system. You can usually run the `hostname` command at the operating system level to obtain the default value of the hostname used in the backup if the original host is available. You can also contact Technical Support to check for the hostname recorded in the backup.

ProtectPoint specific troubleshooting tips

Review the following information about troubleshooting issues that you might encounter in ProtectPoint operations with the database application agent.

ProtectPoint operation might fail due to configuration issues

Before you start a ProtectPoint operation, ensure that the database and software are configured properly.

A ProtectPoint operation might fail due to the following configuration issues:

- The `ddbsm` daemon or service is not running.
- Issues exist in the database application agent configuration file.
- In a snapshot backup, the database does not reside on snapshotable LUNs, such as VMAX devices in a ProtectPoint VMAX backup.
- The required Solutions Enabler software is not installed, especially when a VMAX `symapi` operation is involved.

The follow examples show some of the corresponding error messages:

```
Failed to set up SymApi handle: Unable to get installed Solutions
Enabler version, vmax error code 1, Solutions Enabler is not
installed, /usr/lib64/libsymmlvm64.so: cannot open shared object file:
No such file or directory
```

Application-specific operations also produce specific types of error messages:

- A DB2 ProtectPoint operation produces the following error message:

```
An error occurred while accessing media "libddbboostdb2.so".
Reason code: "11".
```

- An Oracle ProtectPoint operation produces one of the following error messages:

```
> RMAN-03002: failure of backup command at 04/27/2017 06:18:23
> ORA-27203: skgfpqb: sbtpcquerybackup returned error
> ORA-19511: non RMAN, but media manager or vendor specific failure, error text:
> pb_init() failed with: 3 (113:501:111)

RMAN-00571: =====
RMAN-00569: ===== ERROR MESSAGE STACK FOLLOWS =====
RMAN-00571: =====
RMAN-03002: failure of backup command at 06/23/2017 16:08:53
RMAN-06416: PROXY ONLY was specified and some files could not be proxy copied
```

Determine the cause of the error on the particular system, and resolve the issue before you rerun the ProtectPoint operation.

Restore might fail due to an incompatible file system or volume manager version

A restore might fail when the target file system is incompatible with the backed-up file system version.

The restore might fail when the file system or volume manager version on the target host is earlier than the version on the backed-up file systems. For example, with Veritas Volume Manager, the failed restore displays the following error message:

```
Disk group has no valid configuration copies
```

As a workaround, ensure that the file system or volume manager version on the target host is the same as or later than the version on the backed-up file systems, and then rerun the restore. The file system and volume manager documentation provides details.

Troubleshooting the ProtectPoint for VMAX restores

You can set the optional parameter `RESTORE_FROM_DD_ONLY` in the configuration file to assist with troubleshooting of issues with the ProtectPoint for VMAX restore operations. Use this parameter for debugging purposes only.

When you set the `RESTORE_FROM_DD_ONLY` parameter to the default value of `FALSE`, the database application agent first tries to restore a ProtectPoint for VMAX backup from the latest local snapshot on VMAX. If the local snapshot is not available, then the database application agent restores from the Data Domain system. The parameter setting of `TRUE` specifies to restore from the Data Domain system only.

Manual cleanup of FAST.X restore devices after a failed restore of a ProtectPoint for VMAX backup

When a ProtectPoint for VMAX restore fails due to errors or a user-initiated cancel request, the restore might not perform an automatic cleanup, including a cleanup of the restore devices state. In such a case, some manual cleanup steps might be required so that the devices can be used in the next restore. To enable the next restore operation to use the FTS devices, the devices must not have any locks and must be in the Not Ready state.

You can perform the following manual cleanup steps to enable the FTS devices for the next restore operation.

1. To list the FTS devices that can be used for a restore operation, run the following command:

```
symmsg -sid <SymID> show <VMAX_RESTORE_FASTX_SG_name>
```

For example:

```
symmsg -sid 758 show ledma179_sg02
...
Devices (3):
{
-----
Sym          Device          Cap
Dev   Pdev Name   Config      Attr  Sts   (MB)
-----
00260  /dev/sdn      TDEV                NR    8629
00261  /dev/sdy      TDEV                NR    8629
00262  /dev/sdab     TDEV                NR    8629
}
```

2. To check if the devices are locked, run the following command:

```
symdev -sid <SymID> -lock -devs <devSymID1:devSymIDn> list
```

3. To release a lock on the devices, run the following command:

```
symdev -sid <SymID> -lock -devs <devSymID1:devSymIDn> release
```

4. To check whether a device is in the Image Refresh state, run the following command:

```
symdev -sid <SymID> show <devSymID> | grep -i status
```

For example:

```
symdev -sid 758 show 0260 | grep -i status
Device Status   : Not Ready   (NR)   (Image Refresh)
...
```

5. To move the device out of the Image Refresh state, run the following command:

```
symdev -sid <SymID> cancel_image_refresh <devSymID>
```

For example:

```
symdev -sid 758 cancel_image_refresh 260
Execute a 'Cancel Image Refresh' Device operation for device
'260' (y/[n])? y

'Cancel Image Refresh' Device operation successfully completed
for the device.
```

Note

- The cancel operation does not work on a locked device. Release the lock on the device if needed.
- Ensure that the devices are not used by any running ProtectPoint restore operations. The cancellation interferes with ongoing operations and may cause unexpected behavior.

6. To move the device to the Not Ready state for use in the next restore, run the following command:

```
symdev -sid <SymID> not_ready <devSymID>
```


CHAPTER 5

DD Boost Operations on DB2 Systems

This chapter includes the following topics:

- [Overview of DD Boost operations in a DB2 environment](#).....142
- [Configuration of DD Boost operations in a DB2 environment](#)..... 144
- [Performing DD Boost backups and recovery with the DB2 CLP](#)..... 150
- [Performing DD Boost backups and restores with the DB2 GUI](#)..... 158
- [Performing DD Boost backup data recovery with utility programs](#).....158
- [Preparing for DB2 disaster recovery](#)..... 167
- [DB2 DPF requirements for DD Boost operations](#)..... 168
- [DB2 HADR requirements for DD Boost operations](#).....168
- [DB2 pureScale requirements for DD Boost operations](#).....170
- [DB2 troubleshooting tips for DD Boost operations](#).....174

Overview of DD Boost operations in a DB2 environment

The database application agent is integrated with the DB2 interfaces for third-party media management vendors to enable DB2 data backups, restores, transaction log archiving, and backup deletion (pruning).

You can perform a backup or recovery with the product on a DB2 database server by running one of the supported DB2 backup or recovery tools:

- DB2 Command Line Processor (CLP) with the commands `db2 backup`, `db2 restore`, `db2 rollforward`, and `db2 recover`
- IBM Data Studio GUI (DB2 10.1 and later)
- DB2 Control Center GUI (DB2 9.x)

You can use these tools in cooperation with the database application agent to perform the following DB2 operations:

- Online and offline backups
- Full, incremental, and delta backups
- Archived log backups
- Recovery of a database to the current time or a specific point-in-time
- Recovery to the original location or an alternate location
- Backup and recovery of databases, tablespaces, and archived logs
- Backup deletion (pruning)

The product maintains online backup indexes on the Data Domain system, which are in addition to the DB2 history file. During backups, the product creates backup entries in the online indexes, which provide the information required to restore the backed-up data.

During an offline backup, no other application can use the database that is backed up. The restore of an offline backup does not require a rollforward operation. For a recoverable database, as configured with `logarchmethn`, you must specify the `without rolling forward` option explicitly to skip the rollforward operation. Without the rollforward operation, the restore of an offline backup brings the database back to its state at the backup time.

During an online backup, other applications can connect to the database that is backed up. Online backups are only available for recoverable databases. The restore of an online backup requires a rollforward operation, which uses the active or archived logs to restore to either a specific point-in-time or the end of the active logs.

The DB2 documentation provides details about these features and the recovery strategy for databases.

The troubleshooting section at the end of this chapter provides details about limitations in the DD Boost operations with the database application agent in a DB2 environment.

DD Boost DB2 backup processes

A DB2 database backup includes the following process interactions.

1. The database administrator initiates the backup by running the `db2 backup` command, IBM Data Studio GUI, or DB2 Control Center GUI.

2. The DB2 software loads the shared library used by the database application agent.

Note

IBM documentation refers to this library as a vendor library.

3. The database application agent reads the configuration file and initializes the data transport media.
4. The DB2 software sends the database or log data to the database application agent, which uses the DD Boost interface to send the data to the Data Domain system for storage.

DD Boost DB2 restore processes

A DB2 database restore includes the following process interactions.

1. The database administrator initiates the restore by running the DB2 CLP, the IBM Data Studio GUI, or the DB2 Control Center GUI.
2. The DB2 software loads the shared library used by the database application agent.
3. The database application agent reads the configuration file and initializes the data transport media.
4. The DB2 software requests the backup from the database application agent.
5. The database application agent performs the following tasks:
 - a. Queries the index on the Data Domain system to locate the backup data.
 - b. Retrieves the backup data from the Data Domain system.
 - c. Sends the backup data to the DB2 server.

DB2 backups of transaction logs

DB2 software provides two types of transaction logging:

- Circular logging is the default behavior when you create a DB2 database. With this type of logging, each full backup deletes the transaction logs. You can restore only full backups. Circular logging supports only full offline backups of databases.
- Archived logging supports online backups and rollforward recovery. With this type of logging, the transaction logs are retained as archived logs that can be managed by the database application agent. You can recover a database or tablespace to a specific point-in-time by using the rollforward feature. You can recover the archived logs and apply the transactions in the logs in sequence until the specified point, to the end of the backup or the end of the logs.

Ensure that the DB2 archived transaction logs are backed up, for example, by configuring the automatic backup of DB2 transaction logs as described in [Configuring DB2 transaction log archiving](#) on page 146.

Note

For a database that was dropped and then re-created with the same name, ensure that you correctly configure its archived log backups. [DB2 issues due to incorrect log retrieval](#) on page 178 provides details on how to prevent issues with archived log backups for such a re-created database.

After a product software update or deployment of a new vendor library `libddbboostdb2.x`, ensure that the DB2 database is restarted or reactivated if it has been activated, which cleans up the image of the old loaded library in the DB2 log manager process. [Configuring DB2 transaction log archiving](#) on page 146 provides details.

Configuration of DD Boost operations in a DB2 environment

You must complete the required configurations of the database application agent to enable the DD Boost operations in a DB2 environment. The following topics provide the product configuration details.

The troubleshooting section at the end of this chapter provides details about limitations in the DD Boost operations with the database application agent in a DB2 environment.

Integrating the product into the DB2 environment

After the product installation on Windows, verify that the system `%PATH%` environment variable includes the software installation directory:

```
C:\Program Files\DPSAPPS\DBAPPAGENT\bin
```

Note

The directory pathname can include spaces, but there cannot be spaces before or after the pathname.

Restart the DB2 database manager after the product installation.

Configuring the DB2 parameters for DD Boost operations

You must set the required parameters for DB2 operations in the configuration file used by the database application agent.

[Setting up the configuration file](#) on page 78 describes the common parameters and how to set parameters in the configuration file. [Configuring restores of replicated backups](#) on page 90 also describes the parameters and requirements for the restores of replicated backups.

Ensure that the configuration file also contains the appropriate DB2 parameters from the following table. For each parameter, the table lists the section heading of the configuration file section that contains the parameter.

After the configuration file is set up, ensure that the required lockbox procedures have been performed as described in [Configuring the lockbox](#) on page 103.

Table 19 DB2 parameters for DD Boost operations

Parameter: DB2_ALIAS

Table 19 DB2 parameters for DD Boost operations (continued)

Section: [GENERAL]

Specifies the name of the source DB2 database that was used for the backup of the database or archived logs.

Optional for use with the manual backup retrieval program, `ddbmdb2adutil`. Corresponds to the `-a <database>` command option of the program. Use the command option or set this `DB2_ALIAS` parameter in the configuration file.

Valid values:

- Undefined (default).
- Valid name of the DB2 database that was used for the database or log backup.

Parameter: DB2INSTANCE**Section: [GENERAL]**

Specifies the name of the source DB2 instance that was used for the backup and contains the archived logs to be retrieved.

Mandatory for a restore to a database of a different instance and for a recovery and rollforward operation to a database of a different instance.

Optional for use with the manual backup retrieval program, `ddbmdb2adutil`. Corresponds to the `-I <DB2_instance>` command option of the program. Use the command option or set this `DB2INSTANCE` parameter in the configuration file.

Note

Ensure that you set the parameter to the correct value. The DB2 documentation provides details.

Ensure that you set the parameter in the `logarchoptn` configuration file as the parameter is mandatory for a log retrieval operation. In accordance with IBM DB2 restrictions, the length of the `logarchoptn` setting must be less than 30 characters, including the leading symbol `@`.

Valid values:

- Undefined (default).
- Valid name of the source DB2 instance that contains the logs to be retrieved.

Parameter: DB2_NUM_NODE**Section: [GENERAL]**

Specifies the node number or partition number of the node for which the backup was performed. In a pureScale environment, specifies the member ID.

Optional for use by the `ddbmdb2adutil` program to retrieve a database backup or log backup.

Valid values:

- 0 (default).
- Valid node number or partition number of the node for which the backup was performed. Valid member ID in a pureScale environment.

Parameter: SOURCE_CLIENT**Section: [GENERAL]**

Specifies the source client hostname under which the backup was indexed.

Table 19 DB2 parameters for DD Boost operations (continued)

<p>This parameter is used in a rollforward recovery of the destination database after a restore to a different client where a rollback is possible.</p> <p>Optional for a rollforward after a redirected restore.</p> <p>Valid values:</p> <ul style="list-style-type: none"> Undefined (default). If not specified, the <code>CLIENT</code> value is used. Valid hostname of the source client as used in the backup.
<p>Parameter: SOURCE_DBNAME</p> <p>Section: [GENERAL]</p> <p>Specifies the name of the source DB2 database that was originally used for the backup of the archived logs.</p> <p>Mandatory for a database recovery and rollforward operation to a different database.</p> <p>Valid values:</p> <ul style="list-style-type: none"> Undefined (default). Valid name of the source DB2 database that was originally used for the backup of the database or archived logs.

Configuring DB2 transaction log archiving

When you configure the archiving of DB2 transaction logs, the database application agent performs the log backups based on DB2 database policy settings. The product has no control over when the logs are backed up or how often. The DB2 server starts the backup when a transaction log becomes full and when the database is active.

To deactivate and activate the DB2 database, run the following commands:

```
db2 deactivate database <database_name>
db2 activate database <database_name>
```

To list the active DB2 databases, run the following command:

```
db2 list active databases
```

Complete the required steps to configure the DB2 transaction log archiving with the database application agent.

Procedure

1. Create a configuration file for the archived log backup, which can be the same as or different than the configuration file used for a database backup and restore.
2. Configure the database with the command and options appropriate for the client operating system:

- On UNIX:

```
db2 update db cfg for sample using logarchmeth1
'vendor:library_path/libddboostdb2.so' logarchopt1
'@pathname/db2_tlogs.cfg'
```

- On Windows:

```
db2 update db cfg for sample using logarchmeth1 'vendor:C:
\Program Files\DPSAPPS\DBAPPAGENT\bin\libddboostdb2.dll'
logarchopt1 '@pathname\db2_tlogs.cfg'
```

where:

- *library_path* is the directory pathname of the library used by the database application agent for DB2 operations.
- *sample* is the name of the database to be backed up.
- *pathname/db2_tlogs.cfg* or *pathname\db2_tlogs.cfg* is the complete pathname of the DB2 configuration file. Do not specify a relative pathname.

NOTICE

The `logarchoptn` setting is used to fill in the `COMMENT` field in the DB2 history archived log records, which is restricted to a maximum length of 30 characters. This value is then used during the pruning of the archived logs and during the rollforward stage in a `db2 recover` operation. For the operation to succeed, the `logarchoptn` setting must be less than 30 characters in length, including the leading symbol `@`.

3. When you complete the configuration, perform an offline backup as required by IBM DB2. [Performing DD Boost backups and recovery with the DB2 CLP](#) on page 150 includes an example of how to run a DB2 backup.

Configuring DB2 backup deletion

The automatic deletion or pruning of the database recovery history by the DB2 database manager is based on the setting of the configuration parameters `num_db_backups` and `rec_his_retentn`. When you set the DB2 database configuration parameter `auto_del_rec_obj` to on, the DB2 database manager automatically performs the following maintenance operations when both of the `num_db_backups` and `rec_his_retentn` values are exceeded:

- Prunes the database history.
- Deletes the corresponding backup images, load copy images, and log files. These objects are generally referred to as DB2 recovery objects.

These maintenance operations can be performed as part of a backup.

Ensure that the DB2 backup history and the backup configuration and storage are synchronized. Otherwise, the automatic backup object removal might fail with the following error in the DB2 diagnostic log:

```
MESSAGE : SQL2062N An error occurred while accessing media "".
Reason code:
      "".
```

When the DB2 backup history and the backup configuration and storage are not synchronized, the DB2 vendor operations can become degraded. A backup can also become blocked when the required pruning cannot be completed.

Ensure that the following requirements have been met:

- The original configuration file, specified in `logarchoptn` as recorded in the DB2 history file, exists on the system with the required full pathname that will be used in the backup object removal. Without this configuration file, the removal fails.
- It is recommended that you do *not* use `ddbmadmin` to manually remove the backup object when the removal can be managed by the DB2 retention policy. The removal through `ddbmadmin` does not clean up the DB2 history.

If you use `ddbmadmin` for manual deletion, ensure that the DB2 history is also synchronized with the vendor backup storage.

You can perform the following steps to enable the automatic deletion of DB2 backups.

Procedure

1. Set the DB2 database configuration `vendoropt` parameter to the pathname of the configuration file for the DB2 database or tablespace whose backups are to be deleted. For example:

```
db2 update db cfg for sample using vendoropt '@pathname/  
db2_dbbda.cfg'
```

where:

- *sample* is the name of the database or tablespace whose backups are to be deleted.
 - `pathname/db2_dbbda.cfg` is the complete pathname of the configuration file.
2. Enable automatic pruning of the recovery history:

```
db2 update db cfg for sample using num_db_backups n  
db2 update db cfg for sample using rec_his_retentn  
number_of_days
```

3. Enable the automatic deletion of physical backup images and log files:

```
db2 update db cfg for sample using auto_del_rec_obj on
```

where *sample* is the name of the database whose backups are to be deleted.

Note

Without this step, the recovery history pruning removes entries only in the DB2 history file and does not remove the associated backups.

You can also prune the history file and delete the backups manually by using the `db2 prune` command.

The DB2 documentation provides details about the automatic deletion of DB2 recovery objects and the `db2 prune` command.

Preventing deletion of DB2 backup images

The automatic deletion and manual pruning might accidentally remove backup images that are still needed, especially when the backup images retention management relies only on the automatic deletion configuration and manual pruning of recovery objects.

If you want to exclude certain backups from automatic deletion and retain those backups for a longer time, update the status of the associated entries in the recovery history file to `do_not_delete` (acronym X). For example, run the following command:

```
db2 update history EID 10 with status X
```

```
DB20000I The UPDATE HISTORY command completed successfully.
```

When you perform manual pruning, do not use the `with force` option for the status setting to be effective. If you use this option, the backup entries are still pruned.

When you no longer need to keep the backup images, update the status of the entries back to normal, for example, active (A) or expired (E). This updated status enables the pruning and removal of those backup images according to the automatic deletion configuration in place for other backup images. The DB2 documentation provides more details about protecting recovery objects from being deleted.

Estimating the Data Domain resource usage on DB2 systems

The following topics provide additional guidelines and best practices related to the Data Domain resource usage for DB2 systems.

In a DB2 multisession backup and restore, the sessions typically run in parallel and each session acquires its own Data Domain capacity and streams resources.

In a cluster environment, such as a Database Partitioning Feature (DPF) environment, the backups can be run for each node or for all the nodes in parallel.

In an online database backup, the archive log backups typically run in parallel with the database backup.

For the operations that can run in parallel, take into account their Data Domain resource usage for the total usage estimate in the corresponding timeframe.

When an archive log backup fails due to the lack of a Data Domain resource (or any other cause), it fails silently and the DB2 server retries the backup. Ensure that you monitor the `db2diag.log` file to fix any issues.

Capacity usage on DB2 systems

Consider the information in this topic if you want to estimate the amount of space that the Data Domain system requires for backups.

The estimated capacity requirement should include the size of possible archive log backups, based on the level of database activity and the configured size of the log. You can configure the size of the log through the DB2 database configuration. When the database activity increases, more logs can be archived and backed up.

If the storage capacity of the Data Domain system is exceeded, the backup operation fails. The database application agent generates the following type of error message in the operational log `ddboost_db2.message.log`:

```
163542 12/8/2016 11:54:16 AM (pid2640) SYSTEM critical Unable to
write to a file due to reaching the hard quota limit.
The error message is: [5194] [ 2640] [10476] Thu Dec 08 11:54:16 2016
ddp_write() failed Offset 109051904, BytesToWrite 524288,
BytesWritten 0 Err: 5194-Hard Quota Exceeded >
```

Streams usage on DB2 systems

Note

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

DB2 database operations use the following numbers of streams:

- For a database backup, the number of used streams is typically equal to the number of sessions plus 1.
- For an archive log backup, typically 2 streams are used.
- For a database restore and recovery, the number of used streams is typically 1.5 x the number of sessions.
- In a multinode environment, such as a DPF environment, multiply the estimated streams number by the number of nodes that perform the backup or restore.
- For a rollforward operation, typically 2 streams are used.

For example, in a four-session backup, the streams usage estimate for the backup must be at least 5. For an online backup, the streams estimate must be incremented by 2.

If the Data Domain system runs out of streams during a backup, the database application agent generates the following error message in the DB2 operational and debug log:

```
153004 05/06/2016 08:43:05 AM (pid25976) SYSTEM critical Unable to
write to a file because the streams limit was exceeded.
The error message is: [5519] [25976] [139683900024608] Fri May 6
08:43:05 2016
ddp_write() failed Offset 0, BytesToWrite 524288, BytesWritten 0 Err:
5519-Exceeded streams limit
```

Performing DD Boost backups and recovery with the DB2 CLP

You can run the DB2 CLP to perform backups, restores, and recovery with the database application agent. The DB2 documentation provides details about the DB2 CLP commands and options.

Performing DB2 backups with the DB2 CLP

You can perform a DB2 backup after you have completed the backup configurations in [Configuration of DD Boost operations in a DB2 environment](#) on page 144.

You can run the appropriate `db2 backup` command to perform a DB2 backup.

Note

The database application agent includes support of the `dedup_device` option for backups with the `db2 backup` command. The `dedup_device` option optimizes the format of data that is backed up to a deduplication device.

For example on UNIX, run the following command:

```
db2 backup db sample online load 'library_path/libddbboostdb2.so'
open n sessions options '@pathname/db2_ddbda.cfg' dedup_device
```

where:

- *sample* is the name of the database to back up.
- `online` specifies to perform an online backup. The default backup type is an offline backup.
- *library_path* is the directory pathname of the library used by the database application agent for DB2 operations.
- *n* is the number of concurrent backup sessions.
- *pathname/db2_ddbda.cfg* is the pathname of the DB2 configuration file as described in [Configuring the DB2 parameters for DD Boost operations](#) on page 144.

The DB2 documentation provides more details about how to use the `db2 backup` command for manual backups.

Performing load operations with the copy yes option

The database application agent supports the DB2 `load` command with the `copy yes` option. The `load` command efficiently loads large amounts of data at the page level into a DB2 table. The IBM documentation provides complete details about the `load` command and its features and options.

The `copy yes` option of the `load` command specifies that a copy of the loaded data is saved (in the form of a database backup piece) to a specified location and can be used in the database recovery. The `load copy` image can be stored in the same location as archive log backups (recommended) or in any other location:

- You specify the location of the `load copy` image through the `LOAD` or `TO` option.
- You specify the location of the archive log backups through the database configuration parameter `logarchmethn`.

A `load` operation with the `copy yes` option allows the database to be recovered through the `load copy` event time without any special handling. The `copy yes` option is valid only if rollforward recovery is enabled.

A `load` operation on a recoverable database without the `copy yes` option puts the database into a backup pending state after the operation completes.

After you have enabled rollforward recovery for a database and configured DB2 backups with the database application agent, you can run the `load` command with the `copy yes` option. The operation saves a copy of the DB2 database changes (during the `load` operation) to a specified vendor through the `load <library_name>` option, or to a specified directory or device through the `TO` option. A subsequent rollforward recovery of the database loads the copy of the saved changes directly into the database.

The same logging and debugging procedures apply to `load copy` operations as for other DB2 DD Boost operations with the database application agent.

Configuration requirements for the load with copy yes option

Before you run the `load` command with the `copy yes` option, ensure that you meet the following configuration requirements:

- You have installed and configured the database application agent according to the instructions in the preceding chapters.
- You have completed all the required post-installation procedures, including the lockbox configuration procedure.
- You have ensured that the database is recoverable by configuring at least one of the `logarchmethn` database parameters and then performing an offline full database backup.

For example, the following command configures the `logarchmeth1` parameter:

```
db2 update db cfg for <database_name> using logarchmeth1
'vendor:/opt/dpsapps/dbappagent/lib/lib64/libddbboostdb2.so'
logarchopt1 '@/space1/cfg/dd.cfg'
```

- You have set the `vendoropt` parameter to the pathname of the configuration file for the database to be recovered. The `vendoropt` configuration is used in the backup during the `load` time and in a search of the `load copy` backup image during the recovery time.

For example, the following command sets the `vendoropt` parameter:

```
db2 update db cfg for <database_name> using vendoropt '@/
space1/cfg/dd.cfg'
```

Performing the load operation with the copy yes option

After you have met the configuration requirements, you can run the `load` operation with the `copy yes` option to save the loaded data. For example, the following operation loads the data in delimited ASCII format (DEL) from the file `/space1/tmp/load.txt` into table `table1` and to the database application agent through the vendor library as specified by the `copy yes` option:

```
db2 "load from /space1/tmp/load.txt of DEL insert into table1 copy
yes LOAD /opt/dpsapps/dbappagent/lib/lib64/libddbboostdb2.so"
```

Note

Due to IBM issue IT08141, a warning about the number of sessions that are used in the operation might appear in the debug log. This issue does not affect the backup of the `load copy` image or the subsequent recovery that uses the backup image. The issue was first fixed in DB2 9.7 Fix Pack 11. The [IBM website](#) provides more details about issue IT08141.

You can run the DB2 `list history` command to verify the `load` operation and view the saved `load copy` record in the DB2 recovery history:

```
db2 list history load all for db <database_name>
```


In the following example, the restore operation restores the last backup of the database SAMPLE that was performed before the load operation. The rollforward operation locates and loads the load copy image (saved by the load operation) and any other archive logs directly into the database:

```
db2 restore db SAMPLE LOAD /opt/dpsapps/dbappagent/lib/lib64/
libddboostdb2.so options @/home/db2inst1/dd.cfg taken at
20160704120000
db2 rollforward db SAMPLE to end of logs and stop
```

Displaying save information from the load copy backup

When you run a load operation with the copy yes option, the loaded data is saved as a database backup in the vendor namespace. The backup type is classified as LOAD_COPY, as defined by the DB2 vendor and administrative APIs.

You can run the ddbmadmin command in verbose mode (with the -v option) to view the save set and save file information from the load copy backup. [Configuring the display and deletion of save set information](#) on page 117 provides details about the ddbmadmin command and its options.

The load copy backup type is saved in the XBSA metadata of the backup. You can run the ddbmadmin -f -v command to view this backup type in the backup's save file information.

The following example ddbmadmin -v -f command displays the save file information from a load copy backup, including the LOAD_COPY level:

```
ddbmadmin -v -f -b "07/05/2016 12:36:05 PM" -e "07/05/2016 12:36:05
PM" -n db2 -z /space1/mycfg/dd.cfg
```

```
141540:ddbadmin:The parameter 'DB2 ALIAS' is being ignored.
/JT02/NODE0000:/DB_BACKUP.20160705123605.1, application = db2 (27),
date = 1467736565 Tue 05 Jul 2016 12:36:05 PM EDT.
version=1, objectowner= DB2, objectname=/JT02/NODE0000 /DB_BACKUP.
20160705123605.1, createtime=Tue 05 Jul 2016 12:36:05 PM EDT,
copytype=3 BSACopyType_BACKUP, copyId=1467736565.1467736566,
restoreOrder=1467736565.1, objectsize=0.0, resourcetype=database,
objecttype=4 BSAObjectType_DATABASE, objectstatus=2
BSAObjectStatus_ACTIVE, description=database app
agent_v40:DB2_v970:LOAD_COPY:JT02:TEQ, objectinfo=jt971:0.
```

The following example ddbmadmin -v -s command displays the save set information from the load copy backup:

```
ddbmadmin -v -s -e now -n db2 -z /space1/mycfg/dd.cfg
```

```
Record file = /tangdl/dd4/bu-today.lss.emc.com/27/2.0/meta_rec/JT971/
JT02/NODE0000/DBIMG/_ts10k_146773/1467736565. rec.
client = bu-today.lss.emc.com, date and time = 07/05/2016 12:36:05
PM, size = 279008, ssid = 1467736565, name = DB2:/JT02/NODE0000
ssid=00ec186e-00000011-00000000-577belfa-577belf5-68009e56
(1467736565), date and time=07/05/2016 12:36:05 PM (1467736565),
host=bu-today.lss.emc.com, name=DB2:/JT02/NODE0000, continuedfrom=0,
level=incr, sflags=0, size=279008,files=1, insert=07/05/2016,
create=07/05/2016, complete=07/05/2016, browse=forever,
retent=02/07/2106 01:28:15 AM,clientid=0, attrs=\
*ACTUAL_HOST: bu-today.lss.emc.com;
*ss data domain backup cloneid: 1467736565;
*ss data domain dedup statistics: "v1:1467736565:279752:67159:2231";
```

```
index subspace: JT971/JT02/NODE0000/DBIMG;
record file name: /tangd1/dd4/bu-today.lss.emc.com/27/2.0/meta_rec/
JT971/JT02/NODE0000/DBIMG/_ts10k_146773/1467736565.rec;
, clones=0
```

Performing DB2 restores with the DB2 CLP

You can run the appropriate `db2 restore` command to perform a DB2 data restore to either the same DB2 application host or a different host.

A DB2 restore can restore the data to the original database or to a different database under the same or different DB2 instance.

Note

On AIX with Data Domain Fibre Channel (DFC), a DB2 multistream restore might fail or become suspended. You can resolve this issue by increasing the device count of the AIX client on Data Domain as described in [DB2 multistream restore and rollforward might fail on AIX with DFC](#) on page 174.

Performing DB2 restores to the same instance

You can perform the steps to restore the DB2 data to the original DB2 instance.

When you recover the data to a point-in-time, note the timestamp of the backup to restore.

Note

The `db2 restore` command without a timestamp always uses the latest database backup, even when there is a tablespace backup after the database backup. To restore the latest backup that is a tablespace backup, use the full timestamp.

You can run the `db2 restore` command with the appropriate options. For example:

- On UNIX, run the following command:

```
db2 restore db sample load 'library_path/libddbboostdb2.so' open
n sessions options '@pathname/db2_ddbda.cfg' taken at
yyyymmddhhmmss into sample2
```

where:

- *sample* is the name of the database to be restored.
- *library_path* is the directory pathname of the library used by the database application agent for DB2 operations.
- *n* is the number of restore sessions if the database application agent used multiple sessions for the backup.
- *pathname/db2_ddbda.cfg* is the pathname of the DB2 configuration file.
- *yyyymmddhhmmss* is the timestamp of the backup to restore.

Skip the `taken at` parameter if you are restoring only the most recent backup of a database.

- *sample2* is the new name of the database if you are restoring to a different database name.

Skip the `into` parameter if you are restoring the database to the original database name.

- On Windows, run the following command:

```
db2 restore db sample load 'C:\Program Files\DPSAPPS\DBAPPAGENT
\bin\libddboostdb2.dll' open n sessions options '@pathname
\db2_ddbda.cfg' taken at yyyymmddhhmmss into sample2
```

If the timestamp of the backup is unknown, find the timestamp by querying all the backups with the following command:

```
db2 list history backup all for sample
```

where *sample* is the name of the database to be restored.

You can also view the backup timestamp and type in the output of the `db2admin -f -v` command.

Performing DB2 restores to a different instance

You can perform the steps to restore the DB2 data to a different DB2 instance. The following steps include examples of commands and parameter settings.

Procedure

1. From the new instance, generate a redirection script by running the `db2 restore` command with the `redirect generate script` option:

- On UNIX, run the following command:

```
db2 restore db sample load 'library_path/libddboostdb2.so'
options '@pathname/db2_ddbda.cfg' taken at yyyymmddhhmmss
redirect generate script 'pathname/my_redirect.ddl'
```

where:

- `pathname/my_redirect.ddl` is the complete pathname of the generated redirection script.
 - The other command line options are the same as described in [Performing DB2 restores to the same instance](#) on page 154.
- On Windows, run the following command:

```
db2 restore db sample load 'C:\Program Files\DPSAPPS
\DBAPPAGENT\bin\libddboostdb2.dll' options '@pathname
\db2_ddbda.cfg' taken at yyyymmddhhmmss redirect generate
script 'pathname\my_redirect.ddl'
```

Note

Ensure that the new instance has read and write permission to the script.

2. Edit the generated script, and define the following parameters as applicable:
 - `OPTIONS` (mandatory)—Complete pathname of the configuration file used by the database application agent.
 - `ON`—Storage paths of the new database.
 - `DBPATH ON/TO`—Target database directory.

- `INTO`—New database name, if you are redirecting the recovery to a new name.
- `TAKEN AT`—Timestamp of the backup to recover, `yyyymmddhhmmss`, if you are restoring the data to a point-in-time.
- `OPEN SESSIONS`—Number of restore sessions, if the database application agent used multiple sessions for the backup.

For example:

```
OPTIONS '@/bigspace/db2_ddbda.cfg'
ON '/bigspace/db_data'
INTO sample2
```

Note

If the database backup includes DMS tablespaces, then you might need to set the `SET TABLESPACE CONTAINERS` parameter to the appropriate value.

The DB2 documentation provides details.

3. To invoke the redirection script under the redirected different instance where the data is to be restored, run the following command on the DB2 application host:

```
db2 -tvf my_redirect.ddl
```

where `my_redirect.ddl` is the name of the generated redirection script.

Performing DB2 recovery with the DB2 CLP

To recover a DB2 database to either the current time or a specific point-in-time, you can run the `db2 rollforward` command to apply the transaction logs that are stored on the Data Domain system.

If you want to restore and roll forward a DB2 database in a single operation, you can run the `db2 recover` command.

Note

To use rollforward recovery, the database application agent must have backed up the transaction logs. [DB2 backups of transaction logs](#) on page 143 provides details.

On AIX with Data Domain Fibre Channel (DFC), a DB2 multistream rollforward recovery might fail or become suspended. You can resolve this issue by increasing the device count of the AIX client on Data Domain as described in [DB2 multistream restore and rollforward might fail on AIX with DFC](#) on page 174.

You can also perform a recovery through the time of a `load copy yes` event. The rollforward operation locates and loads the `load copy` backup image (saved by the `load copy yes` operation) directly into the database when required.

Performing DB2 recovery with the db2 rollforward command

To apply all the transactions to the end of the logs, run the following command:

```
db2 "rollforward db sample to end of logs and complete"
```

where *sample* is the database name.

To apply the transactions to a specific point-in-time, specify the date and time in the command. For example, run the following command:

```
db2 "rollforward db sample to yyyy-mm-dd-hh.mm.ss using local time and complete"
```

Performing DB2 restore and recovery with the db2 recover command

The `db2 recover` command combines the functions of the `db2 restore` command and `db2 rollforward` command.

Procedure

1. Set the DB2 database configuration `vendoropt` parameter to the pathname of the configuration file for the database to be recovered. For example:

```
db2 update db cfg for sample using vendoropt '@pathname/db2_dbdba.cfg'
```

where:

- *sample* is the name of the database or tablespace to be recovered.
 - *pathname/db2_dbdba.cfg* is the complete pathname of the configuration file.
2. Run the `db2 recover` command with appropriate options.

To apply all the transactions to the end of the logs, run the following command:

```
db2 recover db sample to end of logs
```

To apply the transactions to a specific point-in-time, specify the data and time in the command. For example, run the following command:

```
db2 recover db sample to yyyy-mm-dd-hh.mm.ss using local time
```

The command line options in these examples are the same as described in [Performing DB2 restores to the same instance](#) on page 154.

Note

The `db2 recover` command does not support the `load syntax` or `options syntax` that is available with `db2 backup` and `db2 restore` commands. Instead, the `db2 recover` command uses information in the DB2 history file to determine what file to load during the recovery and uses the `VENDOROPT` variable to pass the options file.

For a dropped database, use the `db2 restore` and `db2 rollforward` commands to perform disaster recovery. You cannot use the `db2 recover` command for dropped databases.

Performing DD Boost backups and restores with the DB2 GUI

You can run the DB2 GUI to perform backups, restores, and recovery with the database application agent. The DB2 documentation provides details about all the GUI procedures.

You can perform a DB2 backup or restore after you have completed the configurations in [Configuration of DD Boost operations in a DB2 environment](#) on page 144.

Specify the product configuration file in the GUI by setting `VENDOROPT` to the value `@configuration_file_pathname`. For example:

```
@d:\db2_ddbda.cfg
```

The product configuration file must contain the parameter settings for the backup or restore as described in [Setting up the configuration file](#) on page 78.

Set the Vendor DLL to the name of the database application agent library in the GUI.

Note

After you select the DLL path with the DB2 Control Center for a Windows client, enclose the path with quotes or use a short file name (8.3 format). Otherwise, the backup returns an error similar to the following example:

```
SQL0104N  An unexpected token "Files\EMC" was found following "<identifier>".
Expected tokens may include:  "INCLUDE".
```

On AIX with Data Domain Fibre Channel (DFC), a DB2 multistream restore and rollforward might fail or become suspended. You can resolve this issue by increasing the device count of the AIX client on Data Domain as described in [DB2 multistream restore and rollforward might fail on AIX with DFC](#) on page 174.

Performing DD Boost backup data recovery with utility programs

You can run the `ddbmdb2adutil` utility or the IBM Optim High Performance Unload (HPU) utility to recover data from DD Boost backups that are performed by the

database application agent. The following topics provide details on how to run the utilities for the DD Boost backup data recovery.

Retrieving DB2 database backups and log backups with the `ddbmdb2adutil` utility

The database application agent provides the stand-alone `ddbmdb2adutil` program to directly retrieve database backups or archive log backups to a local directory. The program does not act through the DB2 server or invoke a DB2 restore or recovery operation. The utility can be run through the command line interface by a user who has the proper permissions to write the retrieved files to the destination directory.

You can run the `ddbmdb2adutil` command with the appropriate command options to retrieve a database backup or log backup that was performed by the database application agent:

- The command object-type option `DB` specifies to retrieve a database backup with a DB2 backup timestamp to a directory on the local host. The retrieved files can be used in a native backup validation tool or check tool. The files can also be used as disk backup images in a restore.
- The command object-type option `LOG` specifies to retrieve archive log backups to a directory on the local host. You can use the prefetched logs for a specific purpose, such as database recovery or log shipping in a high availability environment. For example, you can use the directory that contains the retrieved logs as the overflow log path in a DB2 rollforward operation.

The prefetch helps to reduce the time that is spent in query, search, and retrieval from the vendor storage during a database recovery operation.

Note

- The directory hierarchy structure of the location that contains the retrieved database or archive log backup must follow the DB2 standard for any operations to be performed at that location.
- The `ddbmdb2adutil` program retrieves the backup data. The program does not create any extra directory structures that are required by subsequent operations or clean up after such operations on the retrieved data.

The `ddbmdb2adutil` program retrieves the specified backup to a backup image file with a general 644 permission. The file name uses the DB2 disk-backup naming convention:

- Name of a log backup file:


```
S<nnnnnnn>.LOG
```

For example: `S0000007.LOG`
- Name of a database backup file:
 - For DB2 release earlier than 9.8:


```
<database_alias>.<type>.<instance_name>.NODE<nnnn>.CATN<nnnn>.<timestamp>.<sequence_number>
```

where `<type>` is the backup type:

 - 0—Full database level backup
 - 3—Tablespace level backup

- 4—Backup image generated by the `LOAD COPY TO` command

- For DB2 release 9.8 or later:

`<database_alias>.<type>.<instance_name>.DBPART<nnn>.<timestamp>.<sequence_number>`

For example: Pieces retrieved from a DB2 9.7 three-session full backup of database SAMPLE, instance DB2INST1, and NODE 0, taken at 20160902094203:

```
SAMPLE.0.DB2INST1.NODE0000.CATN0000.20160902094203.001
SAMPLE.0.DB2INST1.NODE0000.CATN0000.20160902094203.002
SAMPLE.0.DB2INST1.NODE0000.CATN0000.20160902094203.003
```

Note

The name of the database backup image is constructed by using the information that is recorded at the backup time. For a DB2 pre-9.8 database backup, the catalog node number is not recorded at the backup time. Therefore, the name of a pre-9.8 database backup file always includes CATN0000. If the catalog node number is different than 0, you must correct CATN0000 in the file name manually before you use the backup file.

When multiple backup files are retrieved, the `ddbmdb2adutil` utility retrieves the files sequentially.

When the `ddbmdb2adutil` program encounters an error during a backup file retrieval:

- For database backups, the program generates an error and exits.
- For archive log backups, the program continues with the next log backup within the specified range.

Requirements for DB2 backup retrieval with the `ddbmdb2adutil` command

Before you use the `ddbmdb2adutil` command for a DB2 backup retrieval, ensure that you have completed the required configurations for restores with the database application agent, as described in the preceding chapter. The lockbox configurations must be complete, and the required parameters must be set in the configuration file.

Certain parameter values from the configuration file can be specified as command options instead with the `ddbmdb2adutil` command. The command options take precedence over the corresponding restore parameters in the configuration file.

The `ddbmdb2adutil` command options and syntax

You can perform the backup retrieval by running the `ddbmdb2adutil` command on the command line. Certain command options are mandatory, including the `-z <configuration_file>` option that specifies the configuration file of the database application agent. You must specify the required parameters by either using the corresponding command line options or setting the parameters in the configuration file for the backup retrieval. The parameters are similar to the parameters that are used for restore operations.

You must run one of the following `ddbmdb2adutil` commands, depending on the object type of the backup retrieval:

- Run the following command for a database backup retrieval:

```
ddbmdb2adutil DB [-a <database>] [-c <client>] [-d
<destination_directory>] [-F] [-I <DB2_instance>] [-N
<node_number>] [-t <backup_timestamp>] -z <configuration_file>
```


- Run the following command for a log backup retrieval:

```
ddbmdb2adutil LOG [-a <database>] [-c <client>] -C <chain_ID> [-d <destination_directory>] [-F] [-I <DB2_instance>] [-N <node_number>] -S <start_log> -E <end_log> -z <configuration_file>
```

The following table describes the `ddbmdb2adutil` command options.

Table 20 Options of the `ddbmdb2adutil` utility for backup image retrieval

Option	Description
DB	Mandatory for a database backup retrieval. Specifies to retrieve a database backup image by timestamp.
LOG	Mandatory for a log backup retrieval. Specifies to retrieve an archive log backup.
-a <database>	Optional. Specifies the name of the database for which the backup was performed. You must either use this command option or set the <code>DB2_ALIAS</code> parameter in the configuration file.
-c <client>	Optional. Specifies the client name under which the backup was performed. The default value is the local hostname. You can either use this command option or set the <code>CLIENT</code> parameter in the configuration file.
-C <chain_ID>	Mandatory for a log backup retrieval. Specifies the log chain ID for the log sequence of the archive log backup. Valid value is between 0 and 9999999, inclusive.
-d <destination_directory>	Optional. Specifies an accessible directory where the retrieved backup files are saved. The directory is created if it does not exist, provided that the user has required permission. The default value is the program working directory.
-E <end_log>	Mandatory for a log backup retrieval. Specifies the upper bound of the log sequence numbers for the archive logs to be retrieved. Valid value is between 0 and 9999999, inclusive. Note You can run the <code>ddbmadmin -f -v -n db2</code> command with the other required options to obtain the list of available archive log backups in the backup storage. The chain and sequence ID numbers are embedded in the file name of the log backup. For example, run the following command to obtain the list of available archive log backups: <pre>ddbmadmin -f -v -n db2 -b <start_time> -e <end_time> -z <configuration_file></pre>
-F	Optional. Specifies to overwrite any existing files in the destination location that have the same names as the retrieved backup files.
-I <DB2_instance>	Optional. Specifies the name of the DB2 instance for which the backup was performed.

Table 20 Options of the ddbmdb2adutil utility for backup image retrieval (continued)

Option	Description
	You must either use this command option or set the <code>DB2INSTANCE</code> parameter in the configuration file.
<code>-N <node_number></code>	<p>Optional. Specifies the node number or partition number of the node for which the backup was performed.</p> <p>In a pureScale environment, it is the member ID. The default value is 0.</p> <p>You can either use this command option or set the <code>DB2_NODE_NUM</code> parameter in the configuration file.</p>
<code>-S <start_log></code>	<p>Mandatory for a log backup retrieval. Specifies the lower bound of the log sequence numbers for the archive logs to be retrieved.</p> <p>Valid value is between 0 and 9999999, inclusive.</p>
<code>-t <backup_timestamp></code>	<p>Optional. Use only for a database backup retrieval. Specifies the DB2 backup timestamp of the backup to be retrieved.</p> <p>Valid value is a full timestamp as <code>yyyymmddhhmmss</code> or a partial timestamp.</p> <p>When this option is not used, the latest available database backup is retrieved. When a partial timestamp is used, the backup with the closest later timestamp is retrieved.</p> <hr/> <p>Note</p> <p>To retrieve the latest backup that is a tablespace backup, use this option with a full timestamp.</p>
<code>-z <configuration_file></code>	<p>Mandatory. Specifies the configuration file that the database application agent must use for the operation.</p> <p>Typically, the configuration file contains the required and optional parameters that are not supported by the command options. The command options take precedence over the corresponding parameter settings in the configuration file.</p> <p>For example:</p> <pre data-bbox="464 1430 1460 1587"> DDBOOST_USER=qa_ost DEVICE_HOST=bu-croco.lss.emc.com DEVICE_PATH=/tangdl/dd01 CLIENT=bu-today.lss.emc.com DB2_ALIAS=SAMPLE DB2INSTANCE=db2inst1 DB2_NODE_NUM=0 </pre>

Examples of DB2 backup retrievals with the ddbmdb2adutil command

The following example shows the `ddbmdb2adutil` command and output from a database backup retrieval:

```
ddbmdb2adutil DB -F -d /tmp/backups -t 20160902094203 -z /tmp/cfg/dd.cfg
```

```
Starting the retrieval operation.
Proceeding with the backup image retrieval with the timestamp
'20160902094203', 3 pieces of backups, and the base file name 'SAMPLE.
0.DB2INST1.NODE0000.CATN0000.20160902094203'.
Successfully retrieved the database backup image file 'SAMPLE.
0.DB2INST1.NODE0000.CATN0000.20160902094203.001'.
Successfully retrieved the database backup image file 'SAMPLE.
0.DB2INST1.NODE0000.CATN0000.20160902094203.002'.
Successfully retrieved the database backup image file 'SAMPLE.
0.DB2INST1.NODE0000.CATN0000.20160902094203.003'.
Successfully retrieved the backup files.
```

The following example shows the `ddbmdb2adutil` command and output from a log backup retrieval:

```
ddbmdb2adutil LOG -d /tmp/backups -z /tmp/cfg/dd.cfg -C 0 -S 0 -E 5
```

```
Starting the retrieval operation.
Successfully retrieved the log file 'C0000000_S0000000.LOG' to the
output file 'S0000000.LOG'.
Successfully retrieved the log file 'C0000000_S0000001.LOG' to the
output file 'S0000001.LOG'.
Successfully retrieved the log file 'C0000000_S0000002.LOG' to the
output file 'S0000002.LOG'.
Successfully retrieved the log file 'C0000000_S0000003.LOG' to the
output file 'S0000003.LOG'.
Successfully retrieved the log file 'C0000000_S0000004.LOG' to the
output file 'S0000004.LOG'.
Invalid status: status = 5.
Could not find or retrieve the backup file 'C0000000_S0000005.LOG'.
No data was written. Removing the empty backup file '/tmp/backups/
S0000005.LOG'.
Successfully retrieved 5 of 6 backup files.
```

For a log backup retrieval, the utility searches for all the available log backups in the specified range.

The following example shows the `ddbmadmin` command output that lists the available archive log backup of the chain 10 in the specified time range. In this output example, the chain number C0000010 and sequence number S0000074 are embedded in the archive log backup file name, `/SAMPLE/NODE0000/DB2LOG/:/C0000010_S0000074.LOG`, for the database `SAMPLE` of instance `db2inst1`:

```
ddbmadmin -f -v -n db2 -b "08/23/2016 07:23:00 AM" -e "08/23/2016
07:24:00 AM" -z /space1/db2_ddbda.cfg | grep C0000010
```

```
/SAMPLE/NODE0000/DB2LOG/:/C0000010_S0000074.LOG, application = db2
(27), date = 1471962220 Tue 23 Aug 2016 07:23:40 AM PDT.
version=1, objectowner= DB2, objectname=/SAMPLE/NODE0000/DB2LOG/ /
C0000010_S0000074.LOG, createtime=Tue 23 Aug 2016 07:23:40 AM PDT,
copytype=3 BSACopyType BACKUP, copyId=1471962220.1471962221,
restoreOrder=1471962220.1, objectsize=0.0, resourcetype=L,
objecttype=2 BSAObjectType_FILE, objectstatus=2
```

```
BSAObjectStatus_ACTIVE, description=database app
agent_v30:DB2_v1051:LOG_IMAGE:SAMPLE, objectinfo=db2inst1:1.
/SAMPLE/NODE0000/DB2LOG7:/C0000010_S0000073.LOG, application = db2
(27), date = 1471962218 Tue 23 Aug 2016 07:23:38 AM PDT.
...
```

You can use the database backup and archive log backup images that are retrieved through the `ddbmbdb2adutil` commands to perform the database restore and rollforward operations.

When debugging is enabled, the database application agent might generate a considerable amount of debugging information on the console during a backup retrieval operation.

The `ddbmbdb2adutil` program uses the same operational and default debug logs that the database application agent uses for other DB2 backup and restore operations:

- Operational and error message log: `ddbboost_db2.messages.log`
- Default debug log: `libddbboostdb2_default.log`

Ensure that the user who runs the program has the required permission to write to the log files. After the program completes, ensure that the DB2 users have the required permission to write to the log files. Otherwise, the logs must be removed from their location prior to the next backup and restore operation.

When debugging is enabled, the `ddbmbdb2adutil` program also produces a debug log with a name that includes the program name, date, and process ID. For example:

```
ddbmbdb2adutil_DB2_2016_07_07.16_30_32.11131.log
```

Limitations with the `ddbmbdb2adutil` command

Consider the following limitations before you run the `ddbmbdb2adutil` command:

- For a `load copy` image, generally in a DB2 recovery, a rollforward through the time of the load event must use the backup from the original location as the backup vendor. Therefore, the `load copy` image that is retrieved to disk may not be used directly in a rollforward operation.
- For a multisession `load copy` backup with a large amount of blob data inserted, the DB2 backup check tool, `db2ckbkp`, might fail to validate the backup image that is retrieved to disk.

Recovering DB2 backup data with the IBM HPU utility

The IBM Optim High Performance Unload (HPU) utility is a high-speed stand-alone utility that unloads database data from a current database or from its backup. The HPU utility can run concurrently with the DB2 database manager, and can access the same physical files as the database manager.

The HPU utility can unload data from a DB2 database or tablespace backup, which can be a full or incremental backup. The HPU utility reads the data directly from the backup image or from the live database file through the SQL engine, bypassing the DB2 database manager. As a result, the utility provides quick and efficient recovery of discrete volumes of data.

For example, you can unload a table that was dropped from a live database (when restore of the database is not an option) by unloading the data from a backup with the `BACKUP CATALOG` option. You can then load the extracted data into the database.

To minimize the impact on a production system, you can use HPU to unload data for dropped or corrupted tables into a non-production system, where you can cleanse and prepare the data for load into the production system.

You can use the HPU utility for a data unload by running the `db2hpu` program with a control file, which enables you to define the options and instructions that the unload uses. The utility unloads data from the backup image to staging files and then writes the data to output files, which can be used with the `load` utility. In a multiple partition environment, the HPU utility supports both single and multiple output files (according to settings in the control file), and can redirect the output file to a remote host or the partition hosts. The utility can write in parallel to output directories across different physical devices.

The IBM DB2 documentation provides details about HPU and the HPU command line options and control file syntax.

The same logging and debugging procedures apply to the DB2 HPU operations as for other DB2 DD Boost operations with the database application agent.

Requirements for DB2 recovery with the HPU utility

You can use the HPU utility to unload and extract a discrete volume of data from a DB2 backup image (created by the database application agent) into an output file. You can then load the data from the output file into a DB2 database.

Ensure that you meet the following requirements before you use the HPU utility for a DB2 data recovery from a database application agent backup:

- The database application agent has been installed and configured according to the instructions in the preceding chapters.
- The DB2 `db2hpu` program has been installed and configured according to the appropriate DB2 documentation.
- You use an offline backup image when possible to help ensure the integrity and consistency of the unloaded data.
- You use an online backup image only when you are certain that no transactions took place during the most recent online backup against the objects that you will unload.
- You use a tablespace backup instead of a full database backup when possible to reduce the size of the backup image being read and enable a faster data unload.
- When a table is dropped from a database, you use the `USING BACKUP CATALOG` option for an unload from a backup that contains the table, or you re-create the table before you run the unload tool.
- You have determined the timestamp of the backup image, for example, 20160718061214. If you do not provide a timestamp, then the utility uses the latest backup.
- You have created the control file for the HPU operation, including a command to extract the data as shown in the following control file examples.
- On AIX, especially when you use the `USING BACKUP CATALOG` option for an unload with a multisession backup, you have ensured that the user process resource limit for data segment is set sufficiently high. You can set the resource limit through either of the following methods:
 - Configure the `db2hpu` program to run with the no memory limit option.
 - Configure the system resource for data segment memory limit by running `ulimit -d` or by editing the default setting in the system configuration file.
- You do not use the HPU utility with `libddbboostdb2` over a Data Domain Fibre Channel (FC) network. Due to a known limitation of the DD Boost library over FC with child processes, the database application agent does not support the HPU utility operations over an FC network connection.

Example 1: Recovery of database tables with the HPU utility

The following control file, `/home/tmp/hpu.ctl`, includes the required commands to unload and extract all the table data from the TEST2 database.

```
GLOBAL CONNECT TO TEST2 DB2 NO;
UNLOAD TABLESPACE
QUIESCE NO
LOCK NO
USING BACKUP DATABASE TEST2 LOAD "/opt/dpsapps/dbappagent/lib/lib64/
libddbboostdb2.so" OPTIONS OPEN 3 SESSIONS @/home/cfg/dd/dd.cfg
TAKEN AT 20160627164046;
SELECT * FROM test;
OUTFILE("/home/tmp/outfile.txt" REPLACE)
FORMAT DEL;
```

In this example control file:

- `/home/cfg/dd/dd.cfg` is the DB2 configuration file pathname for the database application agent.
- `20160627164046` is the timestamp of the backup image.
- `/home/tmp/outfile.txt` is the output file to which the data will be extracted.

You can then run the `db2hpu` command with the control file to unload and extract the data with the HPU utility. For example, the following command uses the example control file and generates the output file, `/home/tmp/outfile.txt`, that contains the extracted data:

```
db2hpu -f /home/tmp/hpu.ctl -i db2inst1
```

Example 2: Recovery of a tablespace with the HPU utility

The following control file, `C:\tmp\cfg\hpu.ctl`, includes the required commands to unload and extract the dropped tablespace TB02 by using the backup catalog.

```
GLOBAL DB2 NO;
USING BACKUP CATALOG HPU2 LOAD "C:\PROGRA~1\EMCDDB~1\DA\bin
\libddbboostdb2.dll" OPTIONS @C:\tmp\cfg\db2ddp.cfg TAKEN AT
20160718061214;
UNLOAD TABLESPACE TB02
QUIESCE NO
LOCK NO
OUTFILE("tmp02.txt" REPLACE)
FORMAT DEL;
USING BACKUP DATABASE HPU2 LOAD "C:\PROGRA~1\EMCDDB~1\DA\bin
\libddbboostdb2.dll" OPTIONS @C:\tmp\cfg\db2ddp.cfg TAKEN AT
20160718061214;
```

In this example control file:

- `C:\tmp\cfg\db2ddp.cfg` is the DB2 configuration file pathname for the database application agent.
- `20160718061214` is the timestamp of the backup image.
- `tmp02.txt` is the output file to which the data will be extracted.

You can then run the `db2hpu` command with the control file to unload and extract the data with the HPU utility. For example, the following command uses the example

control file and generates the output file, `tmp02.txt`, that contains the extracted data:

```
db2hpu -f C:\tmp\cfg\hpu.ct1 -i DB2
```

Preparing for DB2 disaster recovery

For a comprehensive disaster recovery plan, you must ensure that you can reconstruct the computing environment and all the DB2 server files associated with maintaining data on the application host.

Use the following guidelines to prepare for a disaster recovery of the DB2 server host:

- Maintain accurate and complete records of the network and system configurations. Keep all the original software media and the following items in a safe location:
 - Original operating system media and patches
 - Device drivers and device names
 - File system configuration
 - IP addresses and hostnames
- Ensure that you have a current full backup of the database and all the archived logs required for a rollforward operation.
- Save a copy of the product configuration file used for the DB2 backups of the database and archived logs.
- Confirm that the parameter setting of the corresponding options file, such as `logarchopt1`, from the source database image is valid on the destination host.

Plan to perform the following tasks during a disaster recovery.

Procedure

1. Set up the product configuration file to be used during the recovery, including the following parameter settings:
 - `CLIENT=source_client_hostname`
 - `DB2INSTANCE=source_database_instance`
 - `SOURCE_DBNAME=source_database_name`

[Configuring the DB2 parameters for DD Boost operations](#) on page 144 provides details.

Note

In a restore to a different host, where a rollforward is required and a rollback phase might occur, ensure that you set both `SOURCE_CLIENT` and `CLIENT` to point to the correct hosts. [DB2 troubleshooting tips for DD Boost operations](#) on page 174 provides more details.

2. Re-create the lockbox on the database host. [Configuring the lockbox](#) on page 103 provides details about the lockbox.
3. Ensure that all the required database and log backup images are available.

4. Run a `db2 restore` command to restore the database. For example:

```
db2 restore database database_name load 'library_path/
libddbboostdb2.so' options '@pathname/db2_ddbda.cfg' taken at
yyyymmddhhmmss
```

5. Perform a rollforward operation on the restored database, to the end of the logs or a point-in-time:

```
db2 rollforward database database_name
```

DB2 DPF requirements for DD Boost operations

You can perform DD Boost backups and restores in a DB2 Database Partitioning Feature (DPF) environment.

You must meet the following configuration requirements in a DB2 DPF environment:

- You have set up the database in the DPF environment according to the appropriate DB2 DPF documentation.
- You have installed the database application agent software on each node that will participate in backups or recovery.
- You have completed all the post-installation procedures on each node, including the lockbox configuration procedure for all the participating hosts. The hosts can use either a shared lockbox or individually configured lockboxes.
- If an NFS-shared lockbox is configured, you have followed all the required steps in [Configuring the lockbox in a high-availability environment](#) on page 115.
- You have set the `CLIENT` parameter in the DB2 configuration file to the hostname of the catalog node. [Common parameters](#) on page 80 provides details on the parameter.
- You have ensured that all the partitions have the same parameter settings.

DB2 HADR requirements for DD Boost operations

You can perform DD Boost backups and restores with the database application agent in a DB2 High Availability Disaster Recovery (HADR) environment. You can start a backup on the primary node only. DB2 does not support backups on standby nodes. The database application agent only supports HADR setups where the instance of the HADR database has the same name on all the nodes.

You must meet the following configuration requirements in a DB2 HADR environment:

- You have set up the database in the HADR environment according to the appropriate DB2 documentation.
- You have installed the database application agent software on each node that will participate in backups or recovery, including the standby nodes.
- You have completed all the post-installation procedures on each node, including the lockbox configuration procedure for all the participating hosts. The hosts can use either a shared lockbox or individually configured lockboxes.
- You have set the `CLIENT` parameter in the DB2 configuration file to the valid hostname of one of the nodes for all the backups and recovery. [Common parameters](#) on page 80 provides details on the parameter.

Note

The same `CLIENT` setting must be used in all the backup and recovery operations for the HADR nodes.

- You have ensured that all the nodes have the same parameter settings.
- You have ensured that the instance of the HADR database has the same name on all the nodes.

You can perform a DB2 HADR recovery on a single node or multiple nodes.

Note

As a DB2 requirement before you start an HADR recovery, you must stop HADR and deactivate the database at the recovery nodes. The DB2 documentation provides details about the required procedures.

Recovery of a single failed node

A DB2 HADR recovery of a single node requires the following steps.

1. Deactivate the database at the failed node, and then stop HADR.

If the failed node is the primary node, the node should switch the role with another node. It is a standby node when the recovery occurs.

2. Recover the failed node as if it is a stand-alone database. Run a rollforward operation without the `complete` option, which leaves the database in a rollforward pending state as required for a standby node.
3. Configure the HADR environment settings, if required.
4. Start HADR on the recovered standby node.

Recovery of all the nodes

A DB2 HADR recovery of all the nodes restores the whole HADR setup to a point-in-time, as in a disaster recovery.

A DB2 HADR recovery of all the nodes requires the following steps.

1. Recover the database to all the HADR nodes as if they are stand-alone databases.
For the standby nodes, run a rollforward operation without the `complete` option, which leaves the database in a rollforward pending state as required for a standby node.
2. Configure the HADR environment settings, if required.
3. Start HADR on all the standby nodes.
4. Start HADR on the primary node.
5. Manually back up the database.

DB2 pureScale requirements for DD Boost operations

You can perform DD Boost backups and restores in a DB2 pureScale environment. In the active-active application cluster, multiple database servers known as member nodes operate on a single data partition.

Note

The database application agent supports delta and incremental backups in a DB2 pureScale environment for the DB2 versions that support these types of backups.

In the DB2 pureScale environment, you run a single `db2 backup db` or `db2 restore db` command on any member to perform the database backup or restore on behalf of all the members. The backup produces one backup image for the entire database, saved to the Data Domain storage.

Each DB2 pureScale member processes its own metadata, generates independent transactions, and maintains its own transaction log files.

You must meet the following configuration requirements in a DB2 pureScale environment:

- You have set up the database in the DB2 pureScale environment according to the appropriate IBM documentation.
- You have installed the database application agent software on each host on which a pureScale member node resides that will participate in backups or recovery.
- You have completed all the post-installation procedures on each member host, including the lockbox configuration procedure for all the participating hosts. The hosts can use either a shared lockbox or individually configured lockboxes.
- If an NFS-shared lockbox is configured, you have followed all the required steps in [Configuring the lockbox in a high-availability environment](#) on page 115.
- You have set the `CLIENT` parameter in the DB2 configuration file to the hostname of one of the member nodes. You must select one pureScale member as the node that will have the backup data stored under its client name. `CLIENT` must be set to the same value for each member node. [Common parameters](#) on page 80 provides details on the parameter.
- You have ensured that all the member nodes have the same parameter settings. You can create a single configuration file in a file system folder that is accessible to all the member hosts. Alternately, you can create an identical configuration file on each member host, with the same file pathname on each host.
- You have updated the database configuration for log archiving and recovery as shown in the following example:

```
db2 update db cfg for sample using logarchmeth1 vendor:/opt/
dpsapps/dbappagent/lib/lib64/libddbboostdb2.so
db2 update db cfg for sample using logarchopt1 @/db2sd/ddp/
ddp.cfg
db2 update db cfg for sample using vendoropt @/db2sd/ddp/ddp.cfg
```

In this example, *sample* is the database alias name. The database configuration parameter `logarchmeth1` is set with the DB2 library that archives logs to the Data Domain storage. The `logarchopt1` and `vendoropt` parameters are set to

use the DB2 configuration file `ddp.cfg`, located in the `/db2sd/ddp` folder in shared storage that is accessible to all the pureScale members.

Performing DD Boost backups in a DB2 pureScale environment

After the database application agent is installed and configured, a user on any active member host can run the `db2 backup` command to perform a DD Boost backup of the database for the entire DB2 pureScale environment. For example:

```
db2 backup db sample online load /opt/dpsapps/dbappagent/lib/lib64/
libddboostdb2.so open n sessions options @/db2sd/ddp/ddp.cfg
```

In this example, the `db2 backup` command performs an online backup of the *sample* database to the Data Domain storage by using the settings in the configuration file `/db2sd/ddp/ddp.cfg`.

Performing DD Boost restores in a DB2 pureScale environment

A user on any active member host can run the `db2 restore`, `db2 rollforward`, and `db2 recover` commands to perform a DD Boost restore of the database and roll forward the database to a point-in-time. For example:

```
db2 restore db sample load /usr/lib/libddboostdb2.so open n
sessions options @/db2sd/ddp/ddp.cfg taken at yyyymmddhhmmss
```

In this example, the `db2 restore` command restores the *sample* database from the Data Domain storage by using the settings in the configuration file `/db2sd/ddp/ddp.cfg`.

The `db2 recover` command combines the functions of the `db2 restore` and `db2 rollforward` commands. You must configure the DB2 database configuration `vendoropt` parameter to run the `db2 recover` command.

Restoring between a DB2 pureScale instance and Enterprise Server Edition

Starting with DB2 10.5, you can restore an offline database backup of a DB2 pureScale instance to DB2 Enterprise Server Edition. You can also restore an offline backup of DB2 Enterprise Server Edition to a DB2 pureScale instance. The IBM DB2 documentation provides details and restrictions for these types of restores.

Restoring a backup from a DB2 pureScale instance to Enterprise Server Edition

Perform the following steps to restore an offline database backup from a DB2 pureScale instance to DB2 Enterprise Server Edition, without rollforward support through the transition. The IBM DB2 documentation provides details about the transition.

Procedure

1. On the DB2 Enterprise Server, configure the lockbox for the Data Domain device that stores the backup performed by the database application agent, as described in [Configuring the lockbox](#) on page 103.
2. In the configuration file on the target DB2 Enterprise Server, ensure that the `CLIENT` parameter and other parameter settings are identical to the settings used during the backup in the DB2 pureScale environment.

3. On the target DB2 Enterprise Server, restore the offline backup image from the DB2 pureScale instance.
4. Complete any required changes to the restored database configuration and the DB2 configuration file according to the Enterprise Server environment:
 - Update the settings of the restored database configuration parameters, such as `logarchopt1`, `logarchopt2`, and `vendoropt`, if required for future backups and restores of the restored database on the Enterprise Server.
The database was restored with the original settings of database configuration parameters used in the pureScale environment, which might require updates for the new environment.
 - Update any required parameter settings in the configuration file for future operations, for example, to specify the correct lockbox pathname, Data Domain system hostname, and device pathname.
 - Delete the `CLIENT` parameter setting from the configuration file on the Enterprise Server because this parameter is only required in the pureScale environment.
5. If required, reconfigure the lockbox for future backups of the restored database, for example, to use a different device host or device pathname.
6. Perform a full offline database backup of the restored database.

Restoring a backup from Enterprise Server Edition to a DB2 pureScale instance

Before you perform a backup that you will restore to a pureScale instance, you can run the `db2checkSD` command on the DB2 Enterprise Server to verify that the source database is ready for restore into a pureScale environment. With the verification complete, you can perform a full offline backup of the source database to prepare for the following restore operation. The IBM DB2 documentation provides details about these operations.

Perform the following steps to restore an offline backup from DB2 Enterprise Server Edition to a DB2 pureScale instance.

Procedure

1. On the DB2 pureScale instance, configure a shared lockbox to use with the Data Domain device that stores the backup performed by the database application agent, as described in [Configuring the lockbox](#) on page 103. As an alternative, create a separate lockbox on each pureScale host by using the same local pathname on each host.
2. In the configuration file on the DB2 pureScale instance, ensure that the `CLIENT` parameter setting is the identical to the `CLIENT` setting used during the backup on the DB2 Enterprise Server.
3. On the DB2 pureScale common member (member 0), restore the offline backup image from the DB2 Enterprise Server.
4. To perform the conversion of the database for use in the pureScale environment, run the following `db2checkSD` command on the DB2 pureScale instance:

```
db2checkSD sample -l /tmp/checksd.log -u user_ID -p password
```

5. Complete any required changes to the restored database configuration and the DB2 configuration file according to the pureScale environment:

- Update the settings of the restored database configuration parameters, such as `logarchopt1`, `logarchopt2`, and `vendoropt`, as required.
 - Update any required parameter settings in the configuration file for future operations, for example, to specify the correct lockbox pathname, Data Domain system hostname, and device pathname.
 - Update the settings of any required parameters in the configuration file in the pureScale environment, such as the `CLIENT` parameter setting.
6. Reconfigure the lockbox for future backups of the restored database, for example, to use a different device host or device pathname.
 7. Perform a full offline database backup of the restored database on the pureScale member 0.

Backups and restores of transaction logs in a DB2 pureScale environment

In a DB2 pureScale environment, each member generates independent transactions and maintains its own set of transaction log files. To enable DB2 rollforward recovery, the transaction logs must be backed up.

You can configure the archiving of DB2 transaction logs with the database application agent by following the instructions in [Configuring DB2 transaction log archiving](#) on page 146. When you set the `logarchmeth1` or `logarchmeth2` configuration parameter to use the DB2 library, the archived transaction log files are automatically saved to the Data Domain storage.

The database application agent performs the log backups based on DB2 database policy settings. The product has no control over when the logs are backed up or how often. The DB2 server starts the backup when a transaction log becomes full.

During a recovery with the `db2 rollforward` or `db2 recover` command, the archived log files from each pureScale member are retrieved from the Data Domain storage. DB2 merges the retrieved logs and recovers the database to the specified point-in-time.

DB2 and the database application agent use a separate storage location under the device path on the Data Domain system for the archived log files of each pureScale member.

During a rollforward recovery, DB2 retrieves the archived logs through the database application agent from the location that corresponds to the member and applies the transaction logs to the database.

Note

For a database backup, DB2 pureScale always passes the node number 0 to the vendor library, no matter which member performs the backup. As a result, the database backup images are all stored under the `NODE0000/DBIMG/` directory.

Deleting DD Boost backups in a DB2 pureScale environment

You can run the `db2 prune` command with the `and delete` option on any active member in a DB2 pureScale environment to delete entries from the recovery history file and delete the associated archived logs. When you set the database configuration parameter `auto_del_rec_obj` to `on`, the database application agent deletes the index entries and backup image save sets from the Data Domain storage if the history file entry is pruned.

You can configure the automatic deletion or pruning of DD Boost backups in the DB2 pureScale environment by following the instructions in [Configuring DB2 backup deletion](#) on page 147. The IBM DB2 documentation provides more details.

DB2 troubleshooting tips for DD Boost operations

[General troubleshooting tips](#) on page 132 provides common troubleshooting information that applies to the database application agent operations with all the supported databases and applications.

The following topics provide troubleshooting information for DB2 operations with the database application agent.

DB2 multistream restore and rollforward might fail on AIX with DFC

On AIX with Data Domain Fibre Channel (DFC), a DB2 multistream restore and rollforward operation might fail or become suspended.

You can resolve this issue by increasing the setting for the number of DD Boost devices on the Data Domain system to which you are connected. In Data Domain System Manager:

1. Select **Data Management > DD Boost > Fibre Channel**.
2. In the **DD Boost Access Groups** area, select the number of DD Boost devices, to a maximum of 64.

DB2 issue with local hostname resolution

The database application agent must resolve the local hostname during its operations. The local hostname resolution uses the system configuration files. For example, on AIX, the resolution is determined by the host entry in the file `/etc/netsv.conf`. On Linux, the file `/etc/nsswitch.conf` contains the required host entry.

Ensure that the host entry setting in the system configuration file is correct for the intended networking configuration of the host. Otherwise, the database application agent operation might fail unexpectedly, with the following information at the end of the debug log:

```
(pid = 6815758) (10/13/16 08:02:43) setMinDDBoostVersion: Exiting.
(pid = 6815758) (10/13/16 08:02:43) checkDedupSettingsInternal:
Exiting.
(pid = 6815758) (10/13/16 08:02:43) nsrdb2_check_init_input: enabling
index optimizations
:
```

As an example on AIX, when the `/etc/hosts` file contains the effective loopback address `::1 ipv6`, the application agent might assume that the hostname resolution should be IPv6-based. As a result, the application agent might fail to resolve the local hostname when the `netstvc.conf` host contains only an IPv4 setting or the host is not configured for IPv6.

The format of the host entry defines the default order of the resolution mechanism:

```
hosts = <value>[, <value>]
```

The following table lists possible value settings for reference. Use one or more of the following values for the `hosts` keyword.

Table 21 Example values for host entry in system configuration file

Hosts keyword value	Description of keyword value
bind	Uses BIND/DNS services for hostname resolution
local	Searches the local file <code>/etc/hosts</code> for hostname resolution
bind4	Uses BIND/DNS services for IPv4 address resolution only
local4	Searches the local file <code>/etc/hosts</code> for IPv4 address resolution only
bind6	Uses BIND/DNS services for IPv6 address resolution only
local6	Searches the local file <code>/etc/hosts</code> for IPv6 address resolution only

The following site provides more details about hostname resolution on AIX:

http://www.ibm.com/support/knowledgecenter/ssw_aix_61/com.ibm.aix.files/netsvc.conf.htm

DB2 issues with `logarchoptn` setting

Due to a DB2 limitation, only the first 30 characters of the `logarchoptn` setting are stored in the DB2 history file. It is recommended that you configure the `logarchoptn` parameter with a value that has fewer than 30 characters, including the @ symbol.

The following issues might occur If you do not follow the recommendation.

1. Pruning of log backups fails due to the `logarchoptn` value

When the pruning of log backups is configured with the `AUTO_DEL_REC_OBJ` parameter, the original value of `logarchoptn` might not be passed correctly from DB2. As a result, the deletion might fail.

The output of the `db2 prune history` and `delete` command might indicate success but the deletion of log backups might have an issue. Errors might appear in the `db2diag.log` file, and the log entries might still remain in both the DB2 database history and backup storage.

In this case, you can perform the following workaround.

- a. Run the following command to update the `comment` field to a value that has fewer than 30 characters:

```
db2 update history file EID entry-eid with comment
'new_location_of_logarchopt1'
```

For example, run the following command:

```
db2 update history EID 10 with comment '@C:\tmp\other.cfg'
```

- b. Rerun the `db2 prune history` and `delete` command.

2. Recover operation fails due to the `logarchoptn` value

In older versions of DB2 such as 9.x, when the `logarchopt1` setting exceeds 30 characters in length, the `db2 recover db` operation might fail with the following type of error:

```
SQL1268N Roll-forward recovery stopped due to error "SQL1042"
while retrieving log file "S0000001.LOG" for database "TEST" on
node "0"
```

In this case, you can perform the following workaround.

a. Reconfigure the `logarchopt1` parameter by running the following command:

```
db2 update db cfg for database_name using logarchopt1
'@pathname/db2_dbdba.cfg'
```

where:

- *database_name* is the name of the database.
- *pathname/db2_dbdba.cfg* is the complete pathname that contains fewer than 30 characters.

b. Rerun the database recovery. For example, run the following command:

```
db2 recover db database_name
```

c. Back up the database after the recovery completes successfully. For example, run the following command:

```
db2 backup db database_name load 'library_path/
libddbboostdb2.so' options '@pathname2/bk_db2.cfg'
```

This step ensures that future recovery operations, to a point-in-time after this backup time, use the new setting for the `logarchopt1` parameter.

3. Rollforward fails on Windows with DB2 9.7 due to the `logarchopt2` value

On Windows with DB2 9.7, due to a DB2 limitation, a DB2 rollforward operation might crash the DB2 instance if the log query or retrieval uses `logarchmeth2` with the following settings:

- The `logarchmeth2` setting specifies the vendor archive method.
- The `logarchopt2` setting exceeds 50 characters in length.

The rollforward operation failure produces the following type of error message:

```
C:\Program Files\IBM\SQLLIB\BIN> db2 rollforward db sample to
end of logs
```

```
SQL1224N The database manager is not able to accept new requests,
has terminated all requests in progress, or has terminated the
specified request because of an error or a forced interrupt.
SQLSTATE=55032
```

In this case, restart the DB2 instance and reconfigure `logarchopt2` to point to a value of 30 characters or less, as required by the IBM standard. Then restart the rollforward operation.

DB2 pruning issues with AUTO_DEL_REC_OBJ

The following DB2 pruning issues might occur with automatic recovery object deletion that is configured through the `AUTO_DEL_REC_OBJ` parameter.

1. Pruning of a multisession backup leaves entries in the DB2 history

The prune operation of a backup performed in multiple sessions removes all the backup pieces from the Data Domain system but might not remove all the entries in the DB2 database history.

To work around this issue and remove the entries in the DB2 database history, run a `db2 prune history with force option` command without the `delete` option.

2. Pruning of DB2 log history after a rollforward produces errors

If you perform a DB2 restore without the `replace history file` option and then perform a rollforward operation, the log history file might contain duplicated entries for the logs archived after the backup and used in the rollforward. A DB2 history pruning operation with `AUTO_DEL_REC_OBJ` set to `ON` might then produce misleading errors. The errors occur when the DB2 software requests the removal of the same archived log backup multiple times due to the duplicated entries in the history.

Note

This issue does not occur when you perform the DB2 restore with the `replace history file` option.

For example, the `db2 prune history` operation produces the following message in the `db2diag.log` file:

```
2015-05-04-13.39.24.676564-240 E684889E515          LEVEL: Info
PID      : 1027                TID   : 47814394505536PROC : db2sysc
0
INSTANCE: db2inst1            NODE  : 000                DB    : TEST
APPHDL  : 0-5718              APPID: *LOCAL.db2inst1.150504173845
AUTHID   : DB2INST1
EDUID   : 2878                EDUNAME: db2agent (TEST) 0
FUNCTION: DB2 UDB, database utilities, sqluhDeletionReport, probe:
381
MESSAGE : ADM8506I  Successfully deleted the following database
logs "3 - 5" in log chain "1".

2015-05-04-13.39.24.676841-240 E685405E502          LEVEL: Error
PID      : 1027                TID   : 47814394505536PROC : db2sysc
0
INSTANCE: db2inst1            NODE  : 000                DB    : TEST
APPHDL  : 0-5718              APPID: *LOCAL.db2inst1.150504173845
AUTHID   : DB2INST1
EDUID   : 2878                EDUNAME: db2agent (TEST) 0
FUNCTION: DB2 UDB, database utilities, sqluhDeletionReport, probe:
387
MESSAGE : ADM8509N  Unable to delete the database logs "3 - 4" in
log chain "1".
```

To work around this issue, perform the following steps.

- a. To verify that the log entries are removed, check the `db2diag.log` file and the backup index.
- b. Clean up the DB2 history file by running a `db2 prune history...with force option` command without the `and delete option`.

DB2 issues due to incorrect log retrieval

The following DB2 issues are caused by the incorrect retrieval of the backup logs.

1. Log retrieval issue for re-created database

According to IBM DB2 documentation, the destination that contains the archived log backups, known as the archive log path, must contain only the log files that belong to the current database.

If the archive log path was previously used for a database of the same name, which for example was dropped and re-created, the old and new backups should not be stored together. Mixing the new backups with the previous log backups can cause issues during a log pruning and retrieval, as in a rollforward operation or an online backup with the `INCLUDE LOGS` option.

The log retrieval operation might fail with the following error messages in the `db2diag.log` file:

```
Database ID does not match. Extent does not belong to this
database.
Database ID does not match. Extent probably for another database.
```

The dropped database and the re-created database might share the same log sequences and chains. In such a case, the archived log backup pruning for one database might accidentally remove the records of the other database.

To prevent this issue, ensure that you clean up the old log backups if they are no longer needed. If you use the `ddbmadmin` command for the clean-up, refer to the `ddbmadmin` information in [Configuring the display and deletion of save set information](#) on page 117 for the proper deletion options. When the old backups must be retained, it is recommended that you use a new device for the backups of the current database by setting the `DEVICE_PATH` parameter to a new location on the Data Domain system.

2. DB2 issue in rollforward with archived log backup

A DB2 archived log backup might take place in the same timeframe as a log restore during a rollforward operation. For example, to complete a rollforward, uncommitted transactions are rolled back. The rollback action invokes log archiving. In this case, both the log restore and log archiving use the same configuration file.

A redirected recovery to a different client should use two client parameters that refer to different hosts: the source client that has the log to roll forward and the target client that has the log to be backed up. If the recovery uses only one client parameter that points to the source client, the concurrent log backup is saved incorrectly under the source client name. Then a rollforward operation (when needed) of the source database might fail with the following error when the wrong log is retrieved:

```
Database ID does not match. Extent probably for another database.
```

The `SOURCE_CLIENT` parameter prevents this issue by pointing to the source client that has the log backups used to roll forward. `CLIENT` is used to point to the target host under which the archived log backups of the rollback will be stored. If `SOURCE_CLIENT` is not specified, `CLIENT` is used for both the archived log backup and restore.

Database backup might fail when run concurrently with backups of a high number of archived logs

If a database backup, especially one with multiple sessions, starts when a high number of archived logs (more than 300 logs) are ready to be backed up, the database backup might fail.

In this case, you can restart the database backup later when there are fewer logs to be backed up. You can estimate the approximate number of logs ready for backup by comparing the latest log sequence number that was backed up and the next log sequence to be archived.

To estimate the next log sequence number to be archived, query the next active log sequence number from the database configuration. For example:

```
C:\Program Files> db2 get db cfg for testdb1 | grep -i log
```

```
First active log file = S0000559.LOG
```

To estimate the next archived log to be backed up, browse the `db2diag.log` file. For example, the file contains the following information:

```
FUNCTION: DB2 UDB, data protection services, sqlpgArchiveLogFile,
probe:3180
DATA #1 : <preformatted>
Completed archive for log file S0000347.LOG to VENDOR chain 1 from C:
\DB2_01\NODE0000\SQL00001\LOGSTREAM0000\.
```

From these examples, you can calculate the number of archived logs ready to be backed up as: $559 - 347 = 212$ logs.

DB2 operation might generate empty debug logs on Windows

On Windows, certain DB2 operations with the database application agent might generate debug log files with a size of zero bytes.

You can ignore any zero-byte debug logs.

CHAPTER 6

ProtectPoint Operations on DB2 Systems

This chapter includes the following topics:

- [Overview of ProtectPoint operations in a DB2 environment](#)..... 182
- [Configuration of ProtectPoint operations in a DB2 environment](#)..... 183
- [Performing ProtectPoint backups and recovery with the DB2 CLP](#)..... 187
- [Managing and deleting ProtectPoint DB2 backups](#)..... 189
- [Preparing for DB2 disaster recovery](#)..... 190
- [DB2 DPF requirements for ProtectPoint operations](#)..... 192
- [DB2 HADR requirements for ProtectPoint operations](#)..... 197
- [DB2 pureScale requirements for ProtectPoint operations](#)..... 199
- [DB2 troubleshooting tips for ProtectPoint operations](#)..... 210

Overview of ProtectPoint operations in a DB2 environment

The database application agent is integrated with the DB2 interfaces for third-party media management vendors to enable ProtectPoint DB2 backups and restores and the management and deletion of the backups. The IBM DB2 software provides the Advanced Copy Services (ACS) feature that enables ProtectPoint operations on DB2 databases.

You can perform a ProtectPoint backup, restore, query, or deletion with the product on a DB2 database server by running one of the supported DB2 tools:

- DB2 Command Line Processor (CLP) with the commands `db2 backup` and `db2 restore`
- DB2 query and deletion tool `db2acsutil`

You can use these tools in cooperation with the database application agent to perform the following ProtectPoint DB2 operations:

- Online and offline backups
- Full backups of a whole database
- Recovery of a database to the current time or a specific point-in-time
- Recovery to the original location or a different host (same database and instance)
- Backup and recovery of databases only
- Backup query and deletion

Due to DB2 snapshot limitations, the database application agent supports only the ProtectPoint backup and restore of a whole DB2 database. In a ProtectPoint restore, the database and instance must have the same name as in the backup.

The database application agent does not support the ProtectPoint backup and restore of selected DB2 tablespaces, archived logs, or other files. The database application agent also does not support an incremental ProtectPoint backup of DB2 data.

The product maintains online backup indexes on the Data Domain system. During backups, the product creates backup entries in the online indexes, which provide the information required to restore the backed-up data.

The troubleshooting section at the end of this chapter provides details about limitations in the ProtectPoint operations with the database application agent in a DB2 environment.

ProtectPoint DB2 backup processes

A ProtectPoint DB2 database backup includes the following process interactions.

1. The database administrator initiates the backup by running the `db2 backup use snapshot library` command, the IBM Data Studio GUI, or the DB2 Control Center GUI.
2. The DB2 software loads the shared library used by the database application agent, and then invokes the ACS API for the backup tasks.

Note

IBM documentation refers to the library as a vendor library.

3. The database application agent reads the configuration file, and then initializes the connection with the Data Domain system.
4. The DB2 software sends information to the shared library about the database paths to back up, and then the library passes the information to the snapshot agent.
5. The backup workflow proceeds as described in the topic about the ProtectPoint backup workflow or the ProtectPoint with RecoverPoint backup workflow in Chapter 1.

ProtectPoint DB2 restore processes

A ProtectPoint DB2 database restore includes the following process interactions.

1. The database administrator initiates the restore by running the `db2 restore use snapshot library` command, the IBM Data Studio GUI, or the DB2 Control Center GUI.
2. The DB2 software loads the shared library used by the database application agent, and then invokes the ACS API for the restore tasks.
3. The database application agent reads the configuration file, and then initializes the connection with the Data Domain system.
4. The DB2 software requests the backup from the database application agent.
5. The restore workflow proceeds as described in the topic about the ProtectPoint restore workflow or the ProtectPoint with RecoverPoint restore workflow in Chapter 1.

DB2 backups of transaction logs

Ensure that the DB2 archived transaction logs are backed up, for example, by configuring the automatic backup of the transaction logs. The automatic log backup uses the DD Boost workflow, not the ProtectPoint workflow. [DB2 backups of transaction logs](#) on page 143 provides more details.

Configuration of ProtectPoint operations in a DB2 environment

Ensure that the VMAX, XtremIO, RecoverPoint, and Data Domain configurations have been completed according to the ProtectPoint documentation. The required storage resources must be configured and provisioned properly to enable ProtectPoint operations.

Complete the following tasks to enable ProtectPoint operations:

- Ensure that the `ddbcmd` program is started from the `/opt/dpsapps/dbappagent/bin` directory.
- For ProtectPoint for VMAX operations only, ensure that the supported VMAX Solutions Enabler version is installed and configured in local mode on each production host. The online software compatibility guide at <http://compatibilityguide.emc.com:8080/CompGuideApp/> describes the supported versions.

The Solutions Enabler database must be up-to-date on any host where a backup or recovery might run. To update the Solutions Enabler database, run the `symcfg discover` command. The Solutions Enabler documentation provides details.

Ensure that the required gatekeepers are also configured as described in the *ProtectPoint Version 4.0 Primary and Protection Storage Configuration Guide*. Solutions Enabler uses the small gatekeeper devices for communication with the VMAX storage array.

[Database application agent ProtectPoint operations with Data Domain usage limits](#) on page 44 provides general guidelines on the Data Domain usage limit settings for ProtectPoint operations.

To enable the ProtectPoint operations in a DB2 environment, you must complete the required configurations of the database application agent. The following topics provide the product configuration details.

The troubleshooting section at the end of this chapter provides details about limitations in the ProtectPoint operations with the database application agent in a DB2 environment.

Configuring the DB2 parameters for ProtectPoint operations

You must set the required parameters for ProtectPoint DB2 operations in the configuration file used by the database application agent. For example, the configuration file named `db2_ddbda.cfg` contains the following parameter settings for ProtectPoint operations:

```
DDBOOST_USER=qa_ost
DDVDISK_USER=vdisk
DEVICE_HOST=bu-dbe-890.lss.emc.com
DEVICE_PATH=/bu-star1_db2
DEVICE_POOL=IT_data_pool
```

[Setting up the configuration file](#) on page 78 describes the common parameters, ProtectPoint parameters, and how to set the parameters in the configuration file. Other topics in [Product Configuration](#) on page 77 describe the parameters and requirements for the restores of replicated backups and rollback restores.

Ensure that the configuration file also includes `DB2_ACS_LAYOUT_CHECK` parameter if required. [Enforcing the DB2 ACS best practice on log directory layout](#) on page 184 provides details.

After the configuration file is set up, ensure that the required lockbox procedures have been performed as described in [Configuring the lockbox](#) on page 103.

Enforcing the DB2 ACS best practice on log directory layout

To exclude logs in a ProtectPoint backup or to exclude logs in the restore of a ProtectPoint backup that includes logs, the log directories must reside on different disk volumes than the other database paths.

The DB2 ACS best practice recommends to use a dedicated volume group for log paths, with the log paths contained in a snapshot volume that is separate from the database directory and database containers.

Refer to the following IBM documentation:

<http://pic.dhe.ibm.com/infocenter/db2luw/v10r1/topic/com.ibm.db2.luw.admin.ha.doc/doc/c0053158.html>

For ProtectPoint with RecoverPoint, as the snapshots are performed at the consistency group level, the grouping of the database objects must also be performed at that level. To back up or restore a database without the logs, the database log

directories must reside on disks that belong to a different consistency group than the disks of the other database paths.

The `DB2_ACS_LAYOUT_CHECK` parameter specifies whether to enforce the DB2 ACS best practice on the log directory layout of the database during a ProtectPoint backup. Set the parameter in the [GENERAL] section of the configuration file. The following table provides details.

Table 22 DB2 parameter for ProtectPoint operations

Parameter: DB2_ACS_LAYOUT_CHECK

Section: [GENERAL]

Specifies whether to enforce the DB2 ACS best practice on the log directory layout of the database during a ProtectPoint backup.

The best practice requires a dedicated volume group for log paths, with the log paths contained in a different file system volume than the database directory and database containers.

For a ProtectPoint with RecoverPoint backup, the layout enforcement is validated at the consistency group level.

Optional for a ProtectPoint backup.

Valid values:

- TRUE (default) = Enforce the DB2 ACS best practice on the log directory layout. Backups with the `exclude logs` option fail if the log paths are not in a separate file system volume. For ProtectPoint with RecoverPoint backups, the backups with the `exclude logs` option fail if the logs are not in a separate volume that belongs to a separate consistency group.
- FALSE = Do not enforce the DB2 ACS best practice on the log directory layout.

Configuring DB2 transaction log archiving

When you configure the archiving of DB2 transaction logs, the database application agent performs the log backups based on DB2 database policy settings. The product has no control over when the logs are backed up or how often. The DB2 server starts the backup when a transaction log becomes full.

[Configuring DB2 transaction log archiving](#) on page 146 provides details about configuring the automatic backup of DB2 transaction logs, which uses the DD Boost workflow and not the ProtectPoint workflow.

Preparing for DB2 redirected rollback restores of ProtectPoint for VMAX backups

The database application agent 4.0 introduced support for redirected rollback restores of ProtectPoint for VMAX backups to alternate LUNs on an alternate host.

[Configuring rollback restores of ProtectPoint backups](#) on page 97 describes the basic requirements for a redirected rollback restore of a ProtectPoint for VMAX backup.

For a DB2 redirected rollback restore of a ProtectPoint for VMAX backup, ensure that you meeting the following additional requirements:

- You use a consistent backup for the restore. For a recoverable database, you restore the backup by using the `without rolling forward` option.

- You perform the rollback restore of an entire database, including the logs.

Note

The database does not need to exist before the rollback restore.

- The database manager configuration parameter, `DFTDBPATH`, is preferably set to the database path value in the backup. This setting helps with any manual clean-up that might be required after a snapshot restore failure.

Preparing for DB2 ProtectPoint with RecoverPoint backups and rollback restores

With RecoverPoint pre-5.0, the database application agent performs a rollback restore of a DB2 ProtectPoint with RecoverPoint backup at the consistency group level. If the RecoverPoint consistency group being restored contains multiple LUNs, then all those LUNs are overwritten and inaccessible during the rollback restore. Specific requirements apply to the DB2 ProtectPoint with RecoverPoint backups and rollback restores.

Ensure that you follow the requirements and recommendations in [Configuring rollback restores of ProtectPoint backups](#) on page 97.

DB2 supports the backups and restores that exclude log objects when the log objects are classified as in a different group than other database objects. The database application agent performs the grouping during the backup time:

- With the database application agent 2.5, the backup objects grouping for the ProtectPoint with RecoverPoint backup is per file system volume group, in favor of a point-in-time restore.
- With the database application agent 3.0 or later, to support rollback restore with a limitation in retrieving the RecoverPoint version, the backup objects grouping is per consistency group for all RecoverPoint versions.

Ensure that you meet the following requirements for a DB2 ProtectPoint with RecoverPoint backup or rollback restore with the `exclude logs` option:

- The database log LUNs are in a separate dedicated consistency group from the database LUNs.
- The database log LUNs are in a separate dedicated volume group from the database LUNs.

For a rollback restore of a release 2.5 backup and with RecoverPoint pre-5.0, ensure that all the database objects in the same consistency group are included in the restore command. If any LUNs in the backed-up consistency group contain objects that were not included in the backup command, then ensure that you manually unmount those LUNs before the rollback restore and then manually mount the LUNs back after the restore.

Note

- With RecoverPoint pre-5.0, a DB2 ProtectPoint with RecoverPoint backup and rollback restore always occurs at the consistency group level, regardless of which objects are included in the backup command. The backup objects grouping is per consistency group. As a best practice for a DB2 ProtectPoint with RecoverPoint rollback restore, when you perform the backup or rollback restore, do not exclude the logs or any database files that are part of the RecoverPoint consistency group being backed up or restored.
 - Before a rollback restore of a release 2.5 backup or RecoverPoint pre-5.0 backup with the `exclude logs` option (where either the backup or restore uses the option), ensure that the `psrollback.res` file does not list the DB2 log directories.
-

Performing ProtectPoint backups and recovery with the DB2 CLP

You can run the DB2 CLP to perform ProtectPoint backups, restores, and recovery with the database application agent. The DB2 documentation provides details about the DB2 CLP commands and options.

Performing ProtectPoint backups with the DB2 CLP

You can perform a ProtectPoint DB2 backup after you have completed the backup configurations in [Configuration of ProtectPoint operations in a DB2 environment](#) on page 183.

You can run the appropriate `db2 backup use snapshot library` command to perform a ProtectPoint DB2 backup.

DB2 snapshot backups do not support incremental level, tablespace level, or multisession backups. The DB2 documentation provides details about all the unsupported options for snapshot backups.

The default type of ProtectPoint backup is an offline full database backup of all the paths that comprise the database, including all the containers, the local volume directory, the database path, and the primary log and mirror log paths. A ProtectPoint backup uses the `include logs` option by default unless you specify the `exclude logs` option. You can perform an online backup by using the `online` option. The DB2 documentation provides more details.

For example on UNIX, run the following command:

```
db2 backup db sample online use snapshot library /opt/dpsapps/
dbappagent/lib/lib64/libddbostdb2.so options '@pathname/
db2_dbdba.cfg'
```

where:

- *sample* is the name of the database to back up.

- `online` specifies to perform an online backup. The default backup type is an offline backup.
- `pathname/db2_ddbda.cfg` is the pathname of the DB2 configuration file as described in [Configuring the DB2 parameters for ProtectPoint operations](#) on page 184.

Performing ProtectPoint restores with the DB2 CLP

You can run the `db2 restore use snapshot library` command with the appropriate options to perform a ProtectPoint DB2 restore to either the same DB2 application host or a different host.

A ProtectPoint DB2 restore can restore a ProtectPoint backup to the original database.

If you are recovering the data to a point-in-time, note the timestamp of the backup to restore.

DB2 snapshot restores do not support incremental level, tablespace level, or other types of restores. The DB2 documentation provides details about all the unsupported options for snapshot restores.

Before you perform any restores, ensure that you meet the following requirements:

- The numeric user ID (UID) and group ID (GID) of the target database/instance owner matches the original UID and GID captured during the ProtectPoint backup.

A ProtectPoint backup is associated with the original database/instance owner. During the restore of the ProtectPoint backup, the UID and GID of the target database/instance owner must match the original UID and GID. Otherwise, the restore fails because the database/instance owner does not have the permission to access the database objects after they are restored.

- All the file system mount points in the backup are re-created with the proper ownership and permissions.

A ProtectPoint restore does not restore the ownership and permissions of the mount points and the file system directories above them.

- If the database contains symbolic links, then the symbolic links are re-created before you perform a restore.

A ProtectPoint backup does not back up symbolic links.

Perform a ProtectPoint restore by running the `db2 restore use snapshot library` command. For example, run the following command on UNIX:

```
db2 restore db sample use snapshot library /opt/dpsapps/
dbappagent/lib/lib64/libddbboostdb2.so options '@pathname/
db2_ddbda.cfg' taken at yyyymmddhhmmss logtarget include force
```

where:

- *sample* is the name of the database to be restored.
- `pathname/db2_ddbda.cfg` is the pathname of the DB2 configuration file.
- *yyyymmddhhmmss* is the timestamp of the backup to restore.

Skip the `taken at` parameter if you restore only the most recent backup of the database.

If the timestamp of the backup is unknown, you can run the `db2acsutil query` command to find the timestamp. For example:

```
db2acsutil LOAD /opt/dpsapps/dbappagent/lib/lib64/libddboostdb2.so
options '@pathname/db2_ddbda.cfg' query snapshot db sample
```

where:

- `pathname/db2_ddbda.cfg` is the full pathname of the configuration file.
- `sample` is the name of the database to be restored.

[Querying ProtectPoint DB2 backups](#) on page 190 provides more details.

Performing DB2 recovery with the DB2 CLP

You can run the `db2 rollforward` command to apply the transaction logs that are stored on the Data Domain system to recover a DB2 database to either the current time or a specific point-in-time. The rollforward operation uses the DD Boost workflow, not the ProtectPoint workflow. [Performing DB2 recovery with the `db2 rollforward` command](#) on page 157 provides details.

Note

To use rollforward recovery, the database application agent must have backed up the transaction logs. [DB2 backups of transaction logs](#) on page 183 provides details.

The `db2 recover` command does not apply to ProtectPoint backups.

Managing and deleting ProtectPoint DB2 backups

You can use the `db2acsutil` utility to manage the ProtectPoint DB2 backups. You can run the `db2acsutil` command to perform the following operations:

- List the available ProtectPoint DB2 backups that you can use to restore the DB2 database.
- Delete ProtectPoint DB2 backups and release the associated resources.

Note

You cannot use the `db2acsutil` utility to monitor the status of ProtectPoint DB2 backups created with the database application agent.

The IBM DB2 documentation provides details on the `db2acsutil` command and options.

Each ProtectPoint backup is also recorded in the DB2 recovery history, the same as other types of DB2 backups. However, the following operations are not applicable to DB2 snapshot backups:

- Manual pruning of database object backups with the `db2 prune history` and `delete` command.
- Automatic deletion of recovery objects through a configuration with the DB2 parameters `num_db_backups`, `rec_hist_retentn`, and `auto_del_rec_obj`.

Querying ProtectPoint DB2 backups

You can run the `db2acsutil query` command to generate a list of the available ProtectPoint DB2 backups retained in the repository. You can run the command with the `db`, `instance`, or `taken at DB2_timestamp` option. The database application agent does not support the combination of any of these command options.

The following examples show the queries of ProtectPoint backups:

```
db2acsutil LOAD /opt/dpsapps/dbappagent/lib/lib64/libddboostdb2.so
options '@pathname/db2_ddbda.cfg' query snapshot db SAMPLE
db2acsutil LOAD /opt/dpsapps/dbappagent/lib/lib64/libddboostdb2.so
options '@pathname/db2_ddbda.cfg' query snapshot instance db2inst1
db2acsutil LOAD /opt/dpsapps/dbappagent/lib/lib64/libddboostdb2.so
options '@pathname/db2_ddbda.cfg' query snapshot taken at
20150321121212
db2acsutil LOAD /opt/dpsapps/dbappagent/lib/lib64/libddboostdb2.so
options '@pathname/db2_ddbda.cfg' query snapshot older than 7 days
ago instance db2inst1
```

where `pathname/db2_ddbda.cfg` is the full pathname of the configuration file used by the database application agent. The DB2 documentation provides details on the command and options.

Deleting ProtectPoint DB2 backups

You can run the `db2acsutil delete` command with the `taken at yyyyymmddhhmmss` option to delete a ProtectPoint DB2 backup created with the database application agent. The database application agent does not support any other options with the `db2acsutil delete` command.

Before you confirm that a deletion should proceed, check for the name of the instance and database in the `db2acsutil delete` command output. The software deletes the backup entries from the backup indexes.

The `db2acsutil delete` operation does not involve the DB2 recovery history. If required, you can manually prune the DB2 recovery history to keep it synchronized with the `db2acsutil` utility operations.

The following example shows the deletion of a ProtectPoint DB2 backup:

```
db2acsutil LOAD /opt/dpsapps/dbappagent/lib/lib64/libddboostdb2.so
options '@pathname/db2_ddbda.cfg' delete snapshot db SAMPLE taken
at 20150321121212
```

where `pathname/db2_ddbda.cfg` is the full pathname of the configuration file. The DB2 documentation provides details on the command and options.

Preparing for DB2 disaster recovery

For a comprehensive disaster recovery plan, you must ensure that you can reconstruct the computing environment and all the DB2 server files associated with maintaining data on the application host.

Use the following guidelines to prepare for a disaster recovery of the DB2 server host:

- Maintain accurate and complete records of the network and system configurations. Keep all the original software media and the following items in a safe location:
 - Original operating system media and patches
 - Device drivers and device names
 - File system configuration
 - IP addresses and hostnames
- Ensure that you have a current full backup of the database and all the archived logs required for a rollforward operation.
- Save a copy of the configuration file used for the DB2 backups of the database and archived logs.
- Confirm that the parameter setting of the corresponding options file, such as `logarchopt1`, from the source database image is valid on the destination host.

Plan to perform the following tasks during a disaster recovery.

Procedure

1. Set up the configuration file to be used during the recovery, including the following parameter settings:

- `CLIENT=source_client_hostname`
- `DB2INSTANCE=source_database_instance`
- `SOURCE_DBNAME=source_database_name`

The `DB2INSTANCE` and `SOURCE_DBNAME` parameters are only required for rollforward operations that use the DD Boost workflow, not for restores of ProtectPoint database backups.

[Configuring the DB2 parameters for ProtectPoint operations](#) on page 184 provides details.

2. Re-create the lockbox on the database host. [Configuring the lockbox](#) on page 103 provides details about the lockbox.
3. Ensure that all the required database and log backup images are available. You can run the `db2acsutil query` command to obtain a list of available ProtectPoint backups. [Querying ProtectPoint DB2 backups](#) on page 190 provides details on this command.
4. Before you perform any restores, ensure that you meet the ProtectPoint restore requirements described in [Performing ProtectPoint restores with the DB2 CLP](#) on page 188.
5. Run the appropriate `db2 restore` command to restore the database. For example:

```
db2 restore db database_name use snapshot library /opt/
dpsapps/dbappagent/lib/lib64/libddbboostdb2.so options
'@pathname/db2_ddbda.cfg' taken at yyyyymmddhhmmss
```

6. If required, perform a rollforward operation on the restored database, to the end of the logs or a point-in-time:

```
db2 rollforward database database_name
```

DB2 DPF requirements for ProtectPoint operations

In this release, for a DB2 Database Partitioning Feature (DPF) database, you can perform ProtectPoint for VMAX backups and restores. In the DPF environment, a database is partitioned onto multiple nodes or partitions, either on the same host or on multiple hosts. Each physical host can have multiple logical partitions.

Each database partition is part of the database, and consists of its own data, indexes, configuration files, and transaction logs. The data of each partition is managed by the partition itself. In a backup or restore, each partition operates separately and produces its own debug log file.

The database data can be distributed across a partition through the creation of the database tablespace and containers, and through the database partition group. The DB2 documentation provides details about the configuration of database partition containers and groups.

The features and restrictions of ProtectPoint operations that generally apply to DB2 systems also apply in a DPF environment. The following topics provide details of the specific requirements for ProtectPoint operations on DB2 DPF systems.

File system requirements for ProtectPoint operations in a DPF environment

You must meet the following file system requirements in a ProtectPoint DPF environment:

- The database file system is local to the partition host. The DPF databases are not created in an NFS-mounted directory. The devices for the mount points of a partition's file systems are local to the partition host.

Note

ProtectPoint workflows do not support DPF databases with a file system that resides on remote devices.

- The logical partition devices and file systems are dedicated to the partition. The databases must not reside on the same device as the instance owner home or other file system of the host.

Note

The DPF database has a database system directory structure under the common database home directory, such as `<database_home>/<instance_name>/NODEnnnn`, where `NODEnnnn` is the node number of the corresponding partition.

- It is recommended that log devices reside in a separate snapshot unit, such as a volume group or disk, apart from other database component devices. For an exclude log backup and restore, the separate snapshot unit for the log objects is a requirement.

You can update the database configuration (such as `newlogpath`) for a partition by using the option `member <node_ID>` with the following command:

```
db2 update db cfg for db <database_name> member <node_ID> using
<parameter_name> <value>
```

As an example of a DPF database directory structure, the database `sample` is created under the instance `dpf10`, resides in the database home directory `/dbhome`, and has

four partitions on two physical hosts. Each host has two logical nodes. `DFTDBPATH` (the default database path) is set to the absolute pathname `/dbhome`. In this example, the default system directory structure of the partitions is as follows:

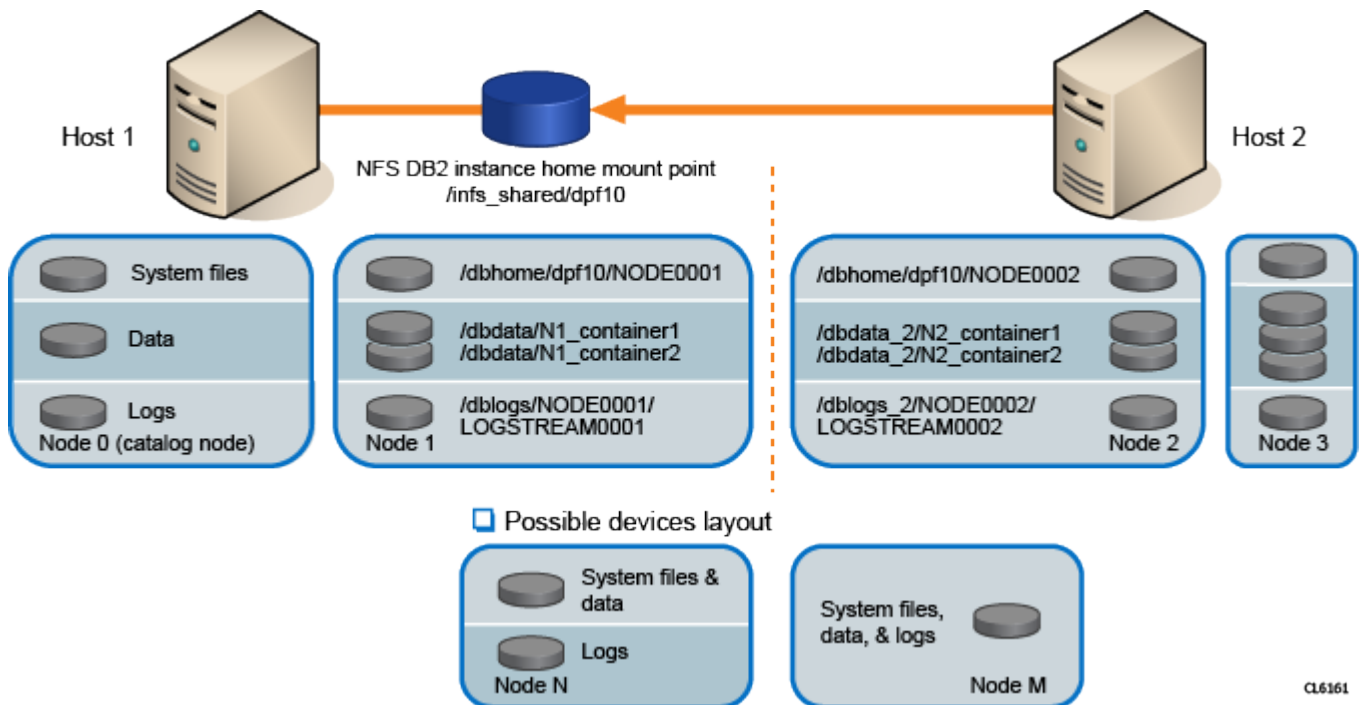
- Host 1:
 - `/dbhome/dpf10/NODE0000`
 - `/dbhome/dpf10/NODE0001`
- Host 2:
 - `/dbhome/dpf10/NODE0002`
 - `/dbhome/dpf10/NODE0003`

The default log directory of the partitions is as follows:

- Host 1:
 - `/dbhome/dpf10/NODE0000/SQL00001/LOGSTREAM0000/`
 - `/dbhome/dpf10/NODE0001/SQL00001/LOGSTREAM0001/`
- Host 2:
 - `/dbhome/dpf10/NODE0002/SQL00001/LOGSTREAM0002/`
 - `/dbhome/dpf10/NODE0003/SQL00001/LOGSTREAM0003`

To perform a ProtectPoint backup of the database including the logs, the subdirectories `NODE000n` and `LOGSTREAM000n` must reside on the local ProtectPoint devices. The following figure illustrates the database file system layout in a DB2 DPF environment.

Figure 10 Database file system layout in a ProtectPoint DPF environment



CL6161

Configuration requirements for ProtectPoint operations in a DPF environment

You must meet the following configuration requirements in a ProtectPoint DPF environment:

- You have properly set up the database DPF environment according to the DB2 documentation.
- You have installed and configured the database application agent software on all the physical hosts of the DPF database. The software is configured properly for all the logical nodes.
- You have ensured that all the nodes have the same parameter settings. You have created a single configuration file in a shared system folder that is accessible to all the nodes. Alternately, you have created an identical configuration file on each partition host, with the same file pathname on each host.

You have set the `CLIENT` parameter in the DB2 configuration file for all the nodes to the same hostname, which is typically the logical node hostname. [Common parameters](#) on page 80 provides details on the parameter.

- The partition hosts use either a shared lockbox or individually configured lockboxes.
- You have set the database manager configuration parameter, `DFTDBPATH` (the default database path), to a proper value. The default value of the parameter is a location under the DB2 instance home directory.

Note

During a restore, `DFTDBPATH` must be set to the value in the backup image.

- The database resides on a dedicated file system layout as described in [File system requirements for ProtectPoint operations in a DPF environment](#) on page 192.

Performing ProtectPoint backups in a DPF environment

DB2 backs up a DPF database per partition. You can perform ProtectPoint backups of a DPF database for a specified list of partitions, with each backup being a full backup of a whole DPF database partition. The DPF archive log backups are performed through the DD Boost workflow, not the ProtectPoint workflow.

You can perform the ProtectPoint backup of a DPF database through either of the following methods:

- Run a Single System View (SSV) backup of multiple partitions from the catalog node as a concurrent backup of specified partitions, by running the `db2 backup` command with the `dbpartitionnums` option.

The status of each partition backup affects the overall status of the SSV backup. When one of the partition backups fails, the whole SSV backup is reported as failed.

For example:

- The following commands are examples of a SSV backup of all partitions:

```
db2 backup db <database_name> ... on all dbpartitionnums ...
```

```
db2 "backup db DB01 on all dbpartitionnums use snapshot
library /opt/dpsapps/dbappagent/lib/lib64/libddboostdb2.so
options @/home/cfg/dd/pp.cfg"
db2 "backup db DB01 on all dbpartitionnums use snapshot
library /opt/dpsapps/dbappagent/lib/lib64/libddboostdb2.so
options @/home/cfg/dd/pp.cfg exclude logs"
```

- The following command is an example of a SSV backup of specific partitions:

```
db2 backup db <database_name> ... on dbpartitionnums (n1,n2,...)

db2 "backup db DB01 on dbpartitionnums (0,1,2) use snapshot
library /opt/dpsapps/dbappagent/lib/lib64/libddboostdb2.so
options @/home/cfg/dd/pp.cfg"
```

- Run individual backups of specified partitions by running the `db2_all` command in sequential or concurrent mode. The status of each partition backup is considered separately. Each backup has its own timestamp, and the failure of a partition backup does not affect any other partition backup.

For example, the following `db2_all` commands perform the backups of specified partitions:

- Sequentially:

```
db2_all "db2 backup db DB01 use snapshot library /opt/dpsapps/
dbappagent/lib/lib64/libddboostdb2.so options @/home/cfg/dd/
pp.cfg"
db2_all "db2 backup db DB01 use snapshot library /opt/dpsapps/
dbappagent/lib/lib64/libddboostdb2.so options @/home/cfg/dd/
pp.cfg"
db2_all "<<<+0< db2 backup db DB01 use snapshot library /opt/
dpsapps/dbappagent/lib/lib64/libddboostdb2.so options @/
home/cfg/dd/pp.cfg"
```

- Concurrently in the background:

```
db2_all "<<<-0<; db2 backup db DB01 use snapshot library /opt/
dpsapps/dbappagent/lib/lib64/libddboostdb2.so options @/
home/cfg/dd/pp.cfg"
```

Performing ProtectPoint restores in a DPF environment

You restore the ProtectPoint backups in a DPF environment by restoring per partition. For example, you can run `db2 restore` within the `db2_all` command to restore one or more partitions.

If the restore includes the catalog partition, then you can restore the catalog partition first, followed by the other partitions. The IBM DB2 documentation provides details about restores in a DPF database environment.

Note

DB2 does not support the `db2 recover` command for the recovery of a snapshot backup.

You can also perform rollforward operations from the catalog node for a single partition or multiple partitions. The rollforward operations are performed through the DD Boost workflow.

For a recovery in a ProtectPoint DPF environment, ensure that you have set the database manager configuration parameter, `DFTDBPATH` (the default database path),

to the database path value in the backup image. This setting helps with the recovery from a failed backup and the clean-up of database partitions in the case of a failed snapshot restore.

When you run the restores of ProtectPoint DPF backups with the `db2_all` command, use the `replace existing` option when applicable because the `db2_all` command does not support interactive input.

For example, the following command performs the restore of partition 0:

```
db2_all "<<+0< db2 restore db DB01 use snapshot library /opt/
dpsapps/dbappagent/lib/lib64/libddbboostdb2.so options @/home/cfg/dd/
pp.cfg taken at 20161007184634 logtarget include force replace
existing"
```

The following commands perform the restore of all partitions except partition 0:

- Sequentially:

```
db2_all "<<-0< db2 restore db DB01 use snapshot library /opt/
dpsapps/dbappagent/lib/lib64/libddbboostdb2.so options @/
home/cfg/dd/pp.cfg taken at 20161007184634 logtarget include
force replace existing"
```

- Concurrently in the background:

```
db2_all "<<-0<; db2 restore db DB01 use snapshot library /opt/
dpsapps/dbappagent/lib/lib64/libddbboostdb2.so options @/
home/cfg/dd/pp.cfg taken at 20161007184634 logtarget include
force replace existing"
```

If a partition restore fails, the restored partition might be damaged. To recover from the failure, try to re-create the partition and restart the restore:

- Set `DB2NODE` to the partition number in the environment, and run the command `db2 terminate` for the change to take effect.
- Run the command `db2 create database <database_name> at dbpartitionnum.`
- If the partition cannot be re-created, manual clean-up might be required. The restore might need to be restarted for all the nodes.

Performing query and deletion operations in a DPF environment

After you perform ProtectPoint backups of DPF databases, you can also perform query and deletion operations for the ProtectPoint DPF backups by using the DB2 query and deletion tool `db2acsutil`.

You can run the `db2acsutil` command to perform query and deletion operations from any partition host for the ProtectPoint DPF backups. The entry for each partition is listed separately:

- You can run a query by specifying the database name, instance name, partition number, and DB2 timestamp. For example, the following command output shows

the query results of the SSV backup of all partitions of a database with a timestamp:

```
db2acsutil load /opt/dpsapps/dbappagent/lib/lib64/
libddbboostdb2.so options @/space1/cfg/pp.cfg query snapshot db
ps03 instance dpf10 taken at 20161007184634
```

Instance	Database	Part	Image Time	Host	First Log
dpf10	PS03	0	20161007184634		3
dpf10	PS03	1	20161007184634		3
dpf10	PS03	2	20161007184634		3
dpf10	PS03	3	20161007184634		3

- You can run a deletion by specifying the DB2 timestamp with the `taken at` option. You can also filter the command output more by specifying the database name, instance name, and `dbpartitionnum <n>` option. For example, the following command shows the deletion results for a ProtectPoint with VMAX backup of the partition 0:

```
db2acsutil load /opt/dpsapps/dbappagent/lib/lib64/
libddbboostdb2.so options @/space1/cfg/pp.cfg delete snapshot db
ps03 instance dpf10 taken at 20161007184634 dbpartitionnum 0
```

```
Instance Database Part Image Time Host First Log
=====
dpf10 PS03 0 20161007184634
Are you sure (y/[n])? y
DD Info Msg: Severity: 3 INFO: Session Created - Host :
ledmd034.lss.emc.com, User: ddvdisk

DD Info Msg: Severity: 3 INFO: Static Image get info request -
Static Image id: 040036601621057ea7f0600372000f000000020a

DD Info Msg: Severity: 3 INFO: Session Created - Host :
ledmd034.lss.emc.com, User: ddvdisk

DD Info Msg: Severity: 3 INFO: Static Image delete request -
Static Image id:
000036601621000f040036601621057ea7f0600372000f000000020a

Deleted.
```

DB2 HADR requirements for ProtectPoint operations

You can perform ProtectPoint backups and restores with the database application agent in a DB2 High Availability Disaster Recovery (HADR) environment. You can start a backup on the primary node only. DB2 does not support backups on standby nodes. The database application agent only supports HADR setups where the instance of the HADR database has the same name on all the nodes. You can configure and back up the archived logs for the database by using the DD Boost operations.

You must meet the following configuration requirements in a DB2 HADR environment:

- You have set up the database in the HADR environment according to the appropriate DB2 documentation.

Note

For ProtectPoint operations, the instance name of all the nodes must be the same, and only one node is allowed per physical host. All the database paths must be the same on all the nodes.

- You have installed the database application agent software on each node that will participate in backups or recovery, including the standby nodes. The software is required on the standby nodes in case role-switching occurs between the nodes and for recovery purposes.
 - You have completed all the post-installation procedures on each node, including the lockbox configuration procedure for all the participating hosts. The hosts can use either a shared lockbox or individually configured lockboxes.
 - You have set the `CLIENT` parameter in the DB2 configuration file to the valid hostname of one of the nodes for all the backups and recovery. [Common parameters](#) on page 80 provides details on the parameter.
-

Note

The same `CLIENT` setting must be used in all the backup and recovery operations for the HADR nodes.

- You have ensured that all the nodes have the same parameter settings.
- You have ensured that the instance of the HADR database has the same name on all the nodes.
- You have ensured that the user ID and group ID of the DB2 users on all the nodes are matched. [Performing ProtectPoint restores with the DB2 CLP](#) on page 188 provides more details on this requirement and the other requirements for restores.
- You have ensured that all the backup and restore LUNs are provisioned correctly to all the nodes as the backups and restores can be started from any node that has the primary role.

You can perform a DB2 HADR recovery on a single node or multiple nodes.

Note

As a DB2 requirement before you start an HADR recovery, you must stop HADR and deactivate the database at the recovery nodes. The DB2 documentation provides details about the required procedures.

Recovery of a single failed node

A DB2 HADR recovery of a single node requires the following steps.

1. Deactivate the database at the failed node, and then stop HADR.

If the failed node is the primary node, the node should switch the role with another node. It is a standby node when the recovery occurs.

2. Recover the failed node as if it is a stand-alone database. Run a rollforward operation without the `complete` option, which leaves the database in a rollforward pending state as required for a standby node.
3. Configure the HADR environment settings, if required.
4. Start HADR on the recovered standby node.

Recovery of all the nodes

A DB2 HADR recovery of all the nodes restores the whole HADR setup to a point-in-time, as in a disaster recovery.

A DB2 HADR recovery of all the nodes requires the following steps.

1. Recover the database to all the HADR nodes as if they are stand-alone databases.
For the standby nodes, run a rollforward operation without the `complete` option, which leaves the database in a rollforward pending state as required for a standby node.
2. Configure the HADR environment settings, if required.
3. Start HADR on all the standby nodes.
4. Start HADR on the primary node.
5. Manually back up the database.

DB2 pureScale requirements for ProtectPoint operations

In a DB2 pureScale environment, you can perform either ProtectPoint for VMAX backups and restores or ProtectPoint with RecoverPoint backups and restores. The environment is an active-active application cluster environment in which multiple database servers known as member nodes operate on a single data partition. The cluster database operates within the IBM General Parallel File Systems (GPFSs) cluster.

Each DB2 pureScale member processes its own metadata, generates independent transactions, and maintains its own transaction log files.

To perform a backup or restore on behalf of all the members in a DB2 pureScale environment, you run a single `db2 backup db` or `db2 restore db` command with the appropriate command options on any active member.

The database application agent supports only the full backup and restore of an entire database, as supported by DB2 snapshot backup and restore operations. The database application agent handles the backups of the archived logs through the DD Boost workflow only.

Overview of ProtectPoint backups and restores of a DB2 pureScale database

A DB2 pureScale database resides on IBM GPFS file systems. A GPFS file system device consists of one or more Network shared disks (NSDs), and an NSD is created by using one hard physical disk. A GPFS file system device has its own management system, which corresponds to traditional file system volume management. The terms *GPFS file system* and *GPFS file system device* are used interchangeably in this chapter.

The DB2 ACS snapshot unit grouping is performed per GPFS file system. A ProtectPoint backup or restore for a pureScale database is performed at the GPFS file system level:

- During a backup, the file system configuration is saved with the snapshot.
- During a restore, the existing file system is exported and the saved file system configuration is imported back to the GPFS global database configuration.

The database application agent supports only rollback restores of the ProtectPoint backups in a DB2 pureScale environment, not other types of restores:

- A rollback restore of a ProtectPoint for VMAX backup can be a regular or redirected rollback restore:
 - A regular rollback restore to the original pureScale cluster system is a LUN-level restore to the original source LUNs.

- A redirected rollback restore to an alternate cluster system is a LUN-level restore of the same VMAX array in a different cluster.
- A rollback restore to an XtremIO system is a RecoverPoint consistency group-level restore, which restores all the source LUNs in a consistency group.

The database application agent supports a redirected restore to relocate a pureScale database from one cluster to another cluster only with ProtectPoint for VMAX backups. [Redirected rollback restores of ProtectPoint for VMAX backups to alternate LUNs in a different cluster](#) on page 205 provides more details.

You must meet specific requirements for a ProtectPoint backup or restore of a pureScale database, including the completion of recommended preprocessing and postprocessing steps. The following topics provide details.

Configuration requirements for ProtectPoint operations in a DB2 pureScale environment

The DB2 pureScale cluster must be configured properly and must meet additional ProtectPoint requirements of the database GPFS layout. The application agent software must be installed and configured properly on all the members.

Ensure that you meet the following configuration requirements for ProtectPoint backup and restore operations in a DB2 pureScale environment:

1. Database configuration:

- You have set up the database in the DB2 pureScale environment according to the IBM documentation. All the database components reside on GPFS file systems on shared disks storage that is accessible to all the members, including the database data, system files, and log paths.

Note

Because a backup and restore are performed at the GPFS file system level, each database must have its own dedicated GPFS devices and file systems, which are separate from the devices and file systems of the instance, other databases, and other components of the GPFS cluster. For RecoverPoint, the consistency groups of the database file systems must be dedicated to the database.

For backups and restores with the `exclude logs` option, the database log directories must reside on different GPFS file systems than other database objects. For ProtectPoint with RecoverPoint operations with the `exclude logs` option, the log directories must also reside on dedicated consistency groups, separate from the consistency groups of other database objects.

This release of the database application agent does not support the DB2 rollback restores of a partial consistency group.

It is recommended by the DB2 ACS best practice to use dedicated file systems for log paths, with the log paths in a separate snapshot volume from the database directory and database containers. [Enforcing the DB2 ACS best practice on log directory layout](#) on page 184 provides details.

-
- The DB2 pureScale database is consistent when it is backed up. During a restore, all the members are up and functioning properly, to enable the export and import of the file systems.
 - It is recommended that you create the database file system with the `automount` option turned off.

If a GPFS file system is monitored and automatically mounted by GPFS management services, such as the Tivoli System Automation for Multiplatforms (SA MP), the automount activity might interfere with rollback restore operations.

In an environment with cluster services where the automount service becomes an issue in rollback restores, you must create the database file system with the `no-automount` option.

2. Database application agent configuration:

- You have installed the database application agent software on each host on which a pureScale member resides that will participate in backups or recovery.
- You have completed all the post-installation procedures on each member host, including the lockbox configuration procedure for all the participating hosts. The hosts use either a shared lockbox or individually configured lockboxes.

If an NFS-shared lockbox is configured, you have followed all the required steps in [Configuring the lockbox in a high-availability environment](#) on page 115.

- You have ensured that all the member nodes have the same parameter settings. You have created a single configuration file in a shared system folder that is accessible to all the member hosts. Alternately, you have created an identical configuration file on each member host, with the same file pathname on each host.

You have set the `CLIENT` parameter in the DB2 configuration file for all members to the same hostname, which is the hostname of one of the member nodes. [Common parameters](#) on page 80 provides details on the parameter.

- You have updated the database configuration for log archiving and recovery by using DD Boost. [Configuring DB2 transaction log archiving](#) on page 146 provides more details.

Maintaining configuration records for ProtectPoint operations

It is recommended that you maintain an up-to-date record of configuration information for the GPFS file systems, NSDs, and physical disks of the databases that you back up and restore. You can use the information during preprocessing and postprocessing steps that might be required to complete the rollback restore operations.

Before each backup and restore, run the appropriate GPFS commands as the root user to obtain the tracking information about the database GPFS file systems, the NSDs, and the hard disk devices. The *IBM GPFS Administration and Programming Reference* provides more details about the GPFS commands:

- To obtain information about the GPFS file systems on which the database is created, run the `mmlsnsd` and `mmlsfs` commands. The following example commands include the database file system `fs_hi`:

```
export PATH=$PATH:/opt/emc/SYMCLI/bin:/usr/lpp/mmfs/bin
mmlsnsd -f fs_hi -X
```

Disk name	NSD volume ID	Device	Devtype	Node name	Remarks

```
gpfs1007nsd 0AF1AF90570FB853 /dev/sdh generic ledmf144.lss.emc.com
gpfs1008nsd 0AF1AF90570FB854 /dev/sdi generic ledmf144.lss.emc.com
```

```
mmlsfs fs_hi -T
```

flag	value	description
-T	/sd_hi	Default mount point

- To retrieve a copy of the database file system configuration to a text file, which can be used in a file system import if required, run the `mmbackupconfig` command. For example:

```
mmbackupconfig fs_hi -o fs_hi.bk.cfg
```

- To obtain information about LUNs in a VMAX environment, run the `sympd list` command. For example:

```
sympd list
```

```
Symmetrix ID: 000196701031
      Device Name          Dir          Device
-----
Physical          Sym  SA :P  Config          Attribute  Sts  Cap
(MB)
-----
/dev/sdb          0089C 01D:009 TDEV          N/Grp'd   RW   6
/dev/sdc          0089D 01D:009 TDEV          N/Grp'd   RW   6
/dev/sdh          008A0 01D:009 TDEV          N/Grp'd   RW  8629
/dev/sdi          008A1 01D:009 TDEV          N/Grp'd   RW  8629
/dev/sdj          008A2 01D:009 TDEV          N/Grp'd   RW  8629
/dev/sdk          008A3 01D:009 TDEV          N/Grp'd   RW  8629
/dev/sdl          008A4 01D:009 TDEV          N/Grp'd   RW  8629
```

- Run the `syminq` command to obtain information about the XtremIO device in a RecoverPoint environment. For example:

```
syminq
```

Device		Product		Device		
Name	Type	Vendor	ID	Rev	Ser Num	Cap (KB)
...						
/dev/sdb		XtremIO	XtremApp	4030	514F0C58C5800019	6291456
/dev/sdc		XtremIO	XtremApp	4030	514F0C58C580001A	6291456

Performing ProtectPoint backups in a DB2 pureScale environment

To perform a ProtectPoint backup of the entire DB2 pureScale database after the database application agent is installed and configured, a DB2 user on an active member can run the `db2 backup` command. For example:

```
db2 backup db sample use snapshot library /opt/dpsapps/
dbappagent/lib/lib64/libddbostdb2.so options @/home/cfg/pp.cfg
```

In this example, the command backs up the *sample* database by using the settings in the configuration file `/home/cfg/pp.cfg`.

Performing ProtectPoint restores in a DB2 pureScale environment

In a DB2 pureScale environment, typically you can run a restore from an active member. Prior to the restore, ensure that you meet the restore requirements as described in the next topic.

For example, in a restore to the original cluster, a DB2 user on an active member can run the following command to perform a rollback restore:

```
db2 restore db <sample> use snapshot library /opt/dpsapps/
dbappagent/lib/lib64/libddbboostdb2.so options @/home/cfg/pp.cfg
taken at <yyyymmddhhmmss> logtarget include force
```

In this example, the command restores the ProtectPoint backup of the *sample* database by using the settings in the configuration file `/home/cfg/pp.cfg`. When the timestamp is not specified, the latest ProtectPoint backup is restored.

ProtectPoint restore requirements in a DB2 pureScale environment

You must meet the DB2 specific requirements for a pureScale restore. The database application agent supports only a rollback restore for a pureScale database. You must also meet the requirements for a rollback restore of a ProtectPoint backup in [Configuring rollback restores of ProtectPoint backups](#) on page 97.

A rollback restore is destructive, and occurs at the GPFS file system level. The target restored GPFS file system configuration is exported, and the backup GPFS configuration is imported back. The target NSD LUNs are overwritten with the original content.

Note

DB2 10.5 or later supports the restore of an offline database backup of a DB2 pureScale instance to DB2 Enterprise Server Edition. DB2 10.5 or later also supports the restore of an offline backup of DB2 Enterprise Server Edition to a DB2 pureScale instance. However, the database application agent does not support these types of restores for ProtectPoint backups in a pureScale environment.

Ensure that you meet the database configuration requirements in [Configuration requirements for ProtectPoint operations in a DB2 pureScale environment](#) on page 200. Ensure that you also meet the following requirements specifically for restores:

- All the pureScale members are up and functioning properly, to enable the export and import of the file systems.
- For ProtectPoint with RecoverPoint restores, the database has a dedicated consistency group. If you must maintain extra LUNs in a RecoverPoint consistency group that are not used by the database, use the post-processing steps in a following topic to maintain the extra LUNs.
- The `RESTORE_TYPE_ORDER=rollback` parameter setting exists in the configuration file of the database application agent. [Configuring rollback restores of ProtectPoint backups](#) on page 97 provides details.
- The `psrollback.res` file lists the files and directories to be excluded from the rollback safety checks. [Configuring rollback restores of ProtectPoint backups](#) on page 97 provides details.
- For a restore to the original cluster, the restore is performed to the original source LUNs. For a redirected restore of a VMAX backup to a different cluster, the restore is performed to different LUNs. [Redirected rollback restores of](#)

[ProtectPoint for VMAX backups to alternate LUNs in a different cluster](#) on page 205 provides details.

In both cases, the file system (with the same name and mount point) must exist and must be mounted. If the file system does not exist, re-create the file system based on the information that was saved during the backup.

- Any NSD or file system naming conflicts, which can cause issues in a file system import, must be resolved before the restore starts.

Note

The original NSD names and physical LUNs are not in use for any other purpose or by other file systems. No free NSD exists with the same name as an original NSD in the backed-up file systems. The *IBM GPFS Administration and Programming Reference* provides details.

After a rollback restore of a ProtectPoint backup, you might need to perform additional postprocessing steps if any of the following conditions exist:

- Physical disks were added to the database file system after the backup time.
- The rollback restore fails in the middle of the operation, when the file system configuration might already be exported from the system.
- The rollback restore of a ProtectPoint with RecoverPoint backup restores the LUNs in the Recoverpoint consistency group that were not part of the backup command.

The following topics provide details about any required postprocessing steps.

Postprocessing due to additional physical disks

In the original backed-up system, if physical disks were added to the database file system after the backup time, the rollback restore of the database backup does not use those disks. However, the disks still carry an NSD signature, which must be scrubbed off to enable the disks to be reused after the restore.

Review the configuration from the backup time and from before the restore time to determine which additional physical disks need to be cleaned up. [Maintaining configuration records for ProtectPoint operations](#) on page 201 provides details about how to obtain the configuration information. Check the configuration of both the GPFS and LVM file systems to ensure that the disks are not being used. Run the `mkfs` command to format the disks and clean up their GPFS related information. For example:

- On Linux, run the `mkfs -t ext2 /dev/sdj` command.
- On AIX, run the `mkfs -V jfs2 /dev/hdisk15` command.

Note

The NSD and the physical disk layout of the existing file system (before it was exported in the restore) was recorded in the temporary output file `/opt/dpsapps/dbappagent/tmp/mmlnsd<file_system_name>_output<dbsmd_pid>`. This NSD and physical disk layout information was produced by the `mmlnsd -f -X` command during the restore.

Postprocessing due to a rollback restore failure

The rollback restore might fail in the middle of the operation, when the file system configuration is already exported from the system. To restart the rollback restore, the file system must be either imported back or re-created.

First, check if the file system configuration is exported. For example, run the `mmlsrusd -f file_system_name -X` command. If the file system is not exported, it is not recognized by the command.

If the file system is exported, then either import back or re-create the file system:

- To import back the file system configuration that was exported during the rollback restore, run the following command:

```
mmimportfs <file_system_name> -i <input_file>
```

where:

- *file_system_name* is the name of the restored file system.
- *input_file* is the file that contains the output of the `mmexportfs` command, as run on the file system `<file_system_name>` during the restore. The file was saved as the temporary file `/opt/dpsapps/dbappagent/tmp/mmexportfs<file_system_name><dbsmd_pid>`.

After you import the file system, run the `mmlsrusd -f <file_system_name> -X` command to verify that the file system resumed correctly.

- If there is a reason that the `mmimportfs` command cannot be run to resume the file system, then re-create the file system by using the configuration information that was obtained from the backup time. For example:

```
db2cluster -create -filesystem fs_hi -disk /dev/sdh,/dev/sdi -mount /sd_hi
```

The file system configuration from before the export was recorded during the restore in the temporary file `/opt/dpsapps/dbappagent/tmp/mmlsrusd<file_system_name>_output<dbsmd_pid>`. This configuration can also be used, provided that it is in line with the file system configuration in the backup, as required.

Postprocessing due to extra LUNs in a RecoverPoint consistency group

In a ProtectPoint with RecoverPoint rollback restore, the original contents of all the LUNs in the same consistency group are restored. However, only the GPFS file system that was included in the backup command has its original configuration restored. The consistency groups must be dedicated to the database file systems.

If the consistency group includes extra LUNs that are not part of the backup, then you must manually clean up the LUNs or the file systems that are not part of the backup. For example, the extra LUNs resulted from the addition of new NSDs to the file system after the backup. If you must maintain extra LUNs in a RecoverPoint consistency group that are not used by the database, then back up the configuration of the file systems of those extra LUNs before the database backup. You can then use the configuration to perform an import after the database restore, where all the LUNs of the consistency group are restored. In a restore of the database, you must unmount and export the existing extra file system that was not managed by the database backup. After the database restore is complete, manually import the saved original file system configuration and mount the file system back.

Redirected rollback restores of ProtectPoint for VMAX backups to alternate LUNs in a different cluster

You can perform a relocated restore of a pureScale database to a different cluster by running a redirected rollback restore with the database application agent to alternate

LUNs in the same VMAX array on the target cluster. In the rollback restore, the GPFS file system configuration in the backup is restored to the different cluster, replacing the existing file system.

Note

A rollback restore to the original cluster is performed to the original source LUNs.

Ensure that you meet the general requirements in [Preparing for DB2 disaster recovery](#) on page 190. In addition, ensure that you meet the DB2 specific requirements for a pureScale database restore from one cluster to a different cluster. It is recommended that you use a consistent backup for the restore. When you use an offline backup, perform the restore with the option `without rolling forward`. Review the following topic before you run a redirected rollback restore of a ProtectPoint for VMAX backup.

Requirements for a redirected rollback restore of a ProtectPoint for VMAX backup

Ensure that you meet the following ProtectPoint requirements before you start a redirected rollback restore of a ProtectPoint for VMAX backup in a pureScale environment:

- The target GPFS file system must exist before the restore and must have the same name and mount point as the backed-up file system. Re-create the original GPFS file system layout as required for the rollback restore. Follow the guidelines in [GPFS file system requirements for a redirected rollback restore of a ProtectPoint for VMAX backup](#) on page 207.
- The number of physical devices (LUNs and NSDs) in the target GPFS file system must be the same as the number in the backed-up file system.

Note

The names of the target LUNs and NSDs can be different from those of the backed-up file system. However, there must be no NSD name conflict between the backed-up file system and the alternate cluster. For example, no NSD exists in file systems outside of the target file system with the same NSD name as an NSD to be restored. All conflicts that can cause an import failure must be resolved before the restore.

-
- The size of the target LUNs must be equal to or greater than the size of the original LUNs.
 - The ProtectPoint VMAX devices must be configured and provisioned properly for the rollback restore.
 - The pureScale instance members must be configured properly on the alternate cluster according to the DB2 documentation.
 - The alternate cluster must follow the ProtectPoint restore requirements from the preceding topics in this chapter. The numeric user ID (UID) and group ID (GID) of the target DB2 instance user must match the UID and GID of the original DB2 user, as recorded in the backup.
 - The database application agent software must be installed and configured properly on the pureScale members that perform the rollback restore. The `CLIENT` parameter must be set to the original value, as recorded in the backup.

GPFS file system requirements for a redirected rollback restore of a ProtectPoint for VMAX backup

Complete the following steps to ensure that the GPFS file system layout meets the requirements for a redirected rollback restore of a ProtectPoint for VMAX backup:

1. Obtain the original GPFS layout and configuration of the backed-up file system. It is recommended that you keep a copy of the backed-up file system configuration. You can also retrieve the device layout information that was saved in the backup by running the `ddbmadmin` command.

The following `ddbmadmin` command displays the save set snapshot handle metadata, which contains the GPFS configuration and physical device information that was saved during the backup but not the size of the backed-up devices:

```
ddbmadmin -s -v -b <start_backup_time> -e <end_backup_time> -n
-z <configuration_file><application> [-D 9]
```

For example, the following `ddbmadmin` command sets the `start_backup_time` value based on the DB2 backup timestamp 20160929112916:

```
ddbmadmin -s -v -b 'Sep 29 11:29:14' -e now -n db2 -z /
home/cfg/dd/pp.cfg 2>&1 | tee ../logs/lastbk.log
```

This backup in this example contains two GPFS file systems, `fs_dbData` and `fs_dbLog`. The command output in the `lastbk.log` file lists the configuration information under the backup file name, `ACS*20160929112916*` for each file system.

The command output lists the following information for the `fs_dbData` file system:

- Mount point: `/dbData`
- NSDs: `gpfs1001nsd`, `gpfs1002nsd`
- Physical devices: `dev/sdp`, `/dev/sdq`

The relevant part of the command output with the `fs_dbData` file system information is as follows:

```
Record file = Record file = /su ledmf144_5/ledmf144.lss.emc.com/
27/2.0/meta_rec/$db2_acs:$/_ts10k_147516/1475163078.rec.
client = ledmf144.lss.emc.com, date and time = 09/29/16 11:31:18,
size = 4108, ssid = 1475163078, name =
ACS.DB2INST1.NODE0000.PSTEST1.20160929112916
ssid=245fcf5c-00000011-00000000-57ed33c6-57ed33c6-da00aa56
(1475163078), date and time=09/29/16 11:31:18 (1475163078),
host=ledmf144.lss.emc.com, name
...
%%home%:40_SG_ETCFS:fs_dbData:1:%2fdbData:
%%home%:40_SG_ETCFS:fs_dbData:2:          dev          = /dev/
fs_dbData
%%home%:40_SG_ETCFS:fs_dbData:3:          vfs           = mmfs
%%home%:40_SG_ETCFS:fs_dbData:4:          nodename      = -
%%home%:40_SG_ETCFS:fs_dbData:5:          mount        = mmfs
%%home%:40_SG_ETCFS:fs_dbData:6:          type         = mmfs
%%home%:40_SG_ETCFS:fs_dbData:7:          account     = false
%%home%:50_SG_MOUNT:fs_dbData:rw:mtime:atime:.....:
%%home%:60_SG_DISKS:fs_dbData:1:gpfs1001nsd:
17671680:-1:dataAndMetadata:
0AF1AF9057ED2F4E:nsd:::other:::generic:cmd:.....:system:.....:
%%home%:60_SG_DISKS:fs_dbData:2:gpfs1002nsd:
```

```
17671680:-1:dataAndMetadata:
0AF1AF9057ED2F4F:nsd:::other::generic:cmd:::::system:::::
|GPFS|7|3|5|0|0|/dev/sdp /dev/sdq ";
```

The command output lists the following information for the `fs_dbLog` file system:

- **Mount point:** `/dbLog`
- **NSDs:** `gpfs1003nsd`
- **Physical devices:** `dev/sdr`

The relevant part of the command output with the `fs_dbLog` file system information is as follows:

```
Record file = /su_ledmf144_5/ledmf144.lss.emc.com/27/2.0/meta_rec/
$db2_acs:$/_ts10k_147516/1475163073.rec.
client = ledmf144.lss.emc.com, date and time = 09/29/16 11:31:13,
size = 3404, ssid = 1475163073, name =
ACS.DB2INST1.NODE0000.PSTEST1.20160929112916
...
%%home%:40_SG_ETCFS:fs_dbLog:1:%2FdbLog:
%%home%:40_SG_ETCFS:fs_dbLog:2:          dev          = /dev/
fs_dbLog
%%home%:40_SG_ETCFS:fs_dbLog:3:          vfs          = mmfs
%%home%:40_SG_ETCFS:fs_dbLog:4:          nodename     = -
%%home%:40_SG_ETCFS:fs_dbLog:5:          mount       = mmfs
%%home%:40_SG_ETCFS:fs_dbLog:6:          type        = mmfs
%%home%:40_SG_ETCFS:fs_dbLog:7:          account    = false
%%home%:50_SG_MOUNT:fs_dbLog::rw:mtime:atime:::::::::::
...
%%home%:60_SG_DISKS:fs_dbLog:1:gpfs1003nsd:
17671680:-1:dataAndMetadata:
0AF1AF9057ED3073:nsd:::other::generic:cmd:::::system:::::
|GPFS|7|3|5|0|0|/dev/sdr ";
```

2. Create the target GPFS file systems by using the same mount points.

For example, the following commands create the `fs_dbData` and `fs_dbLog` file systems with the same mount point as used in the backed-up file systems:

```
db2cluster -create -filesystem fs_dbData -disk /dev/sdf,/dev/
sdg -mount /dbData
File system 'fs_dbData' has been successfully created.
```

```
db2cluster -create -filesystem fs_dbLog -disk /dev/sdl -mount /
dbLog
File system 'fs_dbLog' has been successfully created.
```

List the NSDs and file system information to ensure that no conflicts exist with the original systems from step 1. Ensure that each file system has the same number of NSDs and physical disks as in the backed-up file system.

Note

The `db2cluster` command does not include an option to specify the NSD name. If you need to resolve a conflict in NSD naming, you can use the GPFS utility directly to modify the layout at the NSD layer. The IBM GPFS documentation provides more details.

For example, the following `mmlsnsd` and `mmlsfs` commands list the NSDs and file system information for the `fs_dbData` and `fs_dbLog` file systems:

```
mmlsnsd
```

File system	Disk name	NSD servers
db2fs1	gpfs1nsd	(directly attached)
fs_dbData	gpfs1011nsd	(directly attached)
fs_dbData	gpfs1012nsd	(directly attached)
fs_dbLog	gpfs1013nsd	(directly attached)

```
mmlsnsd -X
```

Disk name	NSD volume ID	Device	Devtype	Node name	Remarks
gpfs1011nsd	0AF1AF9957EDA754	/dev/sdf	generic	ledmf153.lss.emc.com	
gpfs1012nsd	0AF1AF9957EDA755	/dev/sdg	generic	ledmf153.lss.emc.com	
gpfs1013nsd	0AF1AF9957EDA785	/dev/sdl	generic	ledmf153.lss.emc.com	
gpfs1nsd	0AF1AF9956E6EC21	/dev/sdd	generic	ledmf153.lss.emc.com	

```
mmlsfs fs_dbData -T -d
```

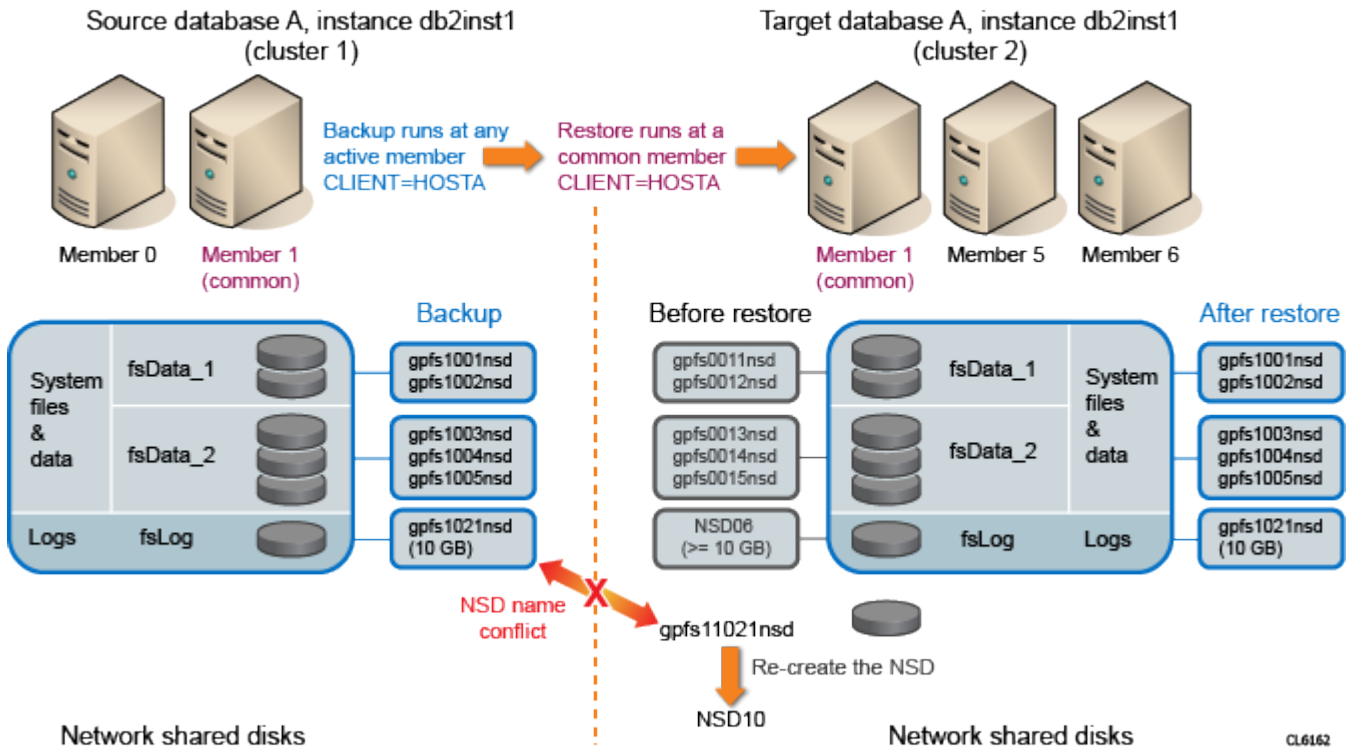
flag	value	description
-d	gpfs1011nsd;gpfs1012nsd	Disks in file system
-T	/dbData	Default mount point

```
mmlsfs fs_dbLog -T -d
```

flag	value	description
-d	gpfs1013nsd	Disks in file system
-T	/dbLog	Default mount point

The following figure illustrates the general requirements of the target file system layout in a redirected rollback restore of a ProtectPoint for VMAX backup to an alternate pureScale cluster.

Figure 11 Target file system layout requirements for a redirected rollback restore to an alternate pureScale cluster



Backups and restores of transaction logs in a DB2 pureScale environment

In a DB2 pureScale environment, each member generates independent transactions and maintains its own set of transaction log files. To enable DB2 rollforward recovery, the transaction logs must be backed up. The transaction logs are always backed up through the DD Boost workflow.

[Configuring DB2 transaction log archiving](#) on page 185 provides more details.

Deleting ProtectPoint backups in a DB2 pureScale environment

As required for snapshot backups, you must use the `db2acsutil` utility to delete the ProtectPoint backups in a DB2 pureScale environment. You cannot use the `db2 prune` command to delete the ProtectPoint backup entries from the recovery history file.

You can run the `db2acsutil` command with the appropriate options to delete the ProtectPoint backups. [Managing and deleting ProtectPoint DB2 backups](#) on page 189 provides details.

DB2 troubleshooting tips for ProtectPoint operations

[General troubleshooting tips](#) on page 132 provides common troubleshooting information that applies to the database application agent operations with all the supported databases and applications.

Errors reported by the snapshot agent or DB2 shared library are treated as fatal and reported as a backup error to the DB2 software.

For ProtectPoint backups and restores, the log and debug files are different than those used for DD Boost operations:

- All the debug messages are recorded in a single file named `libddbboostdb2_acs_DB2_date.time.DB2_pid.log`.
- The only operational logs are the DB2 diagnostic log and the snapshot agent operational log, `ddbsm.log`.

You must set the `DEBUG_LEVEL` parameter to enable debugging.

DB2 ProtectPoint restore might fail with DB2 error code SQL2081N

A DB2 ProtectPoint restore might fail with an SQL message and error code 2 or 3. The failure occurs when the restore starts and checks for other databases that exist on the restore file system. For example:

```
db2 restore db pluto use snapshot library /opt/dpsapps/
dbappagent/lib/lib64/libddbboostdb2.so options @/home/db2inst1/
db2.cfg taken at 20170419130621 LOGTARGET include force
```

```
SQL2081N  A snapshot restore failed because some element or elements
of the database to be restored have the same name as some element or
elements of existing databases. Reason code = "2".
```

In a snapshot restore, a DB2 snapshot utility must uniquely identify the element of the restored database. If the restore directory is also used by another database, the snapshot restore might fail because DB2 cannot identify the restore database.

For example, the DB2 ProtectPoint restore might fail when multiple databases share the same path in the restore directory. When a directory or file has a conventional key value of a DB2 database token, such as `SQL00001` or `LOGSTREAM0000`, DB2 might interpret it as another database residing on the same file system. In such a case, a Protectpoint database restore might fail.

Check the `db2diag.log` file for more information about the issue. The IBM documentation provides more details about the SQL2081N error.

Ensure that the restore file systems, especially the database home and log directories, are dedicated to the database that will be restored. When the restore includes the logs from the backup image, ensure that the log directory is empty.

CHAPTER 7

DD Boost Operations on Oracle Systems

This chapter includes the following topics:

- [Overview of DD Boost operations in an Oracle environment](#)..... 214
- [Configuration of DD Boost operations in an Oracle environment](#)..... 216
- [Migrating an Oracle configuration from DD Boost for RMAN 1.x or later](#).....220
- [Performing DD Boost backups and restores with Oracle RMAN](#).....224
- [Performing DD Boost backups and restores with Oracle Enterprise Manager](#).. 224
- [Performing backups and restores of Oracle CDBs and PDBs](#)..... 225
- [Performing Oracle backup deletion and maintenance operations](#)..... 225
- [Preparing for Oracle disaster recovery](#).....226
- [Oracle RAC and active-passive cluster requirements for DD Boost operations](#).227
- [Oracle troubleshooting tips for DD Boost operations](#).....227

Overview of DD Boost operations in an Oracle environment

The database application agent is integrated with the Oracle RMAN interfaces for third-party media management vendors to enable Oracle data backups, restores, and archived redo log operations. The database application agent also supports Oracle backup deletion and maintenance operations.

You can perform a backup or recovery with the product on an Oracle database server by running one of the supported Oracle backup or recovery tools:

- Oracle Recovery Manager (RMAN) with the `rman` command
- Oracle Enterprise Manager GUI

You can use these tools in cooperation with the database application agent to perform all the operations supported by Oracle RMAN, including the following operations:

- Online and offline backups
- Full and incremental backups
- Archived redo log backups
- Recovery of a database to the current time or a specific point-in-time
- Recovery to the original location or an alternate location
- Backup and recovery of databases, tablespaces, and archived redo logs
- Backup deletion and other maintenance operations

The product maintains online backup indexes on the Data Domain system, which are in addition to the Oracle RMAN catalog. During backups, the product creates backup entries in the online indexes, which provide the information required to restore the backed-up data.

The troubleshooting section at the end of this chapter provides details about limitations in the DD Boost operations with the database application agent in an Oracle environment.

Oracle backup processes

An Oracle backup includes the following process interactions.

1. The database administrator initiates the backup through one of the following methods:
 - To invoke the RMAN backup script, the database administrator runs an `rman` command, such as the following `rman` command:

```
rman target /@SNB catalog rman/rman@catdb cmdfile '/orasnb/
backup.txt'
```

Note

Instead of appearing on the command line, the password could be included with a `connect` command in the RMAN script. In that case, the `rman` command line could be as follows:

```
rman @/orasnb/backup.txt
```

- The database administrator runs the Oracle Enterprise Manager to generate the RMAN backup script and perform the backup operations.
2. The Oracle software loads the Oracle shared library used by the database application agent, as specified by the `SBT_LIBRARY` parameter.
 3. The database application agent reads the configuration file specified by the `CONFIG_FILE` parameter, and then initializes the connection with the Data Domain system, based on the settings in the configuration file.
 4. The Oracle software ensures that each backup piece name is unique by asking the database application agent to check if the backup piece name exists in the database application agent catalog.
 5. If the database application agent responds negatively (as expected), the Oracle software sends the backup pieces to the database application agent through the SBT API.
 6. The database application agent uses the DD Boost interface to send the backup data to the Data Domain system for storage, and catalogs the backup.
 7. The Oracle software asks the database application agent to confirm that the backup is in the catalog, then records the entry in the Oracle catalog and completes the backup.

Oracle restore processes

An Oracle restore includes the following process interactions.

1. The database administrator initiates the restore through one of the following methods:
 - To invoke the RMAN restore script, the database administrator runs an `rman` command, such as the following `rman` command:

```
rman target /@SNB catalog rman/rman@catdb cmdfile '/orasnb/restore.txt'
```

Note

Instead of appearing on the command line, the password could be included with a `connect` command in the RMAN script. In that case, the `rman` command line could be as follows:

```
rman @/orasnb/restore.txt
```

- To generate the RMAN restore script and perform the restore and recovery operations, the database administrator runs the Oracle Enterprise Manager.
2. The Oracle software loads the Oracle shared library used by the database application agent, as specified by the `SBT_LIBRARY` parameter.
 3. The database application agent reads the configuration file specified by the `CONFIG_FILE` parameter, and then initializes the connection with the Data Domain system, based on the settings in the configuration file.
 4. The Oracle software queries, and then requests the backup pieces from the database application agent through the SBT API.

5. To query the catalog and retrieve the backup data from the Data Domain system, the database application agent uses the DD Boost interface.

Oracle backups of archived redo logs

Backups of archived redo logs enable recovery of an Oracle database to its predisaster state. Without these backups, you can recover the database only to the time of the last consistent Oracle backup. In this case, you will lose the transactions that occurred between the time of the last consistent backup and the time of the database corruption.

You might want to perform a full or incremental backup every 24 hours at a minimum, and schedule more frequent backups of only the archived redo logs.

You can back up the archived redo logs by using the appropriate option of the RMAN backup command.

Configuration of DD Boost operations in an Oracle environment

You must complete the required configurations of the database application agent to enable the DD Boost operations in an Oracle environment. The following topics provide the product configuration details.

[Oracle RAC and active-passive cluster requirements for DD Boost operations](#) on page 227 provides additional details on the specific configuration requirements in an Oracle RAC or active-passive cluster environment.

The troubleshooting section at the end of this chapter provides details about limitations in the DD Boost operations with the database application agent in an Oracle environment.

Setting up the configuration file in an Oracle environment

It is recommended that you set the required parameters for Oracle operations in the configuration file used by the database application agent. For example, the configuration file named `oracle_ddbda.cfg` contains the following mandatory parameter settings:

```
DDBOOST_USER=qa_ost
DEVICE_HOST=bu-dbe-890.lss.emc.com
DEVICE_PATH=/bu-star1_ora
```

[Setting up the configuration file](#) on page 78 describes the common parameters and how to set parameters in the configuration file. [Configuring restores of replicated backups](#) on page 90 also describes the parameters and requirements for the restores of replicated backups.

After the configuration file is set up, ensure that the required lockbox procedures have been performed as described in [Configuring the lockbox](#) on page 103.

Creating the RMAN scripts for DD Boost Oracle operations

You must create the required RMAN script for the Oracle backup or restore operations.

Note

In the RMAN script, the % character is not supported in the `FORMAT` string unless the character is used as part of an RMAN substitution variable.

You must set the `SBT_LIBRARY` and `CONFIG_FILE` parameters, either in the configuration file or in the RMAN script:

- Set the `SBT_LIBRARY` parameter to the pathname of the Oracle library used by the database application agent.
- Set the `CONFIG_FILE` parameter to the pathname of the configuration file. Use the correct option if you set `CONFIG_FILE` in the RMAN script:
 - With Oracle 11.2 or later, use the `SBT_PARMS` option.
 - With Oracle 11.1 or earlier, use the `SEND` option.

To optimize the performance of DD Boost operations, use the parameter setting `BLKSIZE=1048576`.

The following examples show the correct parameter settings in the RMAN script:

- On UNIX or Linux, using the `SBT_PARMS` option with Oracle 11.2 or later:

```
ALLOCATE CHANNEL C1 DEVICE TYPE SBT_TAPE PARMS
'BLKSIZE=1048576, SBT_LIBRARY=/opt/dpsapps/dbappagent/lib/lib64/
libddboostora.so, SBT_PARMS=(CONFIG_FILE=/orasnb/
oracle_ddbda.cfg)' FORMAT '%d_%U';
BACKUP DATABASE;

CONFIGURE CHANNEL DEVICE TYPE SBT_TAPE PARMS 'BLKSIZE=1048576,
SBT_LIBRARY=/opt/dpsapps/dbappagent/lib/lib64/libddboostora.so,
SBT_PARMS=(CONFIG_FILE=/orasnb/oracle_ddbda.cfg)';
BACKUP DEVICE TYPE SBT DATABASE FORMAT '%d_%U';
```

- On Windows, using the `SBT_PARMS` option with Oracle 11.2 or later:

```
ALLOCATE CHANNEL C1 DEVICE TYPE SBT_TAPE PARMS
'BLKSIZE=1048576, SBT_LIBRARY=C:\PROGRA~1\DPSAPPS\DBAPPAGENT\bin
\libddboostora.dll, SBT_PARMS=(CONFIG_FILE=D:\orasnb
\oracle_ddbda.cfg)' FORMAT '%d_%U';
BACKUP DATABASE;

CONFIGURE CHANNEL DEVICE TYPE SBT TAPE PARMS 'BLKSIZE=1048576,
SBT_LIBRARY=C:\PROGRA~1\DPSAPPS\DBAPPAGENT\bin
\libddboostora.dll, SBT_PARMS=(CONFIG_FILE=D:\orasnb
\oracle_ddbda.cfg)';
BACKUP DEVICE TYPE SBT DATABASE FORMAT '%d_%U';
```

Note

On Windows, you must use the short Windows pathname in the `SBT_LIBRARY` setting, as shown in the preceding examples. Otherwise, if the pathname contains any spaces, the Oracle software displays a syntax error.

- Using the `SEND` option with Oracle 11.1 or earlier:

```

ALLOCATE CHANNEL C1 DEVICE TYPE SBT_TAPE PARMS
'BLKSIZE=1048576, SBT_LIBRARY=/opt/dpsapps/dbappagent/lib/lib64/
libddbboostora.so FORMAT '%d_%U';
SEND CHANNEL C1 'ENV=(CONFIG_FILE=/orasnb/oracle_ddbda.cfg)';
BACKUP DATABASE;

CONFIGURE CHANNEL DEVICE TYPE SBT_TAPE PARMS 'BLKSIZE=1048576,
SBT_LIBRARY=/opt/dpsapps/dbappagent/lib/lib64/libddbboostora.so';
SEND 'ENV=(CONFIG_FILE=/orasnb/oracle_ddbda.cfg)';
BACKUP DEVICE TYPE SBT DATABASE FORMAT '%d_%U';

```

NOTICE

Consider the following restrictions when you create an RMAN script:

- RMAN multiplexing might have a negative impact on the deduplication ratio. Set `FILESERSET` or `MAXOPENFILES` accordingly.
- `SET` or `BACKUP COPIES` is not supported. If you set `BACKUP COPIES` in the RMAN script, the backup will fail. If you want to use Data Domain replication, refer to [Data Domain replication](#) on page 37.

Configuring operations in an Oracle Data Guard environment

The database application agent supports Oracle Data Guard, which is an Oracle data availability and protection solution for a primary database and one or more standby databases over an IP network. You can configure backup and restore operations with the database application agent in an Oracle Data Guard environment.

In an Oracle Data Guard environment, as transactions occur in the primary database and as Oracle writes redo data to the local redo logs, Data Guard automatically performs the following operations:

- Transfers this redo data to the standby sites.
- Applies the redo data to the standby databases, which synchronizes the standby databases with the primary database.

You can offload RMAN backups of datafiles, archived redo logs, and possibly other files to a physical standby database. You can then use the backups to recover the primary or standby database. RMAN and Data Guard documentation describes how to configure and back up a physical standby database, and use the backups to recover the primary or standby database.

To configure backups and restores with the database application agent in an Oracle Data Guard environment:

1. Follow the instructions in the Oracle documentation about how to set the required RMAN configurations, for example, to use a Recovery Catalog and the `DB_UNIQUE_NAME` parameter.
2. Install and configure the database application agent software on the primary database host, and then on each physical standby database host that is included in the backups and restores.
3. For a backup, create an RMAN script, and then set the parameters in the configuration file for the database application agent to back up data from a physical standby database, which can be used to restore the primary database.

Set the `CLIENT` parameter in the configuration file to a single value that identifies the Data Guard environment, preferably the primary database hostname.

4. For a recovery, create an RMAN script, and then set the parameters in the configuration file for the database application agent to recover the data from a primary or standby database, depending on the restore. Set the `CLIENT` parameter to the same value as used during the backup.

Setting up Oracle Optimized Deduplication

The database application agent supports the configuration of Oracle Optimized Deduplication with a Data Domain appliance.

Use the following commands to administer Oracle Optimized Deduplication at the system level:

- `fileSYS option set app-optimized-compression {none | oracle1}`
- `fileSYS option reset app-optimized-compression`
- `fileSYS option show app-optimized-compression`

The *Data Domain Operating System Command Reference Guide* on the Support website provides details about these commands.

Estimating the Data Domain resource usage on Oracle systems

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

Capacity usage on Oracle systems

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

```
163542 12/09/2016 02:10:00 PM (pid8651) SYSTEM critical Unable to
write to a file due to reaching the hard quota limit.
The error message is: [5194] [ 8651] [139771055018560] Fri Dec 9
14:10:00 2016
      ddp_write() failed Offset 167772160, BytesToWrite 524288,
BytesWritten 0 Err: 5194-Hard Quota Exceeded >
```

Streams usage on Oracle systems

Note

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

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

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

```
153004 05/10/2016 01:42:48 PM (pid6662) SYSTEM critical Unable to
write to a file because the streams limit was exceeded.
```

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

```
RMAN-03009: failure of backup command on CH5 channel at 03/30/2016
16:38:20 ORA-19502: write error on file "CER4_lsrlpqjo_1_1", block
number 33 (block size=8192)
ORA-27030: skgfwr: sbtwrite2 returned error ORA-19511: Error
received from media manager layer, error text: asdf_output_section1()
failed xdr=0x0xb8183f8: bp=0x0xc19c538: send_len=262144: type=12800:
fhand=0x0xc17e688: wrapper=0x(nil): directp=0x0x7f268ad9e000 (1:4:22)
channel CH5 disabled, job failed on it will be run on another channel
```

Although the backup job is reassigned and the backup succeeds, the failed channel generates an incomplete save sets record in the index. You can run the `ddbmadmin -s` command to list the save sets and see the incomplete record. For example:

```
INCOMPLETE: client = ledma178.lss.emc.com, date and time = 03/30/2016
04:38:19 PM (1459370299)
```

The incomplete record does not affect any restore of the backup because the save set is backed up by another channel. You can run the `ddbmadmin -d` command to delete the incomplete save sets. For example:

```
ddbmadmin -d -b 1459370299 -e 1459370299 -n oracle -z
configuration_file
```

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

```
163971 11/23/2016 03:03:24 PM (pid26011) SYSTEM critical Unable to
read from a file because the streams limit was exceeded.
The error message is: [5519] [26011] [140130349429312] Wed Nov 23
15:03:24 2016
      ddp_read() failed Offset 0, BytesToRead 262144, BytesRead 0 Err:
5519-nfs readext remote failed (nfs: Resource (quota) hard limit
exceeded)
```

Migrating an Oracle configuration from DD Boost for RMAN 1.x or later

You must complete the required steps to migrate an Oracle configuration from DD Boost for RMAN 1.x or later to the database application agent 4.0.

Procedure

1. Discontinue all backups performed with DD Boost for RMAN.
2. Install the database application agent 4.0 according to the instructions in [Road map to install or update the software](#) on page 60.

Note

The database application agent installation does not remove or overwrite any of the DD Boost for RMAN files on the same system.

3. Configure the database application agent 4.0 according to the instructions in the preceding topics.
 - a. Create a new lockbox for the database application agent operations. You must register the Data Domain system in the lockbox by using the

`ddbmadmin -P -z configuration_file` command as described in [Configuring the lockbox](#) on page 103.

Do not use the following command for any lockbox operations:

```
send 'set username user password password servername
dd_hostname';
```

- b. Update the RMAN scripts as described in [Updating the RMAN scripts used with DD Boost for RMAN 1.x or later](#) on page 221.
4. Perform an initial full backup of the Oracle database with the database application agent 4.0. Include the control file and archived logs in the backup so that ongoing backups performed with the database application agent have no dependency on backups previously performed with DD Boost for RMAN.
5. Use the correct type of RMAN script for restore operations as described in [Using the correct RMAN script for restore operations](#) on page 222.

You cannot use the database application agent 4.0 to recover backups performed with DD Boost for RMAN. You can recover such backups only by using DD Boost for RMAN.

You can maintain the DD Boost for RMAN software on the Oracle server host as long as needed, and use the software to recover backups previously performed with DD Boost for RMAN.

You cannot run RMAN delete or maintenance commands with the database application agent if the commands refer to backups performed with DD Boost for RMAN. [Performing Oracle backup deletion and maintenance operations](#) on page 225 provides details on the backup deletion and maintenance operations with the database application agent.

NOTICE

With DD Boost for RMAN version 1.1 on a UNIX or Linux system, ensure that `LD_LIBRARY_PATH` is either not set or set to the correct library path before you perform an Oracle backup or restore with the database application agent. For example, set the parameter on a Solaris SPARC system:

```
export LD_LIBRARY_PATH=/opt/dpsapps/dbappagent/lib/
sparcv9:$LD_LIBRARY_PATH
```

If you use Net services in the RMAN connection strings, restart Oracle Listener after you change the `LD_LIBRARY_PATH` setting. [Product Installation](#) on page 59 provides details about the library path used by the database application agent on each UNIX and Linux platform.

Updating the RMAN scripts used with DD Boost for RMAN 1.x or later

The RMAN scripts used for the DD Boost for RMAN operations will not work for the database application agent operations because the scripts refer to the DD Boost for RMAN SBT library and include different parameter settings. You must update these RMAN scripts for the database application agent operations.

Ensure that the RMAN scripts used for Oracle operations conform to the guidelines in [Creating the RMAN scripts for DD Boost Oracle operations](#) on page 217.

Complete the following changes in any existing RMAN script used with DD Boost for RMAN.

Procedure

1. Change the `SBT_LIBRARY` parameter setting to the pathname of the Oracle library used by the database application agent.

For example, the following RMAN commands include incorrect `SBT_LIBRARY` settings that refer to the DD Boost for RMAN SBT library:

- On UNIX or Linux:

```
ALLOCATE CHANNEL C1 TYPE SBT_TAPE PARMS 'BLKSIZE=1048576,
SBT_LIBRARY=<${ORACLE_HOME}/lib/libddobk.so';
```

- On Windows:

```
ALLOCATE CHANNEL C1 TYPE SBT_TAPE PARMS 'BLKSIZE=1048576,
SBT_LIBRARY=<%ORACLE_HOME%>\BIN\libDDobk.dll';
```

2. Stop using the following RMAN command to register a Data Domain system:

```
send 'set username user password password servername
dd_hostname';
```

Ensure that you have created a lockbox for the database application agent operations and registered the Data Domain system in the lockbox, as described in the preceding topic.

3. Update the following parameter settings in the RMAN script for backup, recovery, and any RMAN catalog operations:

Note

For each parameter to be updated, delete the existing parameter setting in the RMAN script. It is recommended that you set the new parameter in the configuration file. Alternatively, you can set the new parameter in the RMAN script.

- Delete the `BACKUP_HOST` setting in the RMAN script. Add the corresponding `DEVICE_HOST` setting in the configuration file.
- Delete the `STORAGE_UNIT` setting in the RMAN script. Add the corresponding `DEVICE_PATH` setting in the configuration file.
- Delete the `ORACLE_HOME` setting in the RMAN script. The database application agent does not require this setting.
- Add the `DDBOOST_USER` setting in the configuration file.
- If you use the configuration file, add the `CONFIG_FILE` setting in the RMAN script. Set `CONFIG_FILE` to the pathname of the configuration file.

[Setting up the configuration file in an Oracle environment](#) on page 216 provides details on the required parameter settings in the configuration file.

Using the correct RMAN script for restore operations

To restore Oracle backups performed with DD Boost for RMAN, you must use an RMAN script as described in the DD Boost for RMAN administration guide. The

database application agent cannot restore backups performed with DD Boost for RMAN.

To restore Oracle backups performed with the database application agent, you must use an RMAN script created for the restores, not an RMAN script created for DD Boost for RMAN restores.

Note

Create a copy of the original recovery script used with DD Boost for RMAN 1.x or later, and modify the script copy for the database application agent. You must include the database application agent 4.0 parameters in the script as required to restore from the database application agent backups. You might need to keep the original recovery script for performing restores from the DD Boost for RMAN backups.

For example, you can use the following RMAN script for an Oracle restore with the database application agent:

```
CONFIGURE CHANNEL DEVICE TYPE SBT_TAPE PARMS 'BLKSIZE=1048576,
SBT_LIBRARY=/opt/dpsapps/dbappagent/lib/lib64/libddboostora.so,
SBT_PARMS=(CONFIG_FILE=/orasnb/oracle_ddbda.cfg)';
RESTORE DEVICE TYPE SBT DATABASE;
```

You can use the RMAN command `RESTORE...PREVIEW` to determine whether a restore will require DD Boost for RMAN backups or backups with the database application agent. The command shows you the required backup pieces and which product performed the backup.

In the `RESTORE...PREVIEW` command output, the DD Boost for RMAN backups are indicated by a media handle (after the label `Media:`) that includes only `<device_path>`. The backups with the database application agent are indicated by a media handle that includes `database app agent <device_path>`.

For example, the following `RESTORE...PREVIEW` command output shows the two types of media handles:

```
RMAN> connect target *
2>
3> RUN {
4> ALLOCATE CHANNEL CH1 TYPE 'SBT_TAPE' PARMS='BLKSIZE=1048576,
SBT_LIBRARY=/opt/dpsapps/dbappagent/lib/lib64/libddboostora.so,
ENV=(CONFIG_FILE=/oracle/SNB/ddbda/ddbda.cfg)';
5> RESTORE DATAFILE 2, 3 PREVIEW;
6> RELEASE CHANNEL CH1;
7> }
connected to target database: SNB (DBID=1230476546, not open)
Starting restore at 25-JUN-14
List of Backup Sets
=====
BS Key   Type LV Size           Device Type Elapsed Time Completion Time
-----
3136    Full  8.13G          SBT_TAPE    00:07:27    25-JUN-14
        BP Key: 4651      Status: AVAILABLE Compressed: NO Tag:
TAG20140625T090250
        Handle: 8jpb9tq_1_1   Media: /oracle_rman_plugin
List of Datafiles in backup set 3136
File LV Type Ckp SCN      Ckp Time  Name
-----
2      Full 99730512 07-MAR-14 /clarspace3/oracle/SNB/sapdata1/
undo_1/undo.data1
BS Key   Type LV Size           Device Type Elapsed Time Completion Time
```

```

-----
3137      Full      272.25M      SBT_TAPE      00:00:08      25-JUN-14
          BP Key: 4652      Status: AVAILABLE      Compressed: NO      Tag:
TAG20140625T091332
          Handle: 8kpbnahs_1_1      Media: database app agent /bu-star1
List of Datafiles in backup set 3137
File LV Type Ckp SCN      Ckp Time      Name
-----
3          Full 99730512      07-MAR-14      /clarspace3/oracle/SNB/sapdata1/
sysaux_1/sysaux.data1
-----

```

Performing DD Boost backups and restores with Oracle RMAN

Before you perform an Oracle backup or restore, ensure that you have completed the required configurations from [Configuration of DD Boost operations in an Oracle environment](#) on page 216.

To perform an Oracle backup or restore on the Oracle server host, you can run the appropriate `rman` command at the command line.

Procedure

1. Log in to the Oracle Server host as the Oracle operating system user.
2. To start the RMAN backup or restore script, run the appropriate `rman` command at the command line.

For example, the RMAN backup and restore scripts are stored in the `full_backup.txt` or `restore.txt` files. To connect to the payroll and rcvcatdb databases, you have configured the Net service. You can run the following commands to perform the Oracle backup and restore operations:

```

rman target sys/oracle@payroll rcvcat rman/rman@rcvcatdb
cmdfile \'/disk1/scripts/full_backup.txt\'
rman target sys/oracle@payroll rcvcat rman/rman@rcvcatdb
cmdfile \'/disk1/scripts/restore.txt\'

```

On Windows systems, you can use the `rman.exe` command to run the RMAN script.

Performing DD Boost backups and restores with Oracle Enterprise Manager

Before you perform an Oracle backup or restore, ensure that you have completed the required configurations from [Configuration of DD Boost operations in an Oracle environment](#) on page 216.

The Oracle Enterprise Manager Backup Management Tools provide a graphical user interface to RMAN, which you can use to perform an Oracle backup or restore on the Oracle server host.

Procedure

1. Log in to the Oracle Server host as the Oracle user.
2. To back up or restore Oracle data by using the GUI, run the Oracle Enterprise Manager Backup Management Tools that run the RMAN script. Set

`SBT_LIBRARY` and `CONFIG_FILE` in the **Media Management Library Parameters** text box.

The Backup Management Tools generate the required RMAN scripts and command and perform the backup and restore operations.

The following figure shows an example of a scheduled backup configuration in Oracle Enterprise Manager.

Figure 12 Scheduled backup settings in Oracle Enterprise Manager

ORACLE Enterprise Manager 11g Database Control

Setup Preferences Help Logout Database

Options Settings Schedule Review

Schedule Customized Backup: Settings

Database orcl
Backup Strategy Customized Backup
Object Type Whole Database

Cancel Back Step 2 of 4 Next

Select the destination media for this backup. You can also override the default backup settings.

Disk
Disk Backup Location C:\app\Administrator\flash_recovery_area

Tape
Media Management Vendor SBT_LIBRARY=C:\PROGRA~1\EMC\DDB~1\DA\bin\libDDBoostora.dll, SBT_PARMS=(MMV) Library Parameters (CONFIG_FILE=C:\PROGRA~1\EMC\DDB~1\DA\config\oracle_ddbda.cfg) FORMAT

View Default Settings Override Default Settings
Changed settings will only apply to the current backup.

Return to Schedule Backup Cancel Back Step 2 of 4 Next

Database | Setup | Preferences | Help | Logout

Copyright © 1996, 2010, Oracle. All rights reserved.
Oracle, JD Edwards, PeopleSoft, and Retek are registered trademarks of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.
[About Oracle Enterprise Manager](#)

Performing backups and restores of Oracle CDBs and PDBs

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

It is recommended that you perform regular backups of the whole CDB to ensure that all PDBs and archived logs are backed up. You can recover the whole CDB or individual PDBs to any point-in-time from a CDB backup. The Oracle documentation provides details.

Performing Oracle backup deletion and maintenance operations

The database application agent does not include any expiration policies. As a result, the backups performed by the database application agent remain on the Data Domain system until the DBA or the Oracle software deletes the backups.

Oracle RMAN provides an Oracle retention policy for backups. An Oracle retention policy is based on the recovery window or on redundancy. RMAN considers a backup

to be obsolete when the backup is no longer required according to the Oracle retention policy. Oracle RMAN checks the retention policy of a backup when you run the `REPORT OBSOLETE` or `DELETE OBSOLETE` command.

You can run `DELETE BACKUP` command to delete backups manually. When a deletion is requested by RMAN, the database application agent deletes the catalog entries and the corresponding backup files on the Data Domain system. If the deletion is successful, the Oracle software then deletes the corresponding backup entry in the Oracle catalog. You can force the deletion by using the `FORCE` option in the `DELETE` command. The RMAN documentation provides details about the `DELETE` command and syntax.

If any orphaned entries remain on the Data Domain system for unexpected reasons, such as a crash during the deletion, you can use the `ddbmadmin` command to purge those files as directed by Technical Support.

Note

Although you can run the `CROSSCHECK` and `DELETE EXPIRED` commands, these commands are not useful for backups performed by the database application agent. The database application agent and the Data Domain system do not expire the backups. If you need to run these commands, narrow down the list of backups against which you run these commands, due to their performance impact on the Data Domain system.

Preparing for Oracle disaster recovery

The database application agent and Oracle support disaster recovery to the extent that RMAN supports the functionality. The Oracle and RMAN documentation provides details on the best practices for disaster recovery.

To prepare an Oracle server for disaster recovery, back up the following minimum list of files:

- Oracle database (all the datafiles)
 - Archived redo logs
 - Control file
 - Initialization parameter file
-

Note

RMAN only backs up the server-managed parameter file (`SPFILE`).

The Oracle documentation provides an exhaustive list of all the files (other than the Oracle database) that you must back up. RMAN does not back up the other files that might be required for disaster recovery, such as the Oracle network files, password file, and so on.

Use the following guidelines to prepare for disaster recovery:

- Institute mirrored control files.
Refer to Oracle documentation for recommendations on whether to institute mirrored online redo logs.
- Back up the archived redo logs frequently between database backups.
- Back up the Recovery Catalog after every target database backup if you have a Recovery Catalog.

To perform an Oracle disaster recovery, complete the required steps.

Procedure

1. To create a configuration file, RMAN script, and lockbox on the disaster recovery host, follow the configuration steps in [Configuration of DD Boost operations in an Oracle environment](#) on page 216.
2. In the configuration file, set the `CLIENT` parameter to the hostname used in the backup.

To obtain the list of client names as described in [Using the `ddbadmin` command to display clients for a device path](#) on page 121, if you do not recall the hostname or you did not set `CLIENT` for the backup, use the `ddbadmin -i` command.

3. Follow the disaster recovery instructions in the RMAN documentation.

Oracle RAC and active-passive cluster requirements for DD Boost operations

You can perform Oracle backups and restores with the database application agent in an Oracle RAC or active-passive cluster environment.

You must meet the following configuration requirements in the Oracle RAC or cluster environment:

- All the nodes can access the configuration file, lockbox, and RMAN script through a shared file system or NFS/CIFS share.

Note

If the configuration file or RMAN script is not accessible through a shared file system or NFS/CIFS share, you must copy and maintain an identical configuration file or RMAN script on each node involved in the backups and restores.

This requirement does not apply to stored RMAN scripts because the Oracle software can access a stored RMAN script from any node.

-
- All the cluster hosts are granted the lockbox access as described in [Configuring the lockbox in a high-availability environment](#) on page 115.
 - You have selected one of the node virtual hostnames or the cluster `SCAN` hostname and set the `CLIENT` parameter to that hostname in the configuration file.
 - In an Oracle RAC environment, if the archived redo logs are not accessible from all the nodes (which is not an Oracle best practice), then you have created an RMAN script in which each channel is allocated to connect to a different node.

Oracle RMAN backs up all the archived redo logs of all the nodes if the logs are located on shared storage.

Oracle troubleshooting tips for DD Boost operations

[General troubleshooting tips](#) on page 132 provides common troubleshooting information that applies to the database application agent operations with all the supported databases and applications.

All the operational and error messages are displayed in the RMAN output. For example:

```
ORA-19506: failed to create sequential file, name="42pultir_1_1",  
parms=""  
ORA-27028: skgfcrcr: sbtbackup returned error  
ORA-19511: Error received from media manager layer, error text:  
    DEVICE_HOST is not set correctly. (105:108:2)
```

Use the following information to perform additional troubleshooting:

- Information in the Oracle `sbtio.log` file.
- Oracle operational logs generated by the database application agent, including error, warning, or any other operational messages.

You must set the `DEBUG_LEVEL` parameter to enable debugging.

CHAPTER 8

ProtectPoint Operations on Oracle Systems

This chapter includes the following topics:

- [Overview of ProtectPoint operations in an Oracle environment](#)..... 230
- [Configuration of ProtectPoint operations in an Oracle environment](#)..... 231
- [Performing ProtectPoint backups and restores with Oracle RMAN](#)..... 238
- [Performing ProtectPoint backups and restores with Oracle Enterprise Manager](#)
..... 239
- [Performing backups and restores of Oracle CDBs and PDBs](#)..... 240
- [Performing Oracle backup deletion and maintenance operations](#)..... 240
- [Preparing for Oracle disaster recovery](#)..... 241
- [Oracle RAC and active-passive cluster requirements for ProtectPoint operations](#)
..... 242
- [ProtectPoint restore and rollback for VCS on Solaris](#)..... 242
- [Oracle troubleshooting tips for ProtectPoint operations](#)..... 248

Overview of ProtectPoint operations in an Oracle environment

The database application agent is integrated with Oracle RMAN through the proxy copy option to enable ProtectPoint backups, restores, and recovery. The database application agent also supports deletion and maintenance operations for ProtectPoint Oracle backups.

You can perform a ProtectPoint backup or recovery with the product on an Oracle database server by running one of the supported Oracle backup or recovery tools:

- Oracle Recovery Manager (RMAN) with the `rman` command
- Oracle Enterprise Manager GUI

You can use these tools in cooperation with the database application agent to perform all the operations supported by Oracle RMAN proxy copy, including the following operations:

- Online and offline backups
- Whole and partial database backups
- Archived redo log backups
- Recovery of a database to the current time or a specific point-in-time
- Recovery to the original location or an alternate location
- Backup and recovery of databases, tablespaces, and archived redo logs
- Backup deletion and other maintenance operations

The database application agent maintains a backup catalog on the Data Domain system. During backups, the database application agent creates backup entries in the backup catalog, which provide the information required to restore the backed-up data. RMAN also keeps track of those backups in its own catalog, the RMAN catalog.

The troubleshooting section at the end of this chapter provides details about limitations in the ProtectPoint operations with the database application agent in an Oracle environment.

ProtectPoint Oracle backup processes

A ProtectPoint Oracle backup includes the following process interactions.

1. The database administrator initiates the backup through one of the following methods:

- To invoke the RMAN backup script, the database administrator runs an `rman` command, such as the following `rman` command:

```
rman target /@SNB catalog rman/rman@catdb cmdfile '/orasnb/
backup.txt'
```

- To generate the RMAN backup script and perform the backup operations, the database administrator runs the Oracle Enterprise Manager.
2. The Oracle software loads the Oracle shared library used by the database application agent, as specified by the `SBT_LIBRARY` parameter.

3. The database application agent reads the configuration file specified by the `CONFIG_FILE` parameter, and then initializes the connection with the Data Domain system, based on the settings in the configuration file.
4. The backup workflow proceeds as described in the topic about the ProtectPoint backup workflow or the ProtectPoint with RecoverPoint backup workflow in Chapter 1.

ProtectPoint Oracle restore processes

A ProtectPoint Oracle restore includes the following process interactions.

1. The database administrator initiates the restore through one of the following methods:
 - To invoke the RMAN restore script, the database administrator runs an `rman` command, such as the following `rman` command:

```
rman target /@SNB catalog rman/rman@catdb cmdfile '/orasnb/restore.txt'
```

- To generate the RMAN restore script and perform the restore and recovery operations, the database administrator runs the Oracle Enterprise Manager.
2. The Oracle software loads the Oracle shared library used by the database application agent, as specified by the `SBT_LIBRARY` parameter.
 3. The database application agent reads the configuration file specified by the `CONFIG_FILE` parameter, and then initializes the connection with the Data Domain system, based on the settings in the configuration file.
 4. The restore workflow proceeds as described in the topic about the ProtectPoint restore workflow or the ProtectPoint with RecoverPoint restore workflow in Chapter 1.

ProtectPoint Oracle backups of archived redo logs

Backups of archived redo logs enable recovery of an Oracle database to its predisaster state. Without these backups, you can recover the database only to the time of the last consistent Oracle backup. In this case, you will lose the transactions that occurred between the time of the last consistent backup and the time of the database corruption.

You might want to perform a full database backup every 24 hours at a minimum, and schedule more frequent backups of only the archived redo logs.

You can back up the archived redo logs by using the appropriate option of the RMAN backup command.

Configuration of ProtectPoint operations in an Oracle environment

Ensure that the VMAX, XtremIO, RecoverPoint, and Data Domain configurations have been completed according to the ProtectPoint documentation. The required storage resources must be configured and provisioned properly to enable ProtectPoint operations.

Complete the following tasks to enable ProtectPoint operations:

- Ensure that the `ddbsmd` program is started from the `/opt/dpsapps/dbappagent/bin` directory.
- For ProtectPoint for VMAX operations only, ensure that the supported VMAX Solutions Enabler version is installed and configured in local mode on each production host. The online software compatibility guide at <http://compatibilityguide.emc.com:8080/CompGuideApp/> describes the supported versions.

The Solutions Enabler database must be up-to-date on any host where a backup or recovery might run. To update the Solutions Enabler database, run the `symcfg discover` command. The Solutions Enabler documentation provides details.

Ensure that the required gatekeepers are also configured as described in the *ProtectPoint Version 4.0 Primary and Protection Storage Configuration Guide*. Solutions Enabler uses the small gatekeeper devices for communication with the VMAX storage array.

[Database application agent ProtectPoint operations with Data Domain usage limits](#) on page 44 provides general guidelines on the Data Domain usage limit settings for ProtectPoint operations.

Note

For ProtectPoint backups, it is recommended that database control files and online redo log files be located on different LUNs than the Oracle datafiles and archived logs. The Oracle documentation describes in the best practices for the database file layout.

You must complete the required configurations of the database application agent to enable the ProtectPoint operations in an Oracle environment. The following topics provide the product configuration details.

[Oracle RAC and active-passive cluster requirements for ProtectPoint operations](#) on page 242 provides additional details on the specific configuration requirements in an Oracle RAC or active-passive cluster environment.

The troubleshooting section at the end of this chapter provides details about limitations in the ProtectPoint operations with the database application agent in an Oracle environment.

Setting up the configuration file in an Oracle environment

It is recommended that you set the required parameters for ProtectPoint Oracle operations in the configuration file used by the database application agent.

For example, the configuration file named `oracle_ddbda_vmax.cfg` contains the following parameter settings for ProtectPoint VMAX operations:

```
DDBOOST_USER=qa_ost
DDVDISK_USER=vdisk
DEVICE_HOST=bu-dbe-890.lss.emc.com
DEVICE_PATH=/bu-star1_ora
DEVICE_POOL=IT_data_pool
```


For example, the configuration file named `oracle_dbbda_recoverpoint.cfg` contains the following parameter settings for ProtectPoint with RecoverPoint operations:

```
DDBOOST_USER=qa_ost
DDVDISK_USER=vdisk
DEVICE_HOST=bu-dbe-890.lss.emc.com
DEVICE_PATH=/bu-star1_ora
DEVICE_POOL=IT_data_pool
RP_MGMT_HOST=RPA_management_hostname
RP_USER=RP_username
```

[Setting up the configuration file](#) on page 78 describes the common parameters, ProtectPoint parameters, and how to set the parameters in the configuration file. Other topics in [Product Configuration](#) on page 77 describe the parameters and requirements for the restores of replicated backups and rollback restores.

After the configuration file is set up, ensure that the required lockbox procedures have been performed as described in [Configuring the lockbox](#) on page 103.

Creating the RMAN scripts for ProtectPoint operations

You must create the required RMAN script for the ProtectPoint Oracle backups or restores.

Note

In the RMAN script, the % character is not supported in the `FORMAT` string unless the character is used as part of an RMAN substitution variable.

You must set the `SBT_LIBRARY` and `CONFIG_FILE` parameters, either in the configuration file or in the RMAN script. You must also specify the appropriate `PROXY` option with the `BACKUP` command in the RMAN script:

- Set the `SBT_LIBRARY` parameter to the pathname of the Oracle library used by the database application agent.
- Set the `CONFIG_FILE` parameter to the pathname of the configuration file. Use the correct option if you set `CONFIG_FILE` in the RMAN script:
 - With Oracle 11.2 or later, use the `SBT_PARMS` option.
 - With Oracle 11.1 or earlier, use the `SEND` option.
- Specify the `PROXY` or `PROXY ONLY` option with the `BACKUP` command:
 - When you specify the `PROXY` option, RMAN performs a DD Boost backup instead of a ProtectPoint backup if the backup cannot be completed through the ProtectPoint workflow. The parameter setting `BLKSIZE=1048576` optimizes a DD Boost backup.
 - When you specify the `PROXY ONLY` option, RMAN terminates the backup with a failure if the backup cannot be completed through the ProtectPoint workflow.

RMAN supports ProtectPoint backups through the proxy copy option. The Oracle documentation provides more details about the proxy copy option.

Note

With a `PROXY` option in the `RMAN BACKUP` command, the Oracle software does not support certain additional options, such as `MAXSETSIZE`, `FILESERSET`, and `DISKRATIO`. The Oracle documentation provides details about the `RMAN` options.

The control file and parameter file are automatically backed up when you back up either the whole database (for example, with `BACKUP PROXY...DATABASE`) or a subset that contains the first datafile (for example, with `BACKUP PROXY...TABLESPACE SYSTEM`). The control file and parameter file are always backed up through the DD Boost workflow, even when the `BACKUP` command includes a `PROXY` option.

You do not need to use a `PROXY` option for the restore of a ProtectPoint backup. The software automatically determines the type of backup being restored.

The following examples show the correct parameter settings in the `RMAN` script for a ProtectPoint backup of all the datafiles and archived logs:

- On UNIX or Linux, using the `SBT_PARMS` option with Oracle 11.2 or later:

If you use manual channels (these channel settings are not persistent in the `RMAN` catalog):

```
RUN (
  ALLOCATE CHANNEL C1 DEVICE TYPE SBT_TAPE PARMS
  'SBT_LIBRARY=/opt/dpsapps/dbappagent/lib/lib64/
  libddbboostora.so, SBT_PARMS=(CONFIG_FILE=/orasnb/
  oracle_ddbda.cfg)' FORMAT '%d %U';
  sql 'ALTER SYSTEM SWITCH LOGFILE';
  BACKUP PROXY ONLY DATABASE;
  sql 'ALTER SYSTEM SWITCH LOGFILE';
  BACKUP PROXY ONLY ARCHIVELOG ALL;
}
```

If you use automatic channels:

```
CONFIGURE CHANNEL DEVICE TYPE SBT TAPE PARMS 'SBT_LIBRARY=/opt/
dpsapps/dbappagent/lib/lib64/libddbboostora.so,
SBT_PARMS=(CONFIG_FILE=/orasnb/oracle_ddbda.cfg)';
sql 'ALTER SYSTEM SWITCH LOGFILE';
BACKUP DEVICE TYPE SBT PROXY ONLY DATABASE FORMAT '%d_%U';
sql 'ALTER SYSTEM SWITCH LOGFILE';
BACKUP DEVICE TYPE SBT PROXY ONLY ARCHIVELOG ALL FORMAT '%d_%U';
```

- On Windows, using the `SBT_PARMS` option with Oracle 11.2 or later:

If you use manual channels (these channel settings are not persistent in the `RMAN` catalog):

```
RUN {
  ALLOCATE CHANNEL C1 DEVICE TYPE SBT_TAPE PARMS 'SBT_LIBRARY=C:
  \PROGRA~1\DPSAPPS\DBAPPAGENT\bin\libddbboostora.dll,
  SBT_PARMS=(CONFIG_FILE=D:\orasnb\oracle_ddbda.cfg)' FORMAT '%d_
  %U';
  sql 'ALTER SYSTEM SWITCH LOGFILE';
  BACKUP PROXY ONLY DATABASE;
  sql 'ALTER SYSTEM SWITCH LOGFILE';
  BACKUP PROXY ONLY ARCHIVELOG ALL;
}
```

If you use automatic channels:

```
CONFIGURE CHANNEL DEVICE TYPE SBT_TAPE PARMS 'SBT_LIBRARY=C:
\PROGRA~1\DPSAPPS\DBAPPAGENT\bin\libddboostora.dll,
SBT_PARMS=(CONFIG_FILE=D:\orasnb\oracle_dbdba.cfg)';
sql 'ALTER SYSTEM SWITCH LOGFILE';
BACKUP DEVICE TYPE SBT PROXY ONLY DATABASE FORMAT '%d_%U';
sql 'ALTER SYSTEM SWITCH LOGFILE';
BACKUP DEVICE TYPE SBT PROXY ONLY ARCHIVELOG ALL FORMAT '%d_%U';
```

Note

On Windows, you must use the short Windows pathname in the `SBT_LIBRARY` setting, as shown in the preceding examples. Otherwise, if the pathname contains any spaces, the Oracle software displays a syntax error.

- Using the `SEND` option with Oracle 11.1 or earlier:

If you use manual channels (these channel settings are not persistent in the RMAN catalog):

```
RUN {
ALLOCATE CHANNEL C1 DEVICE TYPE SBT_TAPE PARMS
'SBT_LIBRARY=/opt/dpsapps/dbappagent/lib/lib64/
libddboostora.so' FORMAT '%d_%U';
SEND CHANNEL C1 'ENV=(CONFIG_FILE=/orasnb/oracle_dbdba.cfg)';
sql 'ALTER SYSTEM SWITCH LOGFILE';
BACKUP PROXY ONLY DATABASE;
sql 'ALTER SYSTEM SWITCH LOGFILE';
BACKUP PROXY ONLY ARCHIVELOG ALL;
}
```

If you use automatic channels:

```
CONFIGURE CHANNEL DEVICE TYPE SBT_TAPE PARMS 'SBT_LIBRARY=/opt/
dpsapps/dbappagent/lib/lib64/libddboostora.so';
SEND 'ENV=(CONFIG_FILE=/orasnb/oracle_dbdba.cfg)';
sql 'ALTER SYSTEM SWITCH LOGFILE';
BACKUP DEVICE TYPE SBT PROXY ONLY DATABASE FORMAT '%d_%U';
sql 'ALTER SYSTEM SWITCH LOGFILE';
BACKUP DEVICE TYPE SBT PROXY ONLY ARCHIVELOG ALL FORMAT '%d_%U';
```

The `BACKUP DATABASE PLUS ARCHIVELOG` command is not recommended because the command backs up the archived log LUNs twice, once before the datafile backup starts and once after the datafile backup finishes. This repeated backup can cause performance issues, as compared to a single backup of the archived logs.

The following example shows the commands in an RMAN script for the restore of a ProtectPoint backup. A `PROXY` option is not required in a restore script:

```
ALLOCATE CHANNEL DEVICE TYPE SBT_TAPE PARMS
'SBT_LIBRARY=/opt/dpsapps/dbappagent/lib/lib64/libddboostora.so,
SBT_PARMS=(CONFIG_FILE=/orasnb/oracle_dbdba.cfg)';
RESTORE DATABASE;
RECOVER DATABASE;
```

Allocating multiple channels in the RMAN scripts

The allocation of multiple RMAN channels in the RMAN script does not control the degree of ProtectPoint backup or restore parallelism. The Oracle software uses only

one of the allocated channels for the ProtectPoint backup or restore, unless you use specific backup options to distribute a backup to multiple channels.

The following type of RMAN script is not recommended for ProtectPoint backups:

```
run {
  allocate channel c1 TYPE SBT_TAPE PARMS
    `SBT_LIBRARY=/opt/dpsapps/dbappagent/lib/lib64/libddboostora.so,
    SBT_PARMS=(CONFIG_FILE=/orasnb/oracle_ddbda.cfg)';
  allocate channel c2 TYPE SBT_TAPE PARMS
    `SBT_LIBRARY=/opt/dpsapps/dbappagent/lib/lib64/libddboostora.so,
    SBT_PARMS=(CONFIG_FILE=/orasnb/oracle_ddbda.cfg)';
  backup proxy
    (tablespace tbs1, tbs2 channel c1)
    (tablespace tbs3, tbs4 channel c2);
}
```

Use the following RMAN script to ensure that the ProtectPoint backup succeeds:

```
run {
  allocate channel c1 TYPE SBT_TAPE PARMS
    `SBT_LIBRARY=/opt/dpsapps/dbappagent/lib/lib64/libddboostora.so,
    SBT_PARMS=(CONFIG_FILE=/orasnb/oracle_ddbda.cfg)';
  backup proxy tablespace tbs1, tbs2, tbs3, tbs4;
  release channel c1;
}
```

You can allocate more than one channel in the RMAN script if you know that some of the datafiles or archived logs do not reside on snapshotable devices. In this case, one channel is used for the ProtectPoint backups and all the other channels are used for the DD Boost backups.

Preparing for restore of archived logs

After you perform a number of Oracle backups, the backed-up archived logs might be in multiple sets of static images. During the restore of Oracle archived logs, a separate set of restore devices must be available to mount each set of static images.

Before you start a restore of archived logs, ensure that you have the required number of restore devices in the DD vdisk device pool. This number of restore devices must be at least equal to the number of VMAX or XtremIO source LUNs multiplied by the number of backups, as required by the specific restore.

Note

The restore devices or LUNs do not need to be dedicated to the Oracle server. The restore devices can be in a pool that is also used for other application or server restores that might run at different times, as long as the devices are masked accordingly.

Ensure that you perform a point-in-time restore of archived logs, not a rollback restore.

For example, the backup policy specifies a daily full database backup and the backup of the archived logs four times a day. To enable the restore from a particular database backup and the application of all the required logs, you need four times the number of VMAX or XtremIO source LUNs where the archived logs are located.

Preparing the Data Domain device for restore on Windows

On Windows in a ProtectPoint with RecoverPoint environment, you must prepare the Data Domain vdisk device before you can restore a ProtectPoint with RecoverPoint backup to an XtremIO array.

After the Data Domain block services have been created for the vdisk device according to the ProtectPoint documentation, complete the following steps.

Procedure

1. To bring the device online, use the Disk Manager.
2. If bringing the device online fails because the device is in an unknown state:
 - a. To take the device offline, use the Windows `diskpart` command.
 - b. To bring the device online, use the Disk Manager.

Preparing for Oracle ProtectPoint with RecoverPoint backups and rollback restores that use RecoverPoint pre-5.0

With RecoverPoint pre-5.0, the database application agent performs a rollback restore of a ProtectPoint with RecoverPoint backup at the consistency group level. If the RecoverPoint consistency group being restored contains multiple LUNs, then all those LUNs are overwritten and inaccessible during the rollback restore. Specific requirements apply to the Oracle ProtectPoint with RecoverPoint backups and rollback restores.

Ensure that you follow the requirements and recommendations in [Configuring rollback restores of ProtectPoint backups](#) on page 97.

Note

With RecoverPoint pre-5.0, a ProtectPoint with RecoverPoint backup and rollback restore always occurs at the consistency group level, regardless of which objects are included in the backup command. As a best practice for the ProtectPoint with RecoverPoint rollback restore, when you perform the backup or rollback restore, do not exclude the logs or any database files that are part of the RecoverPoint consistency group being backed up or restored.

Ensure that you meet the following requirements for the rollback restore of an Oracle ProtectPoint with RecoverPoint backup with RecoverPoint pre-5.0:

- The Oracle control files are in a different RecoverPoint consistency group than the consistency group that is included in the rollback restore. Alternatively, the Oracle control files reside on conventional nonsnapshot devices. Oracle does not support a snapshot (proxy) backup of the control files. Oracle accesses the control files during the rollback restore.
- If an Oracle tablespace must remain online during a rollback restore, then this online tablespace and the tablespace to be restored are located in different RecoverPoint consistency groups.
- The online redo log files are in different devices than the consistency groups that are included in the rollback restore.
- If an Oracle tablespace has a datafile that is located on the LUNs of a RecoverPoint consistency group, then before the rollback restore of the consistency group, the tablespace is turned offline to prevent failure or data corruption.

It is recommended that you allocate the Oracle datafiles and archived log files in separate RecoverPoint consistency groups. This recommendation is for the possible case where you need to perform a rollback restore for only the datafiles or only the archived log files.

Configuring operations in an Oracle Data Guard environment

The database application agent supports Oracle Data Guard, which is an Oracle data availability and protection solution for a primary database and one or more standby databases over an IP network. You can configure backup and restore operations with the database application agent in an Oracle Data Guard environment.

In an Oracle Data Guard environment, as transactions occur in the primary database and as Oracle writes redo data to the local redo logs, Data Guard automatically performs the following operations:

- Transfers this redo data to the standby sites.
- Applies the redo data to the standby databases, which synchronizes the standby databases with the primary database.

You can offload RMAN backups of datafiles, archived redo logs, and possibly other files to a physical standby database. You can then use the backups to recover the primary or standby database. RMAN and Data Guard documentation describes how to configure and back up a physical standby database, and use the backups to recover the primary or standby database.

To configure backups and restores with the database application agent in an Oracle Data Guard environment:

1. Follow the instructions in the Oracle documentation about how to set the required RMAN configurations, for example, to use a Recovery Catalog and the `DB_UNIQUE_NAME` parameter.
2. Install and configure the database application agent software on the primary database host, and then on each physical standby database host that is included in the backups and restores.
3. For a backup, create an RMAN script, and then set the parameters in the configuration file for the database application agent to back up data from a physical standby database, which can be used to restore the primary database. Set the `CLIENT` parameter in the configuration file to a single value that identifies the Data Guard environment, preferably the primary database hostname.
4. For a recovery, create an RMAN script, and then set the parameters in the configuration file for the database application agent to recover the data from a primary or standby database, depending on the restore. Set the `CLIENT` parameter to the same value as used during the backup.

Performing ProtectPoint backups and restores with Oracle RMAN

Before you perform a ProtectPoint Oracle backup or restore, ensure that you have completed the required configurations from [Configuration of ProtectPoint operations in an Oracle environment](#) on page 231.

To perform the backup or restore on the Oracle server host, you can run the appropriate `rman` command at the command line.

Procedure

1. Log in to the Oracle Server host as the Oracle operating system user.
2. To start the RMAN backup or restore script, run the appropriate `rman` command at the command line.

For example, the RMAN backup and restore scripts are stored in the `full_backup.txt` or `restore.txt` files. To connect to the payroll and rcvcatdb databases, you have configured the Net service. You can run the following commands to perform the Oracle backup and restore operations:

```
rman target sys/oracle@payroll rcvcat rman/rman@rcvcatdb
cmdfile \'/disk1/scripts/full_backup.txt\'
rman target sys/oracle@payroll rcvcat rman/rman@rcvcatdb
cmdfile \'/disk1/scripts/restore.txt\'
```

On Windows systems, you can use the `rman.exe` command to run the RMAN script.

Performing ProtectPoint backups and restores with Oracle Enterprise Manager

Before you perform a ProtectPoint Oracle backup or restore, ensure that you have completed the required configurations from [Configuration of ProtectPoint operations in an Oracle environment](#) on page 231.

The Oracle Enterprise Manager Backup Management Tools provide a graphical user interface to RMAN, which you can use to perform the backup or restore on the Oracle server host.

Procedure

1. Log in to the Oracle Server host as the Oracle user.
2. To back up or restore Oracle data by using the GUI, run the Oracle Enterprise Manager Backup Management Tools that run the RMAN script. Set `SBT_LIBRARY` and `CONFIG_FILE` in the **Media Management Library Parameters** text box.

The Backup Management Tools generate the required RMAN scripts and command and perform the backup and restore operations.

The following figure shows an example of a scheduled backup configuration in Oracle Enterprise Manager.

Figure 13 Scheduled backup settings in Oracle Enterprise Manager

ORACLE Enterprise Manager 11g Database Control

Setup Preferences Help Logout

Options Settings Schedule Review

Schedule Customized Backup: Settings

Database orcl
Backup Strategy Customized Backup
Object Type Whole Database

Cancel Back Step 2 of 4 Next

Select the destination media for this backup. You can also override the default backup settings.

Disk
Disk Backup Location C:\app\Administrator\flash_recovery_area

Tape
Media Management Vendor SBT_LIBRARY=C:\PROGRA~1\EMC\DDB~1\DA\bin\libDDBboostora.dll, SBT_PARMS=(MMV) Library Parameters (CONFIG_FILE=C:\PROGRA~1\EMC\DDB~1\DA\config\oracle_ddbda.cfg) F0RMAT

View Default Settings Override Default Settings
Changed settings will only apply to the current backup.

Return to Schedule Backup Cancel Back Step 2 of 4 Next

Database | Setup | Preferences | Help | Logout

Copyright © 1996, 2010, Oracle. All rights reserved.
Oracle, JD Edwards, PeopleSoft, and Retek are registered trademarks of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.
[About Oracle Enterprise Manager](#)

Performing backups and restores of Oracle CDBs and PDBs

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

It is recommended that you perform regular backups of the whole CDB to ensure that all PDBs and archived logs are backed up. You can recover the whole CDB or individual PDBs to any point-in-time from a CDB backup. The Oracle documentation provides details.

Performing Oracle backup deletion and maintenance operations

The database application agent does not include any expiration policies. As a result, the backups performed by the database application agent remain on the Data Domain system until the DBA or the Oracle software deletes the backups.

Oracle RMAN provides an Oracle retention policy for backups. An Oracle retention policy is based on the recovery window or on redundancy. RMAN considers a backup to be obsolete when the backup is no longer required according to the Oracle retention policy. Oracle RMAN checks the retention policy of a backup when you run the `REPORT OBSOLETE` or `DELETE OBSOLETE` command.

You can run `DELETE BACKUP` command to delete backups manually. When a deletion is requested by RMAN, the database application agent deletes the catalog entries and the corresponding backup files on the Data Domain system. If the deletion is

successful, the Oracle software then deletes the corresponding backup entry in the Oracle catalog. You can force the deletion by using the `FORCE` option in the `DELETE` command. The RMAN documentation provides details about the `DELETE` command and syntax.

If any orphaned entries remain on the Data Domain system for unexpected reasons, such as a crash during the deletion, you can use the `ddbmadmin` command to purge those files as directed by Technical Support.

Note

Although you can run the `CROSSCHECK` and `DELETE EXPIRED` commands, these commands are not useful for backups performed by the database application agent. The database application agent and the Data Domain system do not expire the backups. If you need to run these commands, narrow down the list of backups against which you run these commands, due to their performance impact on the Data Domain system.

Preparing for Oracle disaster recovery

The database application agent and Oracle support disaster recovery with the proxy copy option to the extent that RMAN supports the functionality. The Oracle and RMAN documentation provides details on the best practices for disaster recovery.

To prepare an Oracle server for disaster recovery, back up the following minimum list of files:

- Oracle database (all the datafiles)
- Archived redo logs
- Control file
- Initialization parameter file

Note

RMAN only backs up the server-managed parameter file (`SPFILE`).

The Oracle documentation provides an exhaustive list of all the files (other than the Oracle database) that you must back up. RMAN does not back up the other files that might be required for disaster recovery, such as the Oracle network files, password file, and so on.

Use the following guidelines to prepare for disaster recovery:

- Institute mirrored control files.
Refer to Oracle documentation for recommendations on whether to institute mirrored online redo logs.
- Back up the archived redo logs frequently between database backups.
- Back up the Recovery Catalog after every target database backup if you have a Recovery Catalog.

To perform an Oracle disaster recovery, complete the required steps.

Procedure

1. Ensure that the restore LUNs have been provisioned to the disaster recovery host. The ProtectPoint documentation provides details.

2. To create a configuration file, RMAN script, and lockbox on the disaster recovery host, follow the configuration steps in [Configuration of ProtectPoint operations in an Oracle environment](#) on page 231.
3. In the configuration file, set the `CLIENT` parameter to the hostname used in the backup.

To obtain the list of client names as described in [Using the `ddbadmin` command to display clients for a device path](#) on page 121, if you do not recall the hostname or you did not set `CLIENT` for the backup, use the `ddbadmin -i` command.

4. Follow the disaster recovery instructions in the RMAN documentation.

Oracle RAC and active-passive cluster requirements for ProtectPoint operations

You can perform ProtectPoint backups and restores with the database application agent in an Oracle RAC or active-passive cluster environment.

You must meet the following configuration requirements in the Oracle RAC or cluster environment:

- All the nodes can access the configuration file, lockbox, and RMAN script through a shared file system or NFS/CIFS share.

Note

If the configuration file or RMAN script is not accessible through a shared file system or NFS/CIFS share, you must copy and maintain an identical configuration file or RMAN script on each node involved in the backups and restores.

This requirement does not apply to stored RMAN scripts because the Oracle software can access a stored RMAN script from any node.

- All the cluster hosts are granted the lockbox access as described in [Configuring the lockbox in a high-availability environment](#) on page 115.
- You have selected one of the node virtual hostnames or the cluster `SCAN` hostname and set the `CLIENT` parameter to that hostname in the configuration file.
- In an Oracle RAC environment, if the archived redo logs are not accessible from all the nodes (which is not an Oracle best practice), then you have created an RMAN script in which each channel is allocated to connect to a different node.

Oracle RMAN backs up all the archived redo logs of all the nodes if the logs are located on shared storage.

- All the restore LUNs are provisioned correctly to all the nodes as the restore can be started from any cluster node.

ProtectPoint restore and rollback for VCS on Solaris

Use the procedures in the following topics to perform ProtectPoint restore and rollback operations for a VCS system on Solaris.

Performing a ProtectPoint VCS restore

Procedure

1. On the primary VCS node, perform the following steps as the root user.
 - a. List the VCS Service Groups:

```
root:/# hastatus -sum
```

-- SYSTEM STATE			
	System	State	Frozen
A	ledma054	RUNNING	0
A	ledma056	RUNNING	0
-- GROUP STATE			
	Group	System	Probed
AutoDisabled State			
B	ClusterService	ledma054	Y
N		ONLINE	
B	ClusterService	ledma056	Y
N		OFFLINE	
B	oracle_ctl_sg	ledma054	Y
N		ONLINE	
B	oracle_ctl_sg	ledma056	Y
N		OFFLINE	
B	oracle_sg	ledma054	Y
N		ONLINE	
B	oracle_sg	ledma056	Y
N		OFFLINE	
B	vxfen	ledma054	Y
N		ONLINE	
B	vxfen	ledma056	Y
N		ONLINE	

- b. Enable the VCS configuration as Read/Write:

```
root:/# haconf -makerw
```

- c. Freeze the VCS service groups by disabling On line/Off line. Type the following command:

Note

This is an example of a VCS and Oracle configuration.

```
root:/# hagrpf -freeze <oracle_sg> -persistent
```

- d. Confirm the VCS status by typing the following command:

```
root:/# hastatus -sum
```

-- SYSTEM STATE			
	System	State	Frozen
A	ledma054	RUNNING	0

```

A ledma056          RUNNING          0
-- GROUP STATE
-- Group           System           Probed
AutoDisabled      State
B ClusterService  ledma054          Y
N                  ONLINE
B ClusterService  ledma056          Y
N                  OFFLINE
B oracle_ctl_sg   ledma054          Y
N                  ONLINE
B oracle_ctl_sg   ledma056          Y
N                  OFFLINE
B oracle_sg       ledma054          Y
N                  ONLINE
B oracle_sg       ledma056          Y
N                  OFFLINE
B vxfen           ledma054          Y
N                  ONLINE
B vxfen           ledma056          Y
N                  ONLINE

-- GROUPS FROZEN
-- Group
C oracle_ctl_sg
C oracle_sg

-- RESOURCES DISABLED
-- Group           Type           Resource
H oracle_ctl_sg   DiskGroup      oracle_ctl_dg_DG_res1
H oracle_ctl_sg   Mount          oracle_ctl_dg_MNT_res1
H oracle_ctl_sg   Volume        oracle_ctl_dg_VOL_res1
H oracle_sg       DiskGroup      oracle_dg_DG_res1
H oracle_sg       Mount          oracle_dg_MNT_res1
H oracle_sg       Volume        oracle_dg_VOL_res1

```

e. Make the VCS configuration as Read Only. Type the following command:

```
root:/# haconf -dump -makero:
```

2. On the primary VCS node, perform the following steps as the Oracle user.

a. Run the `shutdown` and `startup mount` commands on the Oracle database:

```
a. oracle:/# sqlplus / as sysdba
```

```
b. SQL > shutdown immediate
```

```
c. SQL > startup mount
```

```
d. SQL > exit
```

b. Perform the RMAN restore and recovery.

3. On the primary VCS node, perform the following steps as the root user.

a. Make the VCS configuration Read/Write. Type the following command:

```
root:/# haconf -makerw
```

- b. Unfreeze the service groups, and allow On line/Off line. Type the following command:

```
root:/# hagr -unfreeze <oracle_ctl_sg> -persistent
root:/# hagr -unfreeze <oracle_sg> -persistent
```

- c. Confirm the VCS status. Type the following command:

```
root:/# hastatus -sum
```

```
-- SYSTEM STATE
-- System          State          Frozen
A ledma054         RUNNING       0
A ledma056         RUNNING       0

-- GROUP STATE
-- Group           System          Probed
AutoDisabled      State
B ClusterService ledma054       Y
N                  ONLINE
B ClusterService ledma056       Y
N                  OFFLINE
B oracle_ctl_sg   ledma054       Y
N                  ONLINE
B oracle_ctl_sg   ledma056       Y
N                  OFFLINE
B oracle_sg       ledma054       Y
N                  ONLINE
B oracle_sg       ledma056       Y
N                  OFFLINE
B vxfen           ledma054       Y
N                  ONLINE
B vxfen           ledma056       Y
N                  ONLINE
```

Performing a ProtectPoint VCS rollback

Note

A rollback fails if you change the style of the mpio device name. The rollback to the source LUN is successful. However, the fsck and mount fails. In this scenario, manually mount the FS.

Procedure

1. On the primary VCS node, perform the following steps as the root user.
 - a. List the VCS Service Groups:

```
root:/# hastatus -sum
```

```
-- SYSTEM STATE
-- System          State          Frozen
A ledma054         RUNNING       0
A ledma056         RUNNING       0

-- GROUP STATE
```

```

-- Group          System          Probed
AutoDisabled    State
B ClusterService ledma054        Y
N                ONLINE
B ClusterService ledma056        Y
N                OFFLINE
B oracle_ctl_sg  ledma054        Y
N                ONLINE
B oracle_ctl_sg  ledma056        Y
N                OFFLINE
B oracle_sg      ledma054        Y
N                ONLINE
B oracle_sg      ledma056        Y
N                OFFLINE
B vxfen          ledma054        Y
N                ONLINE
B vxfen          ledma056        Y
N                ONLINE

```

b. Enable the VCS configuration as Read/Write:

```
root:/# haconf -makerw
```

c. Freeze the VCS service groups by disabling On line/Off line. Type the following command:

Note

This is an example of a VCS and Oracle configuration.

```

root:/# hagrps -freeze <oracle_sg> -persistent
root:/# hagrps -freeze <oracle_ctl_sg> -persistent

```

d. Confirm the VCS status, by typing the following command:

```
root:/# hastatus -sum
```

```

-- SYSTEM STATE
-- System          State          Frozen
A ledma054        RUNNING        0
A ledma056        RUNNING        0

-- GROUP STATE
-- Group          System          Probed
AutoDisabled    State
B ClusterService ledma054        Y
N                ONLINE
B ClusterService ledma056        Y
N                OFFLINE
B oracle_ctl_sg  ledma054        Y
N                ONLINE
B oracle_ctl_sg  ledma056        Y
N                OFFLINE
B oracle_sg      ledma054        Y
N                ONLINE
B oracle_sg      ledma056        Y
N                OFFLINE
B vxfen          ledma054        Y
N                ONLINE

```

```

B vxfen          ledma056          Y
N                ONLINE

-- GROUPS FROZEN
-- Group
C oracle_ctl_sg
C oracle_sg

-- RESOURCES DISABLED
-- Group          Type          Resource
H oracle_ctl_sg  DiskGroup  oracle_ctl_dg_DG_res1
H oracle_ctl_sg  Mount      oracle_ctl_dg_MNT_res1
H oracle_ctl_sg  Volume     oracle_ctl_dg_VOL_res1
H oracle_sg      DiskGroup  oracle_dg_DG_res1
H oracle_sg      Mount      oracle_dg_MNT_res1
H oracle_sg      Volume     oracle_dg_VOL_res1

```

e. Make the VCS configuration as Read Only. Type the following command:

```
root:/# haconf -dump -makero:
```

2. On the primary VCS node, perform the following steps as the Oracle user.

a. Run the `shutdown` and `startup mount` commands on the Oracle database:

```
a. oracle:/# sqlplus / as sysdba
```

```
b. SQL > shutdown immediate
```

```
c. SQL > startup mount
```

```
d. SQL > exit
```

b. Perform the RMAN rollback and recovery.

3. On the primary VCS node, perform the following steps as the root user.

a. Make the VCS configuration Read/Write. Type the following command:

```
root:/# haconf -makerw
```

b. Unfreeze the service groups, and allow On line or Off line. Type the following command:

```
root:/# hagrps -unfreeze <oracle_ctl_sg> -persistent
root:/# hagrps -unfreeze <oracle_sg> -persistent
```

c. Confirm the VCS status. Type the following command:

```
root:/# hastatus -sum
```

```

-- SYSTEM STATE
-- System          State          Frozen

```

```

A ledma054          RUNNING          0
A ledma056          RUNNING          0

-- GROUP STATE
-- Group            System            Probed
AutoDisabled      State
B ClusterService  ledma054          Y
N                  ONLINE
B ClusterService  ledma056          Y
N                  OFFLINE
B oracle_ctl_sg   ledma054          Y
N                  ONLINE
B oracle_ctl_sg   ledma056          Y
N                  OFFLINE
B oracle_sg       ledma054          Y
N                  ONLINE
B oracle_sg       ledma056          Y
N                  OFFLINE
B vxfen           ledma054          Y
N                  ONLINE
B vxfen           ledma056          Y
N                  ONLINE

```

Note

The service groups will be faulted, but will come back online in a short time.

Oracle troubleshooting tips for ProtectPoint operations

[General troubleshooting tips](#) on page 132 provides common troubleshooting information that applies to the database application agent operations with all the supported databases and applications.

All the operational and error messages are displayed in the RMAN output. For example:

```

ORA-19506: failed to create sequential file, name="42pultir_1_1",
parms=""
ORA-27028: skgfgcre: sbtbackup returned error
ORA-19511: Error received from media manager layer, error text:
  DEVICE_HOST is not set correctly. (105:108:2)

```

Use the following information to perform additional troubleshooting:

- Information in the Oracle `sbtio.log` file.
- Oracle operational logs generated by the database application agent, including error, warning, or any other operational messages.

ProtectPoint Oracle operations maintain a separate operational log named `ddbsm.log`, which has details about operations and errors.

You must set the `DEBUG_LEVEL` parameter to enable debugging.

Oracle rollback restore to a new database might fail when OMF is enabled

When the Oracle-Managed Files (OMF) database feature is enabled, a rollback restore to a new database might fail.

For example, when you perform a redirected rollback restore to alternate LUNs by using a ProtectPoint for VMAX backup of an Oracle OMF database, the restore might fail with the following error message:

```
ORA-19511: non RMAN, but media manager or vendor specific failure,  
error text:  
A rollback is not possible when doing relocation during a restore.  
Please remove 'rollback' from the RESTORE_TYPE_ORDER parameter or do  
not request relocation. (114:123:2)
```

As a workaround, disable the OMF feature after you restore the spfile of the database and before you restore the control file and data files.

CHAPTER 9

DD Boost Operations on SAP HANA Systems

This chapter includes the following topics:

- [Overview of DD Boost operations in an SAP HANA environment](#)..... 252
- [Configuration of DD Boost operations in an SAP HANA environment](#).....254
- [Performing DD Boost backups, recovery, and deletion with SAP HANA Studio](#)
.....259
- [Performing DD Boost backups and recovery with SAP HANA CLI](#).....265
- [Preparing for SAP HANA disaster recovery](#).....267
- [SAP HANA scale-out requirements for DD Boost operations](#)..... 268
- [SAP HANA troubleshooting tips for DD Boost operations](#)..... 269

Overview of DD Boost operations in an SAP HANA environment

An SAP HANA database holds most of the data in memory but also uses persistent storage on disk. During normal database operations, the data is automatically saved from memory to the disk at regular intervals. All the data changes are also captured in the redo log on disk, which is updated after each committed database transaction. The data on disk must be backed up to ensure protection against a disk failure.

The database application agent is integrated with the SAP HANA `BACKINT` interface to enable SAP HANA database data and redo log backups and restores.

You can perform a backup, recover, inquire, or delete operation with the product on an SAP HANA database server by running one of the supported SAP HANA tools:

- SAP HANA command line interface (CLI) with the `hdbsql` command
- SAP HANA Studio GUI
- SAP DBA Cockpit in The Computing Center Management System (CCMS) GUI

You can use these tools in cooperation with the database application agent to perform the following SAP HANA operations on the single database containers and the multitenant database containers:

- Online backups
- Full database backups
- Redo log backups
- Delta (incremental and differential) backups
- Scheduled full backups and delta backups

Note

SAP HANA SPS 12 or later supports the scheduled backups.

- Recovery of a database to the most recent state, a specific point-in-time, a specific data backup, or a log position
- Recovery to the original host or an alternate host

Note

SAP HANA SPS 09 or later enables you to perform a redirected recovery by using a different SID. The SAP HANA versions that are earlier than SPS 09 enable you to perform a redirected recovery by using the same SID

SAP HANA requires that restore and recovery is performed on a Linux system with the same architecture as the backup system:

- In a replication environment, both the primary and secondary systems must be either Linux x64 or Linux Power PC systems.
- An SAP HANA backup that is performed on a Linux x64 system can be restored only to a Linux x64 system.
- An SAP HANA backup that is performed on a Linux Power PC system can be restored only to a Linux Power PC system.

The troubleshooting section at the end of this chapter provides details about limitations in the DD Boost operations with the database application agent in an SAP HANA environment.

The product maintains online backup indexes on the Data Domain system. During backups, the product creates backup entries in the online indexes, which provide the information required to restore the backed-up data.

SAP HANA backup processes

An SAP HANA database backup includes the following process interactions.

1. The database administrator initiates the backup by running the `hdbsql` command, SAP HANA Studio GUI, or SAP DBA Cockpit in CCMS.
2. The SAP HANA database server runs the `hdbbackint` program, installed as part of the database application agent, and passes a list of pipes to back up.
3. The `hdbbackint` program processes the SAP HANA parameters from the configuration file and starts the child `hdbbackint` processes that back up the required data.
4. The child `hdbbackint` processes send the database data and tracking information to the Data Domain system for storage.

SAP HANA restore processes

An SAP HANA database restore includes the following process interactions.

1. The database administrator initiates the restore by running the SAP HANA Studio GUI.
2. The SAP HANA database server runs the `hdbbackint` program, and then passes a list of pipes to receive the restored data.
3. The `hdbbackint` program processes the SAP HANA parameters from the configuration file, and then starts the child `hdbbackint` processes that restore the required data.
4. The child `hdbbackint` processes perform the following tasks:
 - a. Query the index on the Data Domain system to locate the backup data.
 - b. Retrieve the backup data from the Data Domain system.
 - c. Write the backup data to files on the SAP HANA database server.

SAP HANA backups of redo logs

The SAP HANA database server automatically performs periodic backups of the redo logs for a database. You do not run the log backups with any backup commands as is the case for database backups.

You can use the SAP HANA Studio to configure the frequency of automatic log backups and whether the `BACKINT` interface is used for the log backups as described in [Configuring automatic backups of SAP HANA redo logs](#) on page 257.

Configuration of DD Boost operations in an SAP HANA environment

You must complete the required configurations to enable the DD Boost operations in an SAP HANA environment. The following topics provide the product configuration details.

The troubleshooting section at the end of this chapter provides details about limitations in the DD Boost operations with the database application agent in an SAP HANA environment.

Integrating the product into the SAP HANA environment

The database application agent installation places the `hdbbackint` program in the `/opt/dpsapps/dbappagent/bin` directory. However, SAP HANA requires the program to be accessible from the `/usr/sap/<SID>/SYS/global/hdb/opt` directory.

Note

You must manually create the `opt` subdirectory if it does not exist in the `/usr/sap/<SID>/SYS/global/hdb` directory.

You can either copy the `hdbbackint` file to that directory or create a symbolic link that points from `/usr/sap/<SID>/SYS/global/hdb/opt/hdbbackint` to the actual executable file.

Ensure that the `hdbbackint` file has the required executable permissions for the operating system user `<SID>adm` to run the program, where `<SID>` is the system ID of the SAP HANA system.

Configuring the SAP HANA parameters

You must set the required parameters for SAP HANA operations in the configuration file used by the database application agent.

[Setting up the configuration file](#) on page 78 describes the common parameters and how to set parameters in the configuration file. [Configuring restores of replicated backups](#) on page 90 also describes the parameters and requirements for the restores of replicated backups.

You can optionally set the `PARALLELISM` parameter to specify a multistream backup or restore as described in the following table.

You must set the `PARALLELISM` parameter in the general section of the configuration file, which has the `[GENERAL]` section heading.

After the configuration file is set up, ensure that the required lockbox procedures have been performed as described in [Configuring the lockbox](#) on page 103.

Table 23 SAP HANA parallelism parameter

<p>Parameter: PARALLELISM</p>

<p>Section: [GENERAL]</p>

Table 23 SAP HANA parallelism parameter

Specifies the maximum number of concurrent data streams to send to or from the Data Domain system during a backup or restore, for each `hdbbackint` program that SAP HANA runs for the backup or restore.

SAP HANA SPS 11 introduced the SAP HANA parameter `parallel_data_backup_backint_channels` that enables SAP HANA to split a data backup into multiple channels. The database application agent saves each SAP HANA channel as a separate save set, and each save set uses a single data stream. The database application agent parameter, `PARALLELISM`, limits the maximum number of concurrent data streams, and new streams start when other streams complete. [Streams usage on SAP HANA systems](#) on page 258 provides more details.

The `parallel_data_backup_backint_channels` parameter does not apply to a restore. A restore uses the same number of streams as were used during the backup.

Note

Prior to SAP HANA SPS 11, an SAP HANA backup always has a parallelism of 1 per `hdbbackint` process. If the SAP HANA parameter `parallel_data_backup_backint_channels` is set in SAP HANA SPS 11 or later, the default value of this `PARALLELISM` parameter is 8.

Optional for a restore.

With SAP HANA SPS 11 or later, optional for a backup.

Valid values:

- 8 (default).
- Positive integer number.

Configuring support of SAP HANA 2.0 SPS 00

The database application agent supports SAP HANA 2.0 SPS 00. The SAP HANA 2.0 user documentation provides a complete list of all the backup and recovery features, including details about how to enable and configure the features. The following list highlights several of the new features:

- SAP HANA 2.0 SPS 00 supports data encryption in the persistence layer for both data and log volumes. SAP HANA volume encryption can impact the deduplication rates on the Data Domain systems.
- SAP HANA 2.0 SPS 00 supports the user-configurable interval mode for log backups. The interval mode enables the creation of log backups after the service-specific timeout is reached, instead of when the log segment becomes full. A log backup can include multiple log segments. The interval mode setting limits the number of log backups in high-transaction databases, which reduces the impact on the databases. You can configure the interval mode by setting the following parameter in the SAP HANA `global.ini` file:

`log_backup_interval_mode = service`
- SAP HANA 2.0 requires that the backup catalog backups and transaction log backups be configured separately. In previous SAP HANA releases, the backup catalog was automatically saved as part of the transaction log backups. Ensure that both the backup catalog backups and transaction log backups are configured to use the database application agent through the Backint interface.

You can improve the index query performance by specifying different device paths for the catalog and log backups in separate parameter files. This setting reduces the number of indexes that must be searched during specific inquire or restore operations. For example:

- The device path in the backup catalog parameter file can be set as `DEVICE_PATH=/ <storage_unit>/ <SID>/CATALOG`.
- The device path in the log parameter file can be set as `DEVICE_PATH=/ <storage_unit>/ <SID>/LOGS`.

Enabling the configuration file in SAP HANA Studio

You must specify the location of the configuration file in SAP HANA Studio. On the **Configuration** tab of the **Backup** editor, type the complete pathname of the configuration file in the **Backint Parameter File** text box.

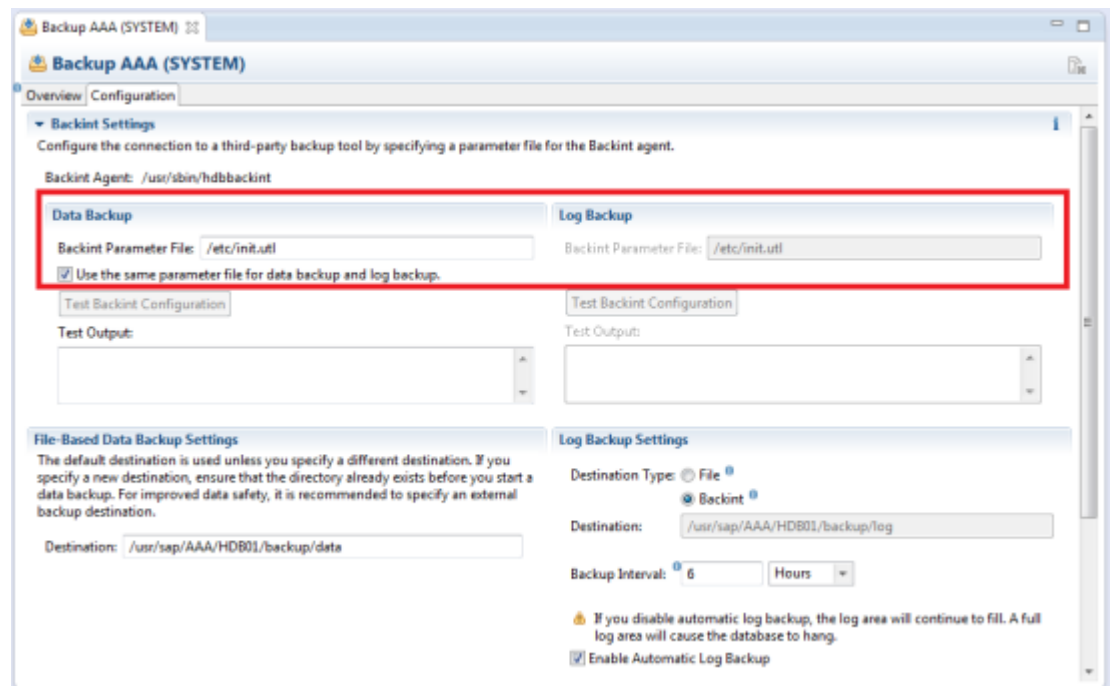
Note

The same configuration file is used for all the SAP HANA CLI and SAP HANA Studio operations.

You can specify separate configuration files for the database backup and log backup. To use the same configuration file for both types of backups, you can select **Use the same parameter file for data backup and log backup**.

The following figure shows a configuration example in SAP HANA Studio 1.0 SPS 5.

Figure 14 Specifying the configuration file in SAP HANA Studio



Configuring automatic backups of SAP HANA redo logs

To configure the automatic backups of SAP HANA redo logs, you must complete the required steps in SAP HANA Studio.

On the **Configuration** tab of the **Backup** editor, complete the settings in the **Log Backup Settings** group box:

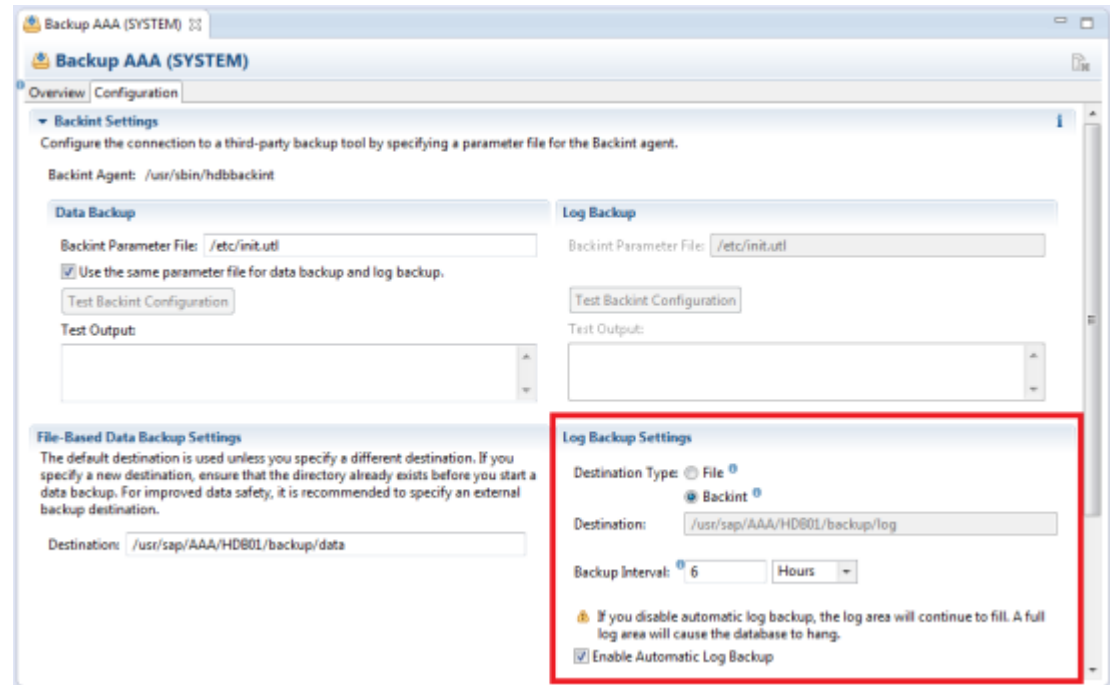
- Select **Enable Automatic Log Backup**.
- For **Destination Type**, select **Backint**.
- For **Backup Interval**, specify the log backup interval. Type the numeric value in the text box, and then select the unit of time, such as **Minutes**, from the menu.

Note

The recommended log backup interval is 30 minutes.

The following figure shows a configuration example in SAP HANA Studio 1.0 SPS 5.

Figure 15 Configuring automatic log backups in SAP HANA Studio



Estimating the Data Domain resource usage on SAP HANA systems

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

Capacity usage on SAP HANA systems

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

```
163542:hdbbackint: Unable to write to a file due to reaching the hard
quota limit.
The error message is: [5194] [ 8920] [139899242542880] Thu Dec  8
12:48:02 2016
      ddp_write() failed Offset 905969664, BytesToWrite 524288,
BytesWritten 0 Err: 5194-Hard Quota Exceeded >
```

Streams usage on SAP HANA systems

Note

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

For SAP HANA, the database application agent requires one Data Domain stream for each backed-up pipe. For example, if an SAP HANA scale-out system has 12 running services, then 12 streams are required to back up the data. Starting with SAP HANA SPS 09, each service can also back up multiple logs for each backup, as controlled by the database parameter `max_log_backup_group_size`.

For a multistream backup starting with SAP HANA SPS 11, the database application agent can use multiple SAP HANA channels to write the backup data for each service. The database application agent uses a separate SAP HANA channel to write each stream of data to the Data Domain system. To specify the number of channels to use for the backup, up to a maximum of 32 channels, you can set the SAP HANA parameter `parallel_data_backup_backint_channels`. SAP HANA opens the corresponding number of pipe files for the backup, and the database application agent saves each stream as a separate save set.

To specify the maximum number of concurrent backup or restore streams, you set the `PARALLELISM` parameter in the `hdbbackint` configuration file. For example, if the `parallel_data_backup_backint_channels` parameter is set to 12 on the SAP HANA server, then 12 streams are used for the backup, which produces 12 save sets. If the `hdbbackint PARALLELISM` parameter is set to 6, then a maximum of 6 streams are backed up concurrently, and new streams start as other streams complete.

A restore uses the same number of streams as the backup, and ignores the `parallel_data_backup_backint_channels` parameter setting.

The SAP HANA storage unit typically uses the following number of streams during a backup and restore:

- If `PARALLELISM` is set in the `hdbbackint` configuration file:
Number of services x `PARALLELISM`
- If `PARALLELISM` is not set:
Number of services x `max_log_backup_group_size`

Due to the design of SAP HANA log backups, an SAP HANA system cannot wait until a stream is available because waiting can negatively affect the database performance.

If the Data Domain system runs out of streams during a backup, the backup fails (although not immediately) with the following error message in the operational log:

```
153004:hdbbackint: Unable to write to a file because the streams
limit was exceeded.
The error message is: [5519] [16805] [140261664245536] Tue May 10
06:45:23 2016
    ddp_write() failed Offset 0, BytesToWrite 317868,
BytesWritten 0 Err: 5519-Exceeded streams limit
```

You can set up the SAP HANA system to use two different storage units for the data backups and the log backups. You complete this setup by creating two different configuration files, one for the data backup and one for the log backup. You must specify a different value for the `DEVICE_PATH` parameter in each configuration file. This configuration enables you to plan the streams usage with other databases on the data storage unit while leaving the logs storage unit available to always accept logs from the SAP HANA system. The configuration also prevents the problem of a log backup using all the available streams and causing other backups to fail.

If the Data Domain system runs out of streams during a restore, then the restore fails (although not immediately) with the following error message in the operational log:

```
163971 11/28/2016 06:55:59 AM hdbbackint SYSTEM critical Unable to
read from a file because the streams limit was exceeded.
The error message is: [5519] [60299] [140167084230432] Mon Nov 28
06:55:59 2016
    ddp_read() failed Offset 192, BytesToRead 262144, BytesRead 0
Err: 5519-nfs readext remote failed (nfs: Resource (quota) hard limit
exceeded)
```

Performing DD Boost backups, recovery, and deletion with SAP HANA Studio

You can run the SAP HANA Studio GUI to perform DD Boost backups, backup deletions, restores, and recovery with the database application agent. The SAP HANA documentation provides details about the SAP HANA Studio procedures.

You can perform operations with SAP HANA Studio after you have completed the configurations in [Configuration of DD Boost operations in an SAP HANA environment](#) on page 254.

Performing DD Boost backups by using SAP HANA Studio

In SAP HANA Studio, you must specify the database for backup and enable the backup to use the `BACKINT` interface.

For example, the **Specify Backup Settings** dialog box appears as follows.

Figure 16 Specifying backup settings in SAP HANA Studio

Specify Backup Settings

Specify the information required for the data backup
Estimated backup size: 1.69 GB.

Backup Type: Complete Data Backup

Destination Type: File

Backup Destination

The default destination is used unless you specify a different destination. If you specify a new destination, ensure that the directory already exists. For improved data safety, we recommend that you specify an external backup destination.

Backup Destination: /usr/sap/VIN/HDB00/backup/data/SYSTEMDB

Backup Prefix: COMPLETE_DATA_BACKUP

i Note that customer-specific changes to the SAP HANA database configuration are not saved as part of the data backup.
More Information: SAP HANA Administration Guide

< Back Next > Finish Cancel

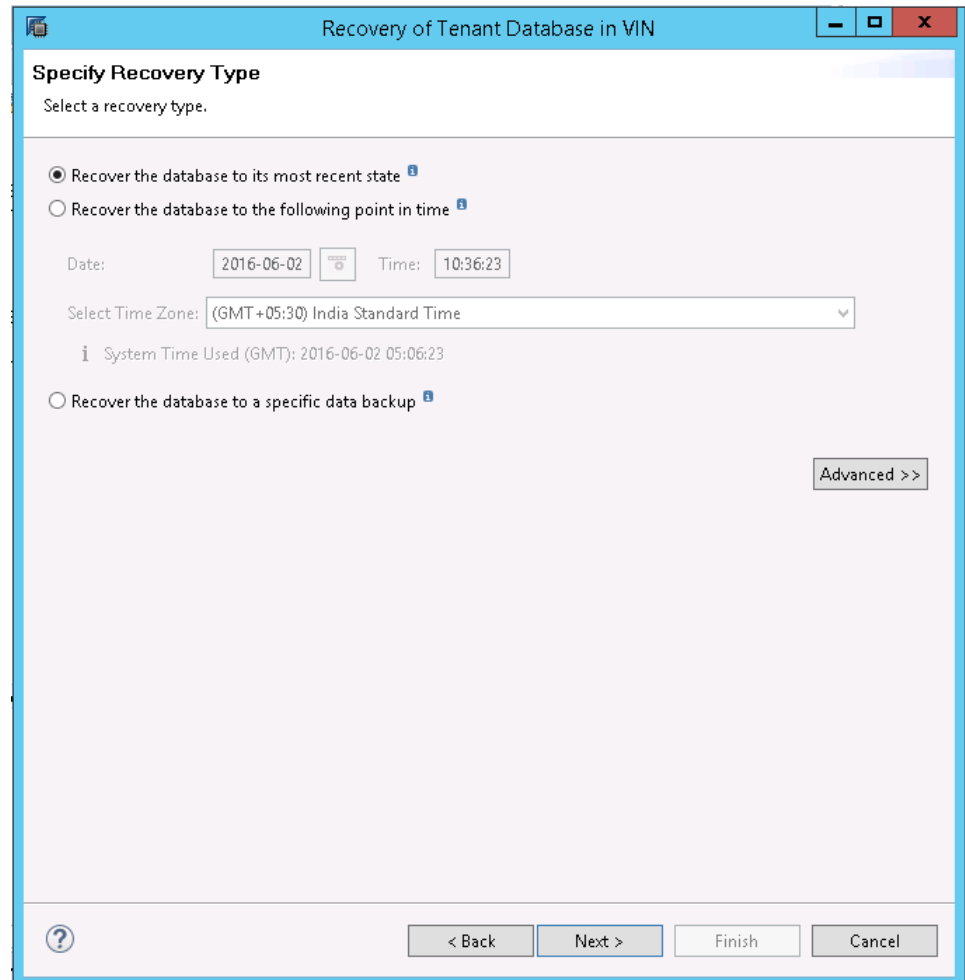
In the **Specify Backup Settings** dialog box, you must select **Backint** for **Destination Type**, and then optionally type a prefix for the backup name in the **Backup Prefix** text box. If you specify a prefix, the backup name will include the prefix, *prefix_databackup_#_#*.

For **Backup Type**, you can select one of the three supported backup types: **Complete Data Backup**, **Differential Backup**, **Incremental Backup**.

Performing DD Boost restore and recovery by using SAP HANA Studio

1. Start SAP HANA Studio.
2. In the recovery GUI, go to the **Specify Recovery Type** page.

Figure 17 Specifying the recovery type in SAP HANA Studio



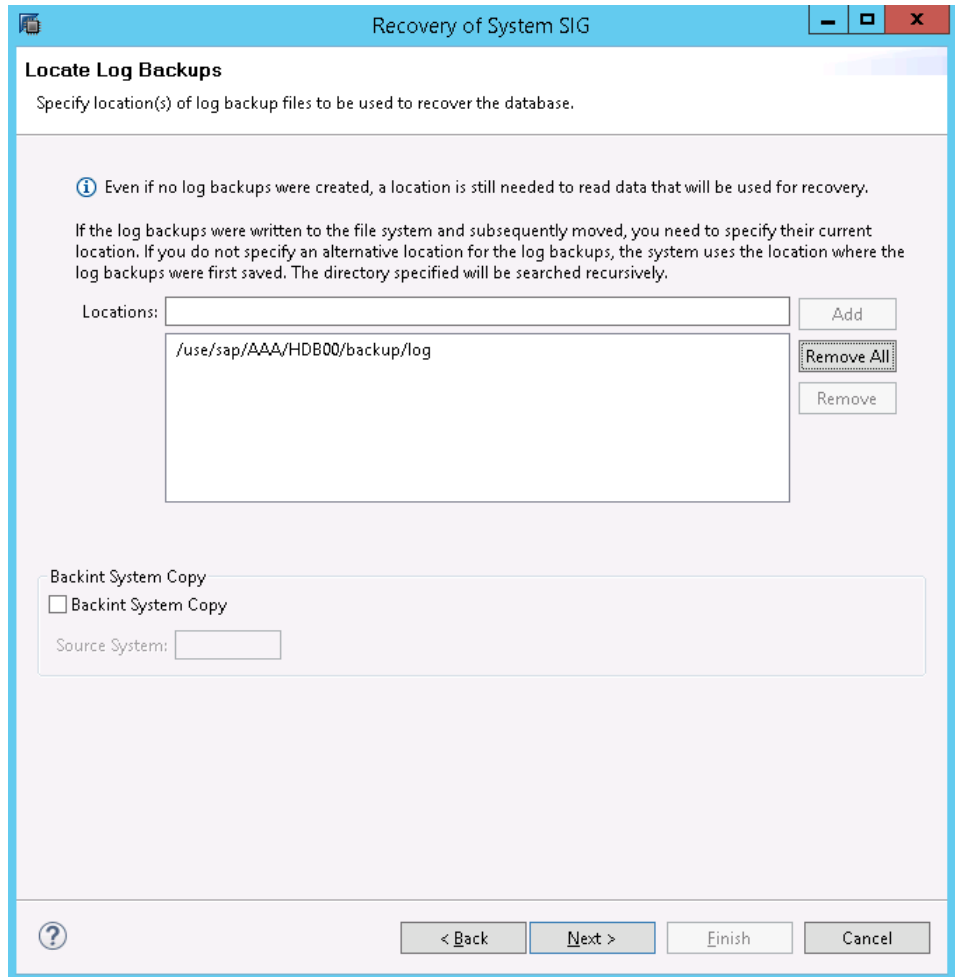
- a. Select one of the following types of recovery for the required SID:
 - **Recover the database to its most recent state**—Recovers the database as close as possible to the current time.
 - **Recover the database to the following point in time**—Recovers the database to a specific point-in-time.
 - **Recover the database to a specific data backup**—Restores only the specified data backup.
 - b. In an exceptional case when a previous recovery has failed, click **Advanced >>**, and then select **Recover the database to the following log position**.
 - c. Click **Next >**.
3. On the **Locate Log Backups** page, specify the locations of the log backup files to use to perform recovery.

The backup locations for the single database containers and the multitenant database containers are as follows:

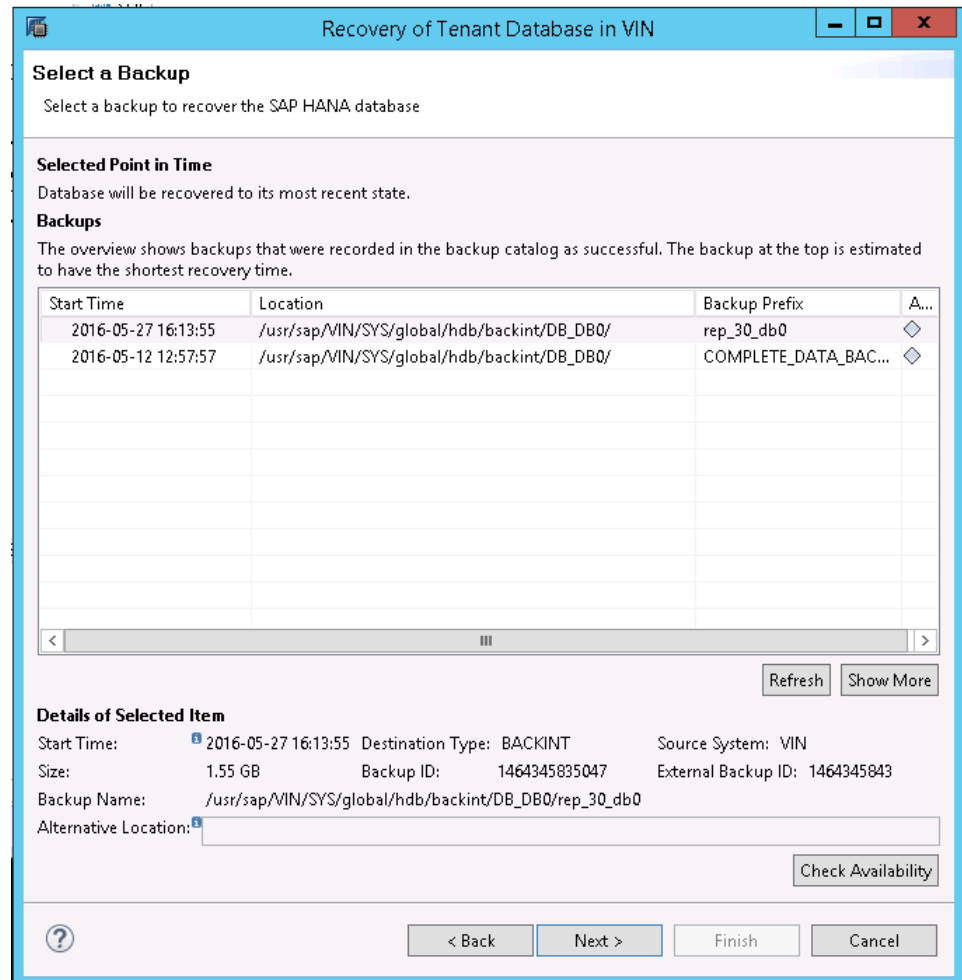
- **Single database container**—`/usr/sap/<SID>/SYS/global/hdb/backint/`
- **Multitenant database container (System DB)**—`/usr/sap/<SID>/SYS/global/hdb/backint/SYSTEMDB`

- **Multitenant database container (Tenant DB)**—`/usr/sap/<SID>/SYS/global/hdb/backup/DB_<tenantDB>`

Figure 18 Locating the log backups in SAP HANA Studio



4. On the **Select Data Backup** page, select the database backup for either restore or recovery.

Figure 19 Selecting the data backup in SAP HANA Studio

To check the availability of a backup, select the backup in the table, and then click **Check Availability**. The **Available** column in the table displays either a green icon if the backup is available or a red icon if the backup is not available.

5. On the **Other Settings** page, select the required options, and then click **Next >**.

6. On the **Review Recovery Settings** page, review the information, and then click **Finish**.

The recovery progress appears for each service that includes the name server, the index server, and the statistics server. A confirmation message appears when the recovery completes.

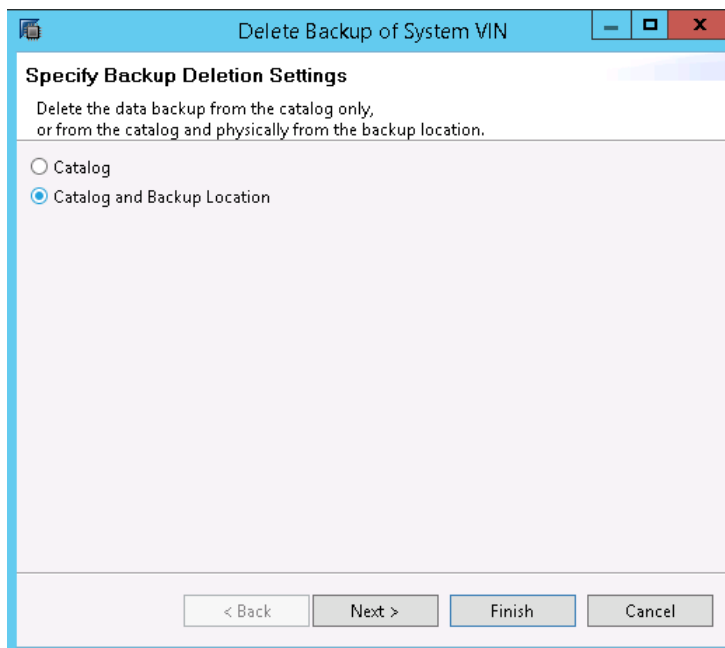
Deleting DD Boost backups by using SAP HANA Studio

By using SAP HANA, you can delete either a backup or the older backups of a backup. Deleting the older backups of a full backup includes deleting the older full, log, and delta backups.

1. Start SAP HANA Studio.
2. In the backup GUI, on the **Backup Catalog** tab:
 - To delete a backup:

- a. Right-click the backup in the table, and then select **Delete Data Backup....**
- b. On the **Specify Backup Deletion Settings** page, select the required type of deletion, and then click **Next >**.
 - **Catalog**—Deletes the backup from the backup catalog only.
 - **Catalog and backup location**—Deletes the backup from both the backup catalog and the Data Domain system.

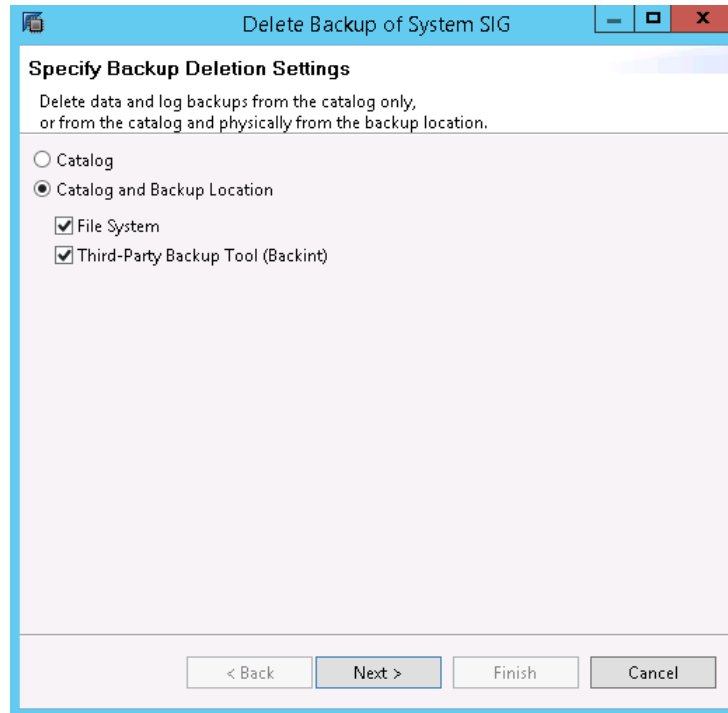
Figure 20 Specifying settings to delete a backup in SAP HANA Studio



- c. On the **Review Backup Deletion Settings** page, review the information, and then click **Finish**.
- To delete the older backups of a backup:
 - a. Right-click the backup in the table, and then select **Delete Older Backups....**
 - b. On the **Specify Backup Deletion Settings** page, select the required type of deletion, and then click **Next >**.
 - **Catalog**—Deletes the backup from the backup catalog only.
 - **Catalog and Backup Location**—Deletes the backup from both the backup catalog and the Data Domain system.

Select the location as **File System** or **Third-Party Backup Tool (Backint)** according to the location of the backups to be deleted.

Figure 21 Specifying settings to delete the older backups of a backup in SAP HANA Studio



- c. On the **Review Backup Deletion Settings** page, review the information and click **Finish**.

Performing DD Boost backups and recovery with SAP HANA CLI

You can run the SAP HANA CLI to perform backups of SAP HANA data.

Starting with SAP HANA SPS 07, you can run the SAP HANA CLI to perform recovery of the SAP HANA backups, and run specific maintenance commands to check the availability and integrity of data and log backups.

You can perform the SAP HANA operations after you have completed the backup configurations in [Configuration of DD Boost operations in an SAP HANA environment](#) on page 254.

Performing DD Boost backups with the SAP HANA CLI

To perform a database backup, you can run the appropriate `hdbsql` command. For example:

```
hdbsql -U <user_key> "backup data using backint ('/usr/sap/<SID>/SYS/global/hdb/backint/<prefix>')"
```

where:

- `<user_key>` is the user store key created with the SAP HANA `hdbuserstore` tool. SAP HANA documentation provides details about the tool.
- `<SID>` is the system ID of the SAP HANA system. `/usr/sap/<SID>/SYS/global/hdb/backint/` is optional in the `hdbsql` command.

- *<prefix>* is an optional prefix for the backup name. If you specify *<prefix>* in the `hdbsql` command, then the backup name will include the prefix, *<prefix>_databackup_#_#*.

The SAP HANA documentation provides details about how to use the `hdbsql` command for backups.

Canceling DD Boost backups with the SAP HANA CLI

You can cancel an SAP HANA backup by using the appropriate SQL command with the backup ID of the running data backup.

Procedure

1. Determine the backup ID of the running data backup by using the monitoring view `M_BACKUP_CATALOG`, which provides an overview of information about backup and recovery activities.

To determine the backup ID, run the following SQL command:

```
select BACKUP_ID from "M_BACKUP_CATALOG" where
entry_type_name = 'complete data backup' and state_name =
'running' order by sys_start_time desc;
```

2. To cancel the running data backup, run the following SQL command:

```
backup cancel backup_ID
```

Checking DD Boost backups with the SAP HANA CLI

To check the integrity and availability of SAP HANA backups, starting with SAP HANA SPS 07, you can run the `hdbbackupcheck` and `hdbbackupdiag` commands.

Using the `hdbbackupcheck` command to check backup integrity

To check the integrity of a data or log backup, you can run the `hdbbackupcheck` command. The command reads part of the backup, checks the metadata for correctness and consistency, and checks the content for any changes.

The following example command checks the integrity of an SAP HANA backup:

```
./hdbbackupcheck -v --backintParamFile /usr/sap/space/space_hana/
init.ddp.utl /usr/sap/AAA/SYS/global/hdb/backint/
foooo_databackup_0_1 -e 1396019304
```

In this command, the backup file name is the pipe name provided by SAP HANA during the backup. The pathname is the location where the pipes were created. The `-e` option provides the external backup ID of the backup file. You can find all these values in the `M_BACKUP_CATALOG_FILES` view.

The SAP Note 1869119 provides details about the `hdbbackupcheck` command.

Using the hdbbackupdiag command to check backup availability

To check the backup availability, you can run the `hdbbackupdiag` command.

The following example command checks the availability of an SAP HANA backup:

```
./hdbbackupdiag --check -i 1393886826664 --useBackintForCatalog
--backintDataParamFile /nsr/res/initAAA.utl --backintLogParamFile
/nsr/res/initAAA-logs.utl
```

Ensure that you use the `--useBackintForCatalog` option.

The SAP Note 1873247 provides more details about the `hdbbackupdiag` command.

Performing DD Boost recovery with the SAP HANA CLI

To perform a recovery of an SAP HANA backup, starting with SAP HANA SPS 07, you can use an SQL command.

Note

The SAP HANA Studio GUI is the recommended method to perform a restore and recovery.

To run the recovery SQL command, you must be logged in to the SAP HANA system as the OS user `<SID>adm`. In a scale-out multinode environment, you must run the command from the master node.

The recovery command must first include the `HDBSettings.sh` and `recoverSys.py` scripts in that order. The `HDBSettings.sh` script sets the environment for recovery. The `recoverSys.py` script processes the recovery SQL command and performs the database recovery.

The following example command performs a recovery of an SAP HANA backup:

```
./HDBSettings.sh recoverSys.py --command="recover database until
timestamp '2014-02-28 00:00:00' using data path
('/usr/sap/AAA/SYS/global/hdb/backint/') using log path
('/usr/sap/AAA/SYS/global/hdb/backint/') using BACKUP_ID
1593520382350" --wait
```

You must include the `--wait` option for proper completion of the `recoverSys.py` script.

The SPS 07 version of the *SAP HANA Administration Guide* provides details about the SQL command syntax of the recovery command.

Preparing for SAP HANA disaster recovery

For a comprehensive disaster recovery plan, you must ensure that you can reconstruct the computing environment and all the SAP HANA server files associated with maintaining data on the application host.

To prepare for a disaster recovery of the SAP HANA server host, use the following guidelines:

- Maintain accurate and complete records of the network and system configurations. Keep all the original software media and the following items in a safe location:
 - Original operating system media and patches
 - Device drivers and device names
 - File system configuration
 - IP addresses and hostnames
- To ensure the recovery of customer-specific settings of the database, back up the SAP HANA configuration files by following the *SAP HANA Administration Guide*. These files are not backed up during database backups.
- Ensure that you have a current full backup of the database and all the redo logs as required for a recovery of the database to the current time.

Perform the following tasks for a disaster recovery.

Procedure

1. Set up the SAP HANA database system to the same configuration as when the data was backed up.
2. Set up the same SAP HANA configuration file for the recovery, as was used to perform backups.

Set the `CLIENT` parameter in the SAP HANA configuration file to the hostname of the host, to which the data was backed up. [Configuring the SAP HANA parameters](#) on page 254 provides details.

3. Re-create the lockbox on the database host. [Configuring the lockbox](#) on page 103 provides details about the lockbox.
4. Ensure that all the required database and log backup images are available.
5. To recover the database and redo logs, run SAP HANA Studio.

SAP HANA scale-out requirements for DD Boost operations

You can perform SAP HANA backups in an SAP HANA scale-out multinode environment.

You must meet the following configuration requirements in an SAP HANA scale-out environment:

- You have set up the database in the scale-out environment according to the appropriate SAP HANA documentation.
- You have installed the database application agent software on each node.
- You have set the `CLIENT` parameter in the SAP HANA configuration file to the hostname of one of the nodes. If the configuration file is not in a shared location, you have set the common `CLIENT` parameter in the configuration file on each node. [Common parameters](#) on page 80 provides details on the parameter.
- You have completed the lockbox configuration for all the participating hosts. The hosts can use either a shared lockbox or individually configured lockboxes. [Configuring the lockbox with the `ddbadmin` command](#) on page 104 provides details.

SAP HANA troubleshooting tips for DD Boost operations

[General troubleshooting tips](#) on page 132 provides common troubleshooting information that applies to the database application agent operations with all the supported databases and applications.

The database application agent maintains the operational and debugging logs in the standard directory `/opt/dpsapps/dbappagent/logs` on Linux, linked to `/var/opt/ddbda/logs`.

You must set the `DEBUG_LEVEL` parameter to enable debugging.

Limitation in dynamic tiering support with SAP HANA

With SAP HANA 1.0 SPS 12 to SAP HANA 2.0 SPS 01 inclusive, the database application agent supports backups and restores for dynamic tiering with Backint to the same extent as supported by SAP.

As a result, the restore of SAP HANA databases with dynamic tiering does not support the use of a `backup_id` during the restore. Only a restore to a specific backup is supported. The SAP Note 2363526 provides details about this limitation.

Note

This limitation in dynamic tiering support does not apply with SAP HANA 2.0 SPS 02 or later.

Limitations in support of SAP HANA 1.0 SPS 09

The database application agent supports SAP HANA 1.0 SPS 09, which includes the following known limitations:

- SAP HANA 1.0 SPS 09 does not support the backup and restore of multitenant databases using `backint`. The SAP Note 2096000 provides details about this limitation.
- The `hdbbackupdiag --check` command does not work for SAP HANA SPS 09 releases 1.00.90 to 1.00.93. The command produces the following error message:

```
ERROR: [110081] Catalog backup log_backup_0_0_0_0 not found
```

A fix for this limitation was introduced in SAP HANA release 1.00.94.

Limitations in support of SAP HANA 2.0 SPS 00

Due to limitations in SAP HANA 2.0 SPS 00, the following issues might occur during the database application agent operations:

- SAP HANA 2.0 SPS 00 REV 00 might report an error during a multitenant database container (MDC) restore operation that is otherwise successful. The SAP Notes 2222121 and 2395530 provide details about this limitation.
- A restore with SAP HANA 2.0 SPS 00 that uses Backint might take longer than expected with a large number of index files. The increased duration of the restore is due to SAP HANA 2.0 requesting the restore of the

`backup_catalog_extension` file, which has not been backed up through Backint. Contact SAP Technical Support for more information.

CHAPTER 10

DD Boost Operations on SAP with Oracle Systems

This chapter includes the following topics:

- [Overview of DD Boost operations in an SAP with Oracle environment.....](#) 272
- [Configuration of DD Boost operations in an SAP with Oracle environment.....](#) 274
- [Performing DD Boost backups and recovery with SAP BR*Tools.....](#) 283
- [Preparing for SAP with Oracle disaster recovery.....](#) 284
- [SAP with Oracle RAC and cluster requirements for DD Boost operations.....](#) 286
- [SAP with Oracle troubleshooting tips for DD Boost operations.....](#) 287

Overview of DD Boost operations in an SAP with Oracle environment

The database application agent is integrated with the SAP BR*Tools `backint` interface and also with the BR*Tools Oracle Recovery Manager (RMAN) interface. This integration enables DD Boost backups, restores, and transaction log archiving in an SAP with Oracle environment.

You can perform a DD Boost backup, restore, or recovery operation with the product on an SAP with Oracle database server by running one of the supported SAP tools:

- BR*Tools command line interface (CLI) with the commands `brbackup`, `brarchive`, `brrestore`, `brrecover`, and `brtools`
- BRGUI
- BR*Tools Studio GUI
- SAP DBA Cockpit in The Computing Center Management System (CCMS) GUI

You can use these tools in cooperation with the database application agent to perform the following operations:

- Online and offline backups
- Full backups of database, tablespace, or datafiles
- Block-level incremental backups through the RMAN interface only
- Archived redo log backups
- Recovery of a database to the current time or a specific point-in-time
- Recovery to the original location or an alternate location
- Backup and recovery of directories
- Oracle ASM operations through the RMAN interface only
- Control parallelism for backups and restores

Note

SAP BR*Tools does not provide backup deletion or other backup maintenance operations. You can use save set deletion tool (provided with the database application agent) to list and delete backups, as described in [Configuring the display and deletion of save set information](#) on page 117.

BR*Tools RMAN backups do not use the Oracle recovery catalog. RMAN backup information is stored in the Oracle control file only, which is backed up during each backup.

The product maintains online backup indexes on the Data Domain system. During backups, the product creates backup entries in the online indexes, which provide the information required to restore the backed-up data.

The troubleshooting section at the end of this chapter provides details about limitations in the DD Boost operations with the database application agent in an SAP with Oracle environment.

SAP with Oracle backup processes

An SAP with Oracle backup includes the following process interactions.

1. The database administrator initiates the backup by running the BR*Tools CLI, BRGUI, BR*Tools Studio GUI, or SAP DBA Cockpit in CCMS.
2. The `brbackup` or `brarchive` program on the SAP with Oracle database server invokes the `backint` or RMAN program, and then passes a list of files or directories to back up.
3. The `backint` program or the shared library of the database application agent reads the parameters from the configuration file, and then initializes the connection with the Data Domain system.
4. The following steps occur for the `backint` backup or RMAN backup:
 - For the `backint` backup:
 - a. The `backint` program starts the child `backint` processes that back up the required files.
 - b. The child `backint` processes send the backup data and tracking information to the Data Domain system for storage by using the DD Boost interface.
 - For the RMAN backup:
 - a. The Oracle software ensures that each backup piece name is unique and sends the backup pieces to the database application agent through the SBT API.
 - b. The database application agent sends the backup data and tracking information to the Data Domain system for storage by using the DD Boost interface.
 - c. The Oracle software asks the database application agent to confirm that the backup is in the catalog of the database application agent, then records the entry in the Oracle catalog and completes the backup.
 - d. The `brbackup` or `brarchive` program invokes the `backint` program to back up the BR*Tools metadata.

SAP with Oracle restore processes

An SAP with Oracle restore includes the following process interactions.

1. The database administrator initiates the restore by running the BR*Tools CLI, BRGUI, or BR*Tools Studio GUI.
2. The `brrestore` or `brrecover` program on the SAP with Oracle database server runs the `backint` or RMAN program, and then passes a list of files or directories to restore.
3. The `backint` program or the shared library of the database application agent reads the parameters from the configuration file, and then initializes the connection with the Data Domain system.
4. The following steps occur for the `backint` restore or RMAN restore:
 - For the `backint` restore:
 - a. The `backint` program starts the child `backint` processes that restore the required files.

- b. The child `backint` processes retrieve the backup data from the Data Domain system to the SAP with Oracle database server.
- For the RMAN restore:
 - a. The Oracle software queries, and then requests the backup pieces from the database application agent through the SBT API.
 - b. The database application agent queries its catalog, and then retrieves the backup data from the Data Domain system by using the DD Boost interface.

After the backup data is restored, the database administrator must recover the database by using the `brrecover` command or the Oracle SQL Plus tool. The SAP with Oracle documentation provides details.

Configuration of DD Boost operations in an SAP with Oracle environment

You must complete the required configurations of the database application agent to enable the DD Boost operations in an SAP with Oracle environment. You can select either the `backint` program or the RMAN program as the backup and restore utility to be used by SAP BR*Tools. The following topics provide the product configuration details.

[SAP with Oracle RAC and cluster requirements for DD Boost operations](#) on page 286 provides additional details on the specific configuration requirements in an SAP with Oracle cluster environment.

The troubleshooting section at the end of this chapter provides details about limitations in the DD Boost operations with the database application agent in an SAP with Oracle environment.

Confirming the environment and file permissions

You must confirm the settings of the database server environment and file permissions before you perform any SAP with Oracle operations.

The DBA operating system group must have read access to datafiles that will be restored. This read access enables the restore of the data by a different database user to a different host.

You can restore only datafiles for which you have read permission, based on the files' operating system permissions at the time that the files were backed up.

On UNIX, the read permission is associated with the user ID (UID) and group ID (GID), not the username or group name. The UID and GID of the user performing the restore must match the IDs associated with the files at backup time.

Enabling administrator privileges for SAP with Oracle restores on Windows

On specific types of Microsoft Windows systems, User Account Control (UAC) is designed to provide additional operating system security by preventing software from being installed or run unless an administrator authorizes the elevated privileges.

On Windows systems with UAC enabled, prior to the start of a restore from the CLI, ensure that administrator privileges are enabled for the user that will perform the restore. Otherwise, the operation might fail.

You can enable administrator privileges from the CLI as follows.

Procedure

1. Right-click the **Command Prompt** icon.
2. Select **Run as administrator** from the list.

Configuring the DD Boost operations with the backint utility

If you want SAP BR*Tools to use the `backint` program for the DD Boost backups and restores, then you must complete the following configurations.

If you prefer that Oracle RMAN be used, then you must complete the configurations in [Configuring the DD Boost operations with Oracle RMAN](#) on page 279.

Integrating the product into the BR*Tools environment for backint

You must complete the required settings in the BR*Tools configuration file `init<DBSID>.sap` to enable the BR*Tools operations to use the `backint` program:

- Define the directory from which BR*Tools calls the `backint` program.

By default, BR*Tools calls `backint` from the `sapexe` directory. You can set the `util_path` parameter to the directory pathname where `backint` is located. For example:

- On UNIX:

```
util_path = /opt/dpsapps/dbappagent/bin
```

- On Windows:

```
util_path = C:\PROGRA~1\DPSAPPS\DBAPPAGENT\bin
```

Note

On Windows, you must specify the short version of the directory pathname `C:\Program Files\DPSAPPS\DBAPPAGENT\bin` because you must not include spaces in file pathnames that you set in `init<DBSID>.sap`. If you include spaces in a pathname setting, the backup will fail. To obtain the short version of a directory, run the `dir /x` command in the parent directory. For example:

```
C:\> dir /x
```

```

      .
      .
11/21/2014  07:38 AM    <DIR>          PROGRA~1    Program Files

```

Here, the short version of “Program Files” is `PROGRA~1`.

- Set the backup medium to use the `backint` program.

Set the `backup_dev_type` parameter to one of the following values:

```
backup_dev_type = util_file
```

or

```
backup_dev_type = util_file_online
```

Set `backup_dev_type = util_file_online` for online backups to decrease the amount of time each tablespace remains in hot backup mode, resulting in a smaller number of generated transaction logs.

You can override this `backup_dev_type` setting with the `-d` option when you perform a BR*Tools operation from the command line. For example:

```
brbackup -d util_file_online
```

- Set the `util_par_file` parameter to the complete pathname of the configuration file as configured in [Configuring the SAP with Oracle parameters for backint](#) on page 276. For example:

```
util_par_file = ?/dbs/init<DBSID>.utl
```

where `?` is `$ORACLE_HOME`.

If you do not specify the complete pathname of this configuration file, the software searches for the file under the following default directory:

- On UNIX: `$ORACLE_HOME/dbs`
- On Windows: `%ORACLE_HOME%\database`

You can override this `util_par_file` setting by specifying the configuration file pathname with the `-r` option when you perform a BR*Tools operation from the command line. For example:

```
brbackup -r pathname/init<DBSID>.utl
```

The SAP BR*Tools documentation provides more details about parameters in the BR*Tools configuration file.

Configuring the SAP with Oracle parameters for backint

You must set the SAP with Oracle parameters for `backint` operations in the configuration file named `init<DBSID>.utl`. You must specify the location of this file in the BR*Tools configuration file as described in [Integrating the product into the BR*Tools environment for backint](#) on page 275.

For example, the configuration file contains the following mandatory parameter settings:

```
DDBOOST_USER=qa_ost
DEVICE_HOST=bu-dbe-890.lss.emc.com
DEVICE_PATH=/bu-star1_ora
```

[Setting up the configuration file](#) on page 78 describes the common parameters and how to set parameters in the configuration file. [Configuring restores of replicated backups](#) on page 90 also describes the parameters and requirements for the restores of replicated backups.

Ensure that the configuration file contains any other required parameters from the following table. For each parameter, the table lists the section heading of the configuration file section that contains the parameter.

After the configuration file is set up, ensure that the required lockbox procedures have been performed as described in [Configuring the lockbox](#) on page 103.

Table 24 SAP with Oracle parameters for DD Boost operations with backint

Parameter: ARCH_LOGS_SEQUENTIAL

Section: [GENERAL]

Specifies whether the `brarchive` program backs up archive logs in alphabetical order to optimize the log removal (cleanup) during an archive log backup.

Setting this parameter to TRUE causes the `brarchive` program to back up archive logs in alphabetical order, which speeds up the log cleanup process during the archive log backup.

The parameter is ignored during a `brbackup` or `brrestore` operation.

Optional for backups of archive logs with `brarchive`.

Valid values:

- FALSE (default).
- TRUE.

Note

When you use the nondefault value TRUE, the load balancing parameter `GROUP_BY_FS` is ignored during the `brarchive` backup.

Parameter: GROUP_BY_FS

Section: [GENERAL]

If you set this parameter to TRUE, then the operation ignores the `SAVESETS` parameter and groups the files by file system instead of file size. This efficient grouping of files can improve the performance of backups, index searches, and restore times.

Note

Setting this parameter to TRUE means that all the files being processed must be visible within the local file system. Windows UNC pathnames must be mapped to a local drive letter.

Optional for a backup.

Valid values:

- FALSE (default).
- TRUE.

Parameter: PARALLELISM

Section: [GENERAL]

Specifies the number of concurrent data streams to send to or from the Data Domain system during a backup or restore, for each `backint` program that SAP Oracle runs for the operation.

Table 24 SAP with Oracle parameters for DD Boost operations with backint (continued)

<p>Note</p> <p>The parallelism value for a backup is reduced if the value is greater than the <code>SAVESETS</code> parameter value.</p> <hr/> <p>Optional for a backup or restore.</p> <p>Valid values:</p> <ul style="list-style-type: none"> • 8 (default). • Positive integer number.
<p>Parameter: RAW_DEVICE_SKIP_BYTES</p> <p>Section: [GENERAL]</p> <p>Specifies to omit unwanted data from recovery on a raw device. Specifies the initial amount of data to skip.</p> <hr/> <p>Note</p> <p>When you adjust the amount of data to skip, you require expert knowledge of the operating system and the volume manager specifications.</p> <hr/> <p>Valid values:</p> <ul style="list-style-type: none"> • Undefined (default). • Valid device name, followed by an equal sign and the amount of data to skip in bytes, kilobytes, megabytes, or gigabytes: <code>RAW_DEVICE_SKIP_BYTES=/raw/dev1=skip[B K M G]</code> <p>Specify multiple devices by separating the devices with a colon. For example, the following setting skips 8 KB for the raw device <code>/dev/rdisk/c2t4d0s5</code>:</p> <pre>RAW_DEVICE_SKIP_BYTES=/dev/rdisk/c2t4d0s5=8K</pre>
<p>Parameter: RAW_DIR</p> <p>Section: [GENERAL]</p> <p>Specifies the directories for raw disk partitions on UNIX only. Any files below these directories are considered to be raw disk partitions.</p> <p>Do not set this parameter to the <code>/dev</code> or <code>/devices</code> directories, which are already treated as raw devices.</p> <hr/> <p>Note</p> <p>The listed partitions are pointed to by the symbolic links under <code>\$SAPDATA_HOME/sapraw/TBS_number/</code>.</p> <hr/> <p>Valid values:</p> <ul style="list-style-type: none"> • Undefined (default). • Directory pathnames of the raw disk partitions, separated by a semicolon (;). For example: <code>RAW_DIR = /oracle/data;/sap/data</code>
<p>Parameter: RELOCATION_DEST</p> <p>Section: [GENERAL]</p>

Table 24 SAP with Oracle parameters for DD Boost operations with backint (continued)

For a relocated restore, specifies a new root directory for SAP datafiles when `SAPDATA_HOME` is changed but the structure of subdirectories under `SAPDATA_HOME` is not changed.

Valid values:

- Undefined (default).
- Same value as the new `SAPDATA_HOME` value.

Parameter: SAVESETS**Section: [GENERAL]**

Specifies the number of save sets created on the Data Domain target. The backup data files are divided into the save sets based on file size.

This parameter is ignored if the `GROUP_BY_FS` parameter is set to `TRUE`.

Valid values:

- 20 (default).
- Positive integer number.

Configuring the DD Boost operations with Oracle RMAN

If you want SAP BR*Tools to use the Oracle RMAN program for the DD Boost backups and restores, then you must complete the following configurations.

Note

The database application agent always uses the `backint` program to back up and restore the BR*Tools metadata, including the configuration files and backup catalog, even if you complete the following RMAN configurations.

If you prefer that `backint` be used for all backups and restores, then you must complete the configurations in [Configuring the DD Boost operations with the backint utility](#) on page 275.

Integrating the product into the BR*Tools environment for RMAN

To enable the BR*Tools operations to use the RMAN program, you must complete the required settings in the BR*Tools configuration file `init<DBSID>.sap`:

- Set the backup medium to use the RMAN program.

Set the `backup_dev_type` parameter to the following value:

```
backup_dev_type = rman_util
```

Note

Do not use the `rman_disk` or `rman_stage` value because these values cause BR*Tools to back up the metadata without using the database application agent.

You can override this `backup_dev_type` setting with the `-d` option when you perform a BR*Tools operation from the command line. For example:

```
brbackup -d rman_util
```

- Set the `SBT_LIBRARY` and `CONFIG_FILE` values in the `rman_parms` parameter:
 - Set `SBT_LIBRARY` to the complete pathname of the database application agent library that is used with RMAN.
 - Set `CONFIG_FILE` to the complete pathname of the configuration file, `init<DBSID>.utl`. [Configuring the SAP with Oracle parameters for RMAN on page 281](#) provides details on setting up the configuration file.

The following examples show the `rman_parms` settings:

- On UNIX or Linux:

```
rman_parms = "SBT_LIBRARY=/opt/dpsapps/dbappagent/lib/lib64/libddboostsapora.so, ENV=(CONFIG_FILE=/db/app/oracer/product/11.2.0/dbhome_1/dbs/initCER.utl)"
```

- On Windows:

```
rman_parms = "SBT_LIBRARY=C:\PROGRA~1\DPSAPPS\DBAPPAGENT\bin\libddboostsapora.dll, ENV=(CONFIG_FILE=D:\app\oracer\product\11.2.0\dbhome_1\database\initCER.utl)"
```

Note

On Windows, you must specify the short version of a directory pathname to avoid including spaces in the pathname in `init<DBSID>.sap`. If you include spaces in a pathname setting, the backup will fail. To obtain the short version of a directory, run the `dir /x` command in the parent directory. For example:

```
C:\> dir /x
```

```

      .
      .
11/21/2014  07:38 AM    <DIR>          PROGRA~1      Program Files

```

Here, the short version of “Program Files” is `PROGRA~1`.

Alternatively, you can set `CONFIG_FILE` in the `rman_send` parameter instead of the `rman_parms` parameter. For example:

```
rman_send="ENV=(CONFIG_FILE=/db/app/oracer/product/11.2.0/dbhome_1/dbs/initCER.utl)"
```

- Set any other `rman_XXX` parameters required to configure the RMAN operations. For example:
 - Set `rman_channels` to the number of concurrent data streams. The default value is 1.
 - Set `rman_filesperset` or `rman_maxopenfiles` to improve the deduplication ratio.

Note

The database application agent supports only the values 0 (default) and 1 for the `rman_copies` parameter, and both values cause the creation of a single backup copy.

The BR*Tools documentation provides details on all the supported RMAN parameters.

- Set the `util_par_file` parameter to the complete pathname of the configuration file as configured in [Configuring the SAP with Oracle parameters for RMAN](#) on page 281. This setting ensures that the database application agent backs up the BR*Tools metadata.
-

Note

The `CONFIG_FILE` and `util_par_file` settings must be the same.

For example:

```
util_par_file = ?/dbs/init<DBSID>.utl
```

where ? is `$ORACLE_HOME`.

If you do not specify the complete pathname of this configuration file, the software searches for the file under the following default directory:

- On UNIX: `$ORACLE_HOME/dbs`
- On Windows: `%ORACLE_HOME%\database`

You can override the `util_par_file` setting by specifying the configuration file pathname with the `-r` option when you perform a BR*Tools operation from the command line. For example:

```
brbackup -r pathname/init<DBSID>.utl
```

The SAP BR*Tools documentation provides more details about parameters in the BR*Tools configuration file.

Configuring the SAP with Oracle parameters for RMAN

You must set the SAP with Oracle parameters for RMAN operations in the configuration file named `init<DBSID>.utl`. You must specify the location of this file in the BR*Tools configuration file as described in [Integrating the product into the BR*Tools environment for RMAN](#) on page 279.

For example, the configuration file contains the following mandatory parameter settings:

```
DDBOOST_USER=qa_ost
DEVICE_HOST=bu-dbe-890.lss.emc.com
DEVICE_PATH=/bu-star1_ora
```

[Setting up the configuration file](#) on page 78 describes the common parameters and how to set parameters in the configuration file. [Configuring restores of replicated backups](#) on page 90 also describes the parameters and requirements for the restores of replicated backups.

After the configuration file is set up, ensure that the required lockbox procedures have been performed as described in [Configuring the lockbox](#) on page 103.

Estimating the Data Domain resource usage on SAP with Oracle systems

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

Capacity usage on SAP with Oracle systems

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

```
163542:backint: Unable to write to a file due to reaching the hard
quota limit.
The error message is: [5194] [ 1472] [2304] Fri Dec 09 00:50:50 2016
                ddp_write() failed Offset 746586112, BytesToWrite 524288,
BytesWritten 0 Err: 5194-Hard Quota Exceeded >
```

Streams usage on SAP with Oracle systems

Note

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

For SAP with Oracle using `backint`, the database application agent uses the `PARALLELISM` parameter to control the number of SAP streams that are used in a backup or restore operation. The `PARALLELISM` parameter specifies the number of concurrent SAP data streams to send to or from the Data Domain system during the backup or restore.

The number of Data Domain read or write streams that the `backint` program uses to connect to the Data Domain system is typically less than or equal to the `PARALLELISM` parameter setting.

If the Data Domain system runs out of streams during an SAP with Oracle backup, the backup operation fails. The database application agent generates the following type of error message in the operational log:

```
153004:backint: Unable to write to a file because the streams limit
was exceeded.
The error message is: [5519] [ 3052] [2536] Thu Dec 08 23:07:51 2016
                ddp_write() failed Offset 0, BytesToWrite 524288,
BytesWritten 0 Err: 5519-Exceeded streams limit
```

If the Data Domain system runs out of streams during an SAP with Oracle restore, the restore operation fails. The database application agent generates the following type of error message in the operational log:

```
63971 12/9/2016 1:11:48 AM backint SYSTEM critical Unable to read
from a file because the streams limit was exceeded.
The error message is: [5519] [ 2244] [2440] Fri Dec 09 01:11:48 2016
                ddp_read() failed Offset 0, BytesToRead 262144, BytesRead
0 Err: 5519-nfs readext remote failed (nfs: Resource (quota) hard
limit exceeded)
```

Performing DD Boost backups and recovery with SAP BR*Tools

You can perform DD Boost operations with SAP BR*Tools after you have completed the backup configurations in [Configuration of DD Boost operations in an SAP with Oracle environment](#) on page 274.

You can run the BR*Tools CLI or a supported GUI to perform DD Boost backup, restore, or recovery operations with the database application agent.

The SAP and Oracle documentation provides details about the available options.

Performing DD Boost backups with BR*Tools

You can run the `brbackup` and `brarchive` commands with the appropriate options to perform database and archived redo log backups. For example:

- Offline backup of the whole database:

```
brbackup -m all -t offline -d util_file
```

- Online backup of a single tablespace:

```
brbackup -m system -t online -d util_file_online
```

If you perform an online backup, you can run the `brarchive` command to back up the archived redo logs:

```
brarchive
```

Note

Without the redo logs, you can recover a database that was backed up online only to the time of its last full backup.

You can run the following command to back up nondatabase files or directories:

```
brbackup -m {sap_dir | ora_dir | all_dir | full_directory_path | nondatabase_file_path}
```

Performing DD Boost restore and recovery with BR*Tools

You can use the `brrestore` command or `brrecover` interface to perform a restore or recovery. The `brrestore` command restores the backed-up database files, and you must run Oracle SQL Plus to recover the database by applying the transaction logs to roll forward the database to a specific point-in-time. The `brrecover` interface automates the restore and recovery process by calling `brrestore` and SQL Plus to perform specific steps as needed.

For example, you can run the following command to restore the entire database:

```
brrestore -u / -m all -d util_file -c force
```

You can run the following command to restore a single tablespace:

```
brrestore -u / -m PSAPSR3DB -d util_file -c force
```

NOTICE

The `brrestore` program used for a restore operation, whether run directly from the command line or indirectly from the `brrecover` command, first deletes all the original files to be restored before the program runs `backint` to restore the backed-up files.

You will lose the original files if `brrestore` or `backint` fails.

Perform one of the following actions to prevent these issues:

- Restore the files to a different location that does not include any files by using the `-m` option and specifying the restore destination, for example:

```
brrestore -m tablespace_name=restore_directory
```

- Use the `brrestore -NFD` option to prevent deletion of the original files by `brrestore`.

To restore data from a directory (nondatabase) backup, use the `brrestore -m non_db` command. When you restore the SAP directory data, prevent deletion of the BR*Tools and `backint` binaries during the restore by performing a relocated restore. For example, run the following command:

```
brrestore -m non_db=restore_directory
```

Performing a DD Boost restore to an alternate host

To optionally restore the data to a different SAP Oracle host (destination host) than the one that was backed up:

- Set the `CLIENT` parameter to the hostname of the host where the data was backed up.
- Follow the disaster recovery steps to re-create the environment and restore the configuration files and BR*Tools logs. [Preparing for SAP with Oracle disaster recovery](#) on page 284 provides details.

Preparing for SAP with Oracle disaster recovery

For a comprehensive disaster recovery plan, you must ensure that you can reconstruct the computing environment and all the SAP Oracle server files associated with maintaining data on the application host.

Use the following guidelines to prepare for a disaster recovery of the SAP with Oracle server host:

- Maintain accurate and complete records of the network and system configurations. Keep all the original software media and the following items in a safe location:
 - Original operating system media and patches
 - Device drivers and device names

- File system configuration
- IP addresses and hostnames
- To ensure the recovery of customer-specific settings of the database, back up the system configuration files by following the SAP with Oracle documentation. These files are not backed up during database backups.
- Ensure that you have a current full backup of the database and all the archived redo logs as required for a recovery of the database to the current time.

To recover from a disaster, you must first restore any lost Oracle and SAP configuration files and lost BR*Tools backup log files according to the following instructions. After you restore these files, you can perform a database point-in-time recovery or a whole database reset according to the instructions in [Performing DD Boost restore and recovery with BR*Tools](#) on page 283.

Restoring the required Oracle and SAP BR*Tools files

The following procedure is a concise version of the disaster recovery steps described in the SAP documentation, modified for the specific requirements of the database application agent. You can perform this procedure on the original host or a new host, both referred to as the destination host.

To restore to a different host than the one backed up, follow the guidelines in [Performing a DD Boost restore to an alternate host](#) on page 284.

Procedure

1. If the entire SAP Oracle system is lost:
 - a. Reinstall all the required SAP and Oracle software components according to the SAP and Oracle documentation.
 - b. Reconfigure the SAP data layout, such as the `SAPDATA_HOME` directory and its subdirectories, to the same state as before the disaster.
2. Configure the BR*Tools configuration file, `init<DBSID>-dr.sap`, and the configuration file of the database application agent, `init<DBSID>-dr.utl`, for recovery on the destination host.
3. Re-create the lockbox on the destination host. [Configuring the lockbox](#) on page 103 provides details about the lockbox.
4. Ensure that the parameters are correctly set for a redirected restore according to [Configuring the SAP with Oracle parameters for backint](#) on page 276.
5. To perform a disaster recovery of profiles and logs on the destination host, start BR*Tools on the host, and then follow the onscreen instructions.

For example, the following steps show how to use BR*Tools for disaster recovery to restore profiles and log files from a `BRBACKUP` backup:

- a. In the `brtools` menu, select **Restore and Recovery** and then **Disaster recovery**.
- b. On the page **BRRECOVER options for disaster recovery**, set the location of the BR*Tools configuration file and the SAP with Oracle configuration file to the file names configured in [step 2](#).

Note

If you are restoring the original BR*Tools configuration file or SAP with Oracle configuration file, ensure that the files used for the disaster recovery have different names or are stored in a different location than the original files to be restored.

- c. On the page **Device type for restoring profiles and log files from BRBACKUP backup**, select **Backup utility**.
- d. On the page **Parameters for restoring profiles and log files from BRBACKUP backup utility backup**, specify the files to restore.
- e. On the page **Restore of profiles and log files from BRBACKUP backup**, select the components that you want to restore.

The SAP documentation provides more details about disaster recovery.

Recovering an SAP Oracle database after disaster

After you restore the correct SAP BR*Tools configuration file and logs on the application host, follow the normal BR*Tools recovery procedure to perform database point-in-time recovery or database reset.

[Performing DD Boost restore and recovery with BR*Tools](#) on page 283 provides the procedure to recover an SAP with Oracle database.

The SAP documentation provides more details about database recovery.

SAP with Oracle RAC and cluster requirements for DD Boost operations

You can perform DD Boost backups and restores in an Oracle RAC or active-passive cluster environment.

Active-passive cluster requirements

You must meet the following configuration requirements in an active-passive cluster environment:

- You have set up the database in the cluster according to the appropriate database server documentation.
- You have installed the database application agent on each node of the cluster that will participate in backups or restores.
- You have set the `CLIENT` parameter in the configuration file to the hostname of the virtual node. [Common parameters](#) on page 80 provides details.

Oracle RAC requirements

The SAP documentation provides details about the Oracle RAC setup. You must meet the following configuration requirements in an Oracle RAC environment:

- One of the Oracle RAC instances, set as the dedicated database (DDB) instance, is used to perform all the database administration tasks. You must be able to administer all the RAC instances from the DDB instance.
- You have installed BR*Tools and the database application agent on the DDB instance host.

- You have created the BR*Tools log directories, such as `SAPBACKUP` and `SAPARCH`, on a shared file system.
- You have correctly configured Oracle SQL Net.
- You have set the required parameters, such as `parallel_instances` and `db_services`, in the SAP initialization file, `init<DBSID>.sap`, located on the DDB instance.

Oracle RAC requirements for the backint interface

All the Oracle RAC backups performed with BR*Tools and the `backint` interface run on a single RAC instance, the DDB instance. The configuration of the DDB instance is the same as for a stand-alone Oracle system.

If you change the DDB instance after some backups were performed with the original DDB instance, set the parameter `CLIENT=original_DDB_instance_hostname` to ensure that all the backups are stored in the same location on the Data Domain system.

[Common parameters](#) on page 80 provides details on the parameter.

Oracle RAC requirements for the RMAN interface

BR*Tools 7.00 to 7.20 patch 30 support RMAN operations on a single instance only, the DDB instance. The configuration of these RMAN backups and restores is the same as for RMAN operations in a stand-alone environment. It is recommended that you set the `CLIENT` parameter to the DDB instance hostname, in case the DDB instance is moved later.

Starting with BR*Tools 7.20 patch 31 and 7.40, the SAP software supports Oracle RAC backups and restores on multiple RAC nodes, also known as distributed RMAN operations. Ensure that you meet the following requirements for these operations:

- You have installed the database application agent on each RAC node that will participate in the backups or restores.
- All the nodes included in the RMAN operations use the same type of network connection, IP or FC, to the Data Domain system.
- You have set the parameter `CLIENT=<DDB_node_name>` in the `init<DBSID>.utl` file. This configuration file is stored in a shared location, accessible to all the RAC nodes.
- You have set the required parameters in the BR*Tools `init<DBSID>.sap` file as described in [Integrating the product into the BR*Tools environment for RMAN](#) on page 279. Ensure that `rman_channels` is set to a value greater than 1, preferably a multiple of the number of RAC nodes. For example, set `rman_channels = 4` for a RAC system with 2 nodes (BR*Tools allocates 2 channels per node). To enable distributed RMAN operations, set `rman_rac_dist = yes`.
- For the distributed RAC operations, you have connected to the database with a specific database username, for example, by using the `-u <system>/<password>` option. The database user must have the `SAPDBA`, `SYSDBA`, and `SYSOPER` roles.

SAP with Oracle troubleshooting tips for DD Boost operations

[General troubleshooting tips](#) on page 132 provides common troubleshooting information that applies to the database application agent operations with all the supported databases and applications.

The database application agent maintains the operational and debugging logs in the standard directories:

- **On UNIX or Linux:** /opt/dpsapps/dbappagent/logs, linked to /var/opt/ddbda/logs
- **On Windows:** C:\Program Files\DPSAPPS\DBAPPAGENT\logs

To enable debugging, you must set the `DEBUG_LEVEL` parameter.

CHAPTER 11

ProtectPoint Operations on SAP with Oracle Systems

This chapter includes the following topics:

- Overview of ProtectPoint operations in an SAP with Oracle environment..... 290
- Configuration of ProtectPoint operations in an SAP with Oracle environment. 291
- Performing ProtectPoint backups and recovery with SAP BR*Tools.....300
- Preparing for SAP with Oracle disaster recovery..... 302
- SAP with Oracle RAC and cluster requirements for ProtectPoint operations...303
- ProtectPoint restore and rollback for VCS on Solaris.....304
- SAP with Oracle troubleshooting tips for ProtectPoint operations..... 310

Overview of ProtectPoint operations in an SAP with Oracle environment

The database application agent is integrated with the SAP BR*Tools `backint` interface to enable ProtectPoint backups, restores, and transaction log archiving in an SAP with Oracle environment.

You can perform a ProtectPoint backup, restore, or recovery operation with the product on an SAP with Oracle database server by running one of the supported SAP tools:

- BR*Tools command line interface (CLI) with the commands `brbackup`, `brarchive`, `brrestore`, `brrecover`, and `brtools`
- BRGUI
- BR*Tools Studio GUI
- SAP DBA Cockpit in The Computing Center Management System (CCMS) GUI

You can use these tools in cooperation with the database application agent to perform the following operations:

- Online and offline backups
- Full backups of database, tablespace, or datafiles
- Archived redo log backups
- Recovery of a database to the current time or a specific point-in-time
- Recovery to the original location or an alternate location
- Backup and recovery of directories

SAP BR*Tools does not provide backup deletion or other backup maintenance operations. You can use the save set deletion tool (provided with the database application agent) to list and delete backups, as described in [Configuring the display and deletion of save set information](#) on page 117.

The product maintains online backup indexes on the Data Domain system. During backups, the product creates backup entries in the online indexes, which provide the information required to restore the backed-up data.

The troubleshooting section at the end of this chapter provides details about limitations in the ProtectPoint operations with the database application agent in an SAP with Oracle environment.

SAP with Oracle backup processes

An SAP with Oracle backup includes the following process interactions.

1. The database administrator initiates the backup by running the BR*Tools CLI, BRGUI, BR*Tools Studio GUI, or SAP DBA Cockpit in CCMS.
2. The `brbackup` or `brarchive` program on the SAP with Oracle database server runs the `backint` program, installed as part of the database application agent, and passes some parameters and a list of files or directories to back up.
3. The `backint` program processes the parameters, including parameters from the configuration file of the database application agent.

4. The `backint` program determines the snapshotable files to back up with the ProtectPoint workflow and the nonsnapshotable files to back up with the DD Boost workflow.
5. Based on the parallelism settings, the `backint` program might start multiple child `backint` processes to back up the nonsnapshotable files.
6. The backup workflow proceeds for the snapshotable files as described in the topic about the ProtectPoint backup workflow or the ProtectPoint with RecoverPoint backup workflow in Chapter 1.
7. The `backint` program uses the DD Boost workflow to back up the Oracle and BR*Tools parameter files, catalog files, and control file.

SAP with Oracle restore processes

An SAP with Oracle restore includes the following process interactions.

1. The database administrator initiates the restore by running the BR*Tools CLI, BRGUI, or BR*Tools Studio GUI.
2. The `brrestore` or `brrecover` program on the SAP with Oracle database server runs the `backint` program and passes some parameters and a list of files or directories to restore.
3. The `backint` program processes the parameters, including parameters from the configuration file of the database application agent.
4. The `backint` program performs an index lookup, and then starts a ProtectPoint workflow to restore the snapshotable files and a DD Boost workflow to restore the nonsnapshotable files.
5. Based on the parallelism settings, the `backint` program might start multiple child `backint` processes to restore the nonsnapshotable files.
6. The restore workflow proceeds for the snapshotable files as described in the topic about the ProtectPoint with VMAX restore workflow or the ProtectPoint with RecoverPoint restore workflow in Chapter 1.

After the data is restored, the database administrator must recover the database by using the `brrecover` command or the Oracle SQL Plus tool. The SAP with Oracle documentation provides details.

Configuration of ProtectPoint operations in an SAP with Oracle environment

Ensure that the VMAX, XtremIO, RecoverPoint, and Data Domain configurations have been completed according to the ProtectPoint documentation. The required storage resources must be configured and provisioned properly to enable ProtectPoint operations.

Complete the following tasks to enable ProtectPoint operations:

- Ensure that the `ddbmsmd` program is started from the `/opt/dpsapps/dbappagent/bin` directory.
- For ProtectPoint for VMAX operations only, ensure that the supported VMAX Solutions Enabler version is installed and configured in local mode on each production host. The online software compatibility guide at <http://compatibilityguide.emc.com:8080/CompGuideApp/> describes the supported versions.

The Solutions Enabler database must be up-to-date on any host where a backup or recovery might run. To update the Solutions Enabler database, run the `symcfg discover` command. The Solutions Enabler documentation provides details.

Ensure that the required gatekeepers are also configured as described in the *ProtectPoint Version 4.0 Primary and Protection Storage Configuration Guide*. Solutions Enabler uses the small gatekeeper devices for communication with the VMAX storage array.

[Database application agent ProtectPoint operations with Data Domain usage limits](#) on page 44 provides general guidelines on the Data Domain usage limit settings for ProtectPoint operations.

Note

For ProtectPoint backups, it is recommended that database control files and online redo log files be located on different LUNs than the Oracle datafiles and archived logs. The Oracle documentation describes in the best practices for the database file layout.

You must complete the required configurations of the database application agent to enable the ProtectPoint operations in an SAP with Oracle environment. The following topics provide the product configuration details.

[SAP with Oracle RAC and cluster requirements for ProtectPoint operations](#) on page 303 provides additional details on the specific configuration requirements in an SAP with Oracle RAC or active-passive cluster environment.

The troubleshooting section at the end of this chapter provides details about limitations in the ProtectPoint operations with the database application agent in an SAP with Oracle environment.

Integrating the product into the BR*Tools environment

To enable the BR*Tools operations to use the `backint` program, you must complete the required settings in the BR*Tools configuration file `init<DBSID>.sap`:

- Define the directory from which BR*Tools calls the `backint` program.
By default, BR*Tools calls `backint` from the `sapexe` directory. You can set the `util_path` parameter to the directory pathname where `backint` is located. For example:

- On UNIX:

```
util_path = /opt/dpsapps/dbappagent/bin
```

- On Windows:

```
util_path = C:\PROGRA~1\DPSAPPS\DBAPPAGENT\bin
```

Note

On Windows, you must specify the short version of the directory pathname `C:\Program Files\DPSAPPS\DBAPPAGENT\bin` because you must not include spaces in file pathnames that you set in `init<DBSID>.sap`. If you include spaces in a pathname setting, the backup will fail. To obtain the short version of a directory, run the `dir /x` command in the parent directory. For example:

```
C:\> dir /x
```

```

      :
      :
11/21/2014  07:38 AM    <DIR>          PROGRA~1      Program Files

```

Here, the short version of “Program Files” is `PROGRA~1`.

- Set the backup medium to use the `backint` program. Set the `backup_dev_type` parameter to one of the following values:

```
backup_dev_type = util_file
```

or

```
backup_dev_type = util_file_online
```

Set `backup_dev_type = util_file_online` for online backups to decrease the amount of time each tablespace remains in hot backup mode, resulting in a smaller number of generated transaction logs.

You can override this `backup_dev_type` setting with the `-d` option when you perform a BR*Tools operation from the command line. For example:

```
brbackup -d util_file_online
```

- Set the `util_par_file` parameter to the location of the configuration file as configured in [Configuring the SAP with Oracle parameters](#) on page 294. For example:

```
util_par_file = ?/dbs/init<DBSID>.utl
```

where `?` is `$ORACLE_HOME`.

If you do not specify the complete pathname of this configuration file, the software searches for the file under the following default directory:

- On UNIX: `$ORACLE_HOME/dbs`
- On Windows: `%ORACLE_HOME%\database`

You can override this `util_par_file` setting by specifying the configuration file pathname with the `-r` option when you perform a BR*Tools operation from the command line. For example:

```
brbackup -r pathname/init<DBSID>.utl
```

The SAP BR*Tools documentation provides more details about parameters in the BR*Tools configuration file.

Confirming the environment and file permissions

You must confirm the settings of the database server environment and file permissions before you perform any SAP with Oracle operations.

The DBA operating system group must have read access to datafiles that will be restored. This read access enables the restore of the data by a different database user to a different host.

You can restore only datafiles for which you have read permission, based on the files' operating system permissions at the time that the files were backed up.

On UNIX, the read permission is associated with the user ID (UID) and group ID (GID), not the username or group name. The UID and GID of the user performing the restore must match the IDs associated with the files at backup time.

Enabling administrator privileges for SAP with Oracle restores on Windows

On specific types of Microsoft Windows systems, User Account Control (UAC) is designed to provide additional operating system security by preventing software from being installed or run unless an administrator authorizes the elevated privileges.

On Windows systems with UAC enabled, prior to the start of a restore from the CLI, ensure that administrator privileges are enabled for the user that will perform the restore. Otherwise, the operation might fail.

You can enable administrator privileges from the CLI as follows.

Procedure

1. Right-click the **Command Prompt** icon.
2. Select **Run as administrator** from the list.

Configuring the SAP with Oracle parameters

You must set the required parameters for SAP with Oracle operations in the configuration file named `init<DBSID>.utl`.

[Setting up the configuration file](#) on page 78 describes the common parameters, ProtectPoint parameters, and how to set the parameters in the configuration file. Other topics in [Product Configuration](#) on page 77 describe the parameters and requirements for the restores of replicated backups and rollback restores.

Ensure that the configuration file contains any other required parameters from the following table. For each parameter, the table lists the section heading of the configuration file section that contains the parameter. You must specify the location of the configuration file in the BR*Tools configuration file as described in [Integrating the product into the BR*Tools environment](#) on page 292.

After the configuration file is set up, ensure that the required lockbox procedures have been performed as described in [Configuring the lockbox](#) on page 103.

NOTICE

The `SNAPSHOT_OBJECTS` parameter setting determines whether to perform a ProtectPoint backup or a DD Boost backup.

Table 25 SAP with Oracle parameters for ProtectPoint operations

Parameter: ARCH_LOGS_SEQUENTIAL

Section: [GENERAL]

Specifies whether the `brarchive` program backs up archive logs in alphabetical order to optimize the log removal (cleanup) during an archive log backup.

Setting this parameter to `TRUE` causes the `brarchive` program to back up archive logs in alphabetical order, which speeds up the log cleanup process during the archive log backup.

The parameter is ignored during a `brbackup` or `brrestore` operation.

Optional for backups of archive logs with `brarchive`.

Valid values:

- `FALSE` (default).
- `TRUE`.

Note

When you use the nondefault value `TRUE`, the load balancing parameter `GROUP_BY_FS` is ignored during the `brarchive` backup.

Parameter: GROUP_BY_FS

Section: [GENERAL]

If you set this parameter to `TRUE`, then the operation ignores the `SAVESETS` parameter and groups the files by file system instead of file size. This efficient grouping of files can improve the performance of backups, index searches, and restore times.

Note

Setting this parameter to `TRUE` means that all the files being processed must be visible within the local file system. Windows UNC pathnames must be mapped to a local drive letter.

Optional for a backup.

Valid values:

- `FALSE` (default).
- `TRUE`.

Parameter: PARALLELISM

Section: [GENERAL]

Specifies the number of concurrent data streams to send to or from the Data Domain system during a DD Boost backup or restore operation, for each `backint` program that SAP Oracle runs for the operation.

This setting applies to a DD Boost backup or restore of specific SAP Oracle objects, such as metadata, that can occur in conjunction with a ProtectPoint operation.

Table 25 SAP with Oracle parameters for ProtectPoint operations (continued)

<p>Note</p> <p>The parallelism value for a backup is reduced if the value is greater than the <code>SAVESETS</code> parameter value.</p> <p>For a ProtectPoint backup or restore, only a single <code>backint</code> process performs the ProtectPoint operations.</p> <hr/> <p>Optional for a backup or restore.</p> <p>Valid values:</p> <ul style="list-style-type: none"> • 8 (default). • Positive integer number.
<p>Parameter: RAW_DEVICE_SKIP_BYTES</p> <p>Section: [GENERAL]</p> <p>Specifies to omit unwanted data from recovery on a raw device. Specifies the initial amount of data to skip.</p> <p>NOTICE</p> <p>When you adjust the amount of data to skip, you require expert knowledge of the operating system and the volume manager specifications.</p> <hr/> <p>Valid values:</p> <ul style="list-style-type: none"> • Undefined (default). • Valid device name, followed by an equal sign and the amount of data to skip in bytes, kilobytes, megabytes, or gigabytes: <code>RAW_DEVICE_SKIP_BYTES=/raw/dev1=skip[B K M G]</code> <p>Specify multiple devices by separating the devices with a colon. For example, the following setting skips 8 KB for the raw device <code>/dev/rdisk/c2t4d0s5</code>:</p> <pre>RAW_DEVICE_SKIP_BYTES=/dev/rdisk/c2t4d0s5=8K</pre>
<p>Parameter: RAW_DIR</p> <p>Section: [GENERAL]</p> <p>Specifies the directories for raw disk partitions on UNIX only. Any files below these directories are considered to be raw disk partitions.</p> <p>Do not set this parameter to the <code>/dev</code> or <code>/devices</code> directories, which are already treated as raw devices.</p> <hr/> <p>Note</p> <p>The listed partitions are pointed to by the symbolic links under <code>\$SAPDATA_HOME/sapraw/<TBS>_<number>/</code>.</p> <hr/> <p>Valid values:</p> <ul style="list-style-type: none"> • Undefined (default). • Directory pathnames of the raw disk partitions, separated by a semicolon (;). For example: <code>RAW_DIR = /oracle/data;sap/data</code>
<p>Parameter: RELOCATION_DEST</p>

Table 25 SAP with Oracle parameters for ProtectPoint operations (continued)**Section: [GENERAL]**

For a relocated restore, specifies a new root directory for SAP datafiles when `SAPDATA_HOME` is changed but the structure of subdirectories under `SAPDATA_HOME` is not changed.

Valid values:

- Undefined (default).
- Same value as the new `SAPDATA_HOME` value.

Parameter: SAVESETS**Section: [GENERAL]**

Specifies the number of save sets created on the Data Domain target. The backup data files are divided into the save sets based on file size.

This parameter is ignored if the `GROUP_BY_FS` parameter is set to TRUE.

Valid values:

- 20 (default).
- Positive integer number.

Parameter: SNAPSHOT_OBJECTS**Section: [GENERAL]**

Specifies whether to perform a ProtectPoint backup or a DD Boost backup of the SAP Oracle database files, archived redo log files, and metadata files. The metadata includes the SAP Oracle profiles and catalogs.

The database application agent performs a DD Boost backup of any files that do not reside on a snapshotable volume and any file type that is not specified in this parameter setting.

Note

It is recommended that you perform a DD Boost backup of the metadata because the metadata files are very small in size and number.

Mandatory for a ProtectPoint backup.

Valid values:

- Undefined (default) = Perform a DD Boost backup of the database files, log files, and metadata.
- One or more of the following values, separated by a comma:
 - DATA = Perform a ProtectPoint backup of the database files.
 - LOGS = Perform a ProtectPoint backup of the archived redo log files.
 - METADATA = Perform a ProtectPoint backup of the metadata.

Note

METADATA must be used in combination with DATA or LOGS or both.

Table 25 SAP with Oracle parameters for ProtectPoint operations (continued)

For example, the following setting specifies a ProtectPoint backup of the database and log files and a DD Boost backup of the metadata:

```
SNAPSHOT_OBJECTS = DATA, LOGS
```

Preparing for restore of archived logs

After you perform a number of SAP with Oracle backups, the backed-up archived logs might be in multiple sets of static images. During the restore of Oracle archived logs, a separate set of restore devices must be available to mount each set of static images.

Before you start a restore of archived logs, ensure that you have the required number of restore devices in the DD vdisk device pool. This number of restore devices must be at least equal to the number of VMAX or XtremIO source LUNs multiplied by the number of backups, as required by the specific restore.

Note

The restore devices or LUNs do not need to be dedicated to the SAP with Oracle server. The restore devices can be in a pool that is also used for other application or server restores that might run at different times, as long as the devices are masked accordingly.

Ensure that you perform a point-in-time restore of archived logs, not a rollback restore.

For example, the backup policy specifies a daily full database backup and the backup of the archived logs four times a day. To enable the restore from a particular database backup and the application of all the required logs, you need four times the number of VMAX or XtremIO source LUNs where the archived logs are located.

Preparing the Data Domain device for restore on Windows

On Windows in a ProtectPoint with RecoverPoint environment, you must prepare the Data Domain vdisk device before you can restore a ProtectPoint with RecoverPoint backup to an XtremIO array.

After the Data Domain block services have been created for the vdisk device according to the ProtectPoint documentation, complete the following steps.

Procedure

1. To bring the device online, use the Disk Manager.
2. If bringing the device online fails because the device is in an unknown state:
 - a. To take the device offline, use the Windows `diskpart` command.
 - b. To bring the device online, use the Disk Manager.

Preparing for rollback restores of SAP with Oracle ProtectPoint backups

Before you perform a rollback restore of a SAP with Oracle ProtectPoint backup, ensure that you have copied the SAPBACKUP directory to an alternate location and set the SAPBACKUP environment variable accordingly.

During a rollback restore, SAP records the list of datafiles for the restore under the SAPBACKUP directory. By default, this directory is located under one of the file

systems that are involved in the rollback restore. Unless you specify a new location for the SAPBACKUP directory prior to the rollback restore, the directory is considered "in use" when the restore departs the existing file systems and the rollback restore fails.

You can perform the following steps to reconfigure the SAPBACKUP directory to a non-default location.

Procedure

1. Copy the existing SAPBACKUP directory to a new location. For example:

```
cd /mnt/oracle/CER
cp -r -f sapbackup /newLocation
```

2. Set the SAPBACKUP environment variable to the new directory location. For example:

```
export SAPBACKUP=/newLocation/sapbackup
```

Preparing for SAP with Oracle ProtectPoint with RecoverPoint backups and rollback restores that use RecoverPoint pre-5.0

With RecoverPoint pre-5.0, the database application agent performs a rollback restore of a ProtectPoint with RecoverPoint backup at the consistency group level. If the RecoverPoint consistency group being restored contains multiple LUNs, then all those LUNs are overwritten and inaccessible during the rollback restore. Specific requirements apply to the SAP with Oracle ProtectPoint with RecoverPoint backups and rollback restores.

Ensure that you follow the requirements and recommendations in [Configuring rollback restores of ProtectPoint backups](#) on page 97.

Note

With RecoverPoint pre-5.0, a ProtectPoint with RecoverPoint backup and rollback restore always occurs at the consistency group level, regardless of which objects are included in the backup command. As a best practice for the ProtectPoint with RecoverPoint rollback restore, when you perform the backup or rollback restore, do not exclude the logs or any database files that are part of the RecoverPoint consistency group being backed up or restored.

Ensure that you meet the following requirements for the SAP with Oracle rollback restore of a ProtectPoint with RecoverPoint backup with RecoverPoint pre-5.0:

- The Oracle control files and online redo logs reside on conventional nonsnapshot devices or in a separate RecoverPoint consistency group.
- If an Oracle tablespace must remain online during a rollback restore, then this online tablespace and the tablespace to be restored are in separate consistency groups.
- The Oracle datafiles and archived log files are in separate consistency groups.
- The BR*Tools catalog files and Oracle datafiles are in separate consistency groups.

Note

It is recommended that you perform a DD Boost backup of the BR*Tools metadata files, which include the catalog files, because the metadata files are very small in size. You can specify the DD Boost backup of the metadata by omitting the METADATA value from the `SNAPSHOT_OBJECTS` parameter setting.

Performing ProtectPoint backups and recovery with SAP BR*Tools

You can perform ProtectPoint operations with SAP BR*Tools after you have completed the backup configurations in [Configuration of ProtectPoint operations in an SAP with Oracle environment](#) on page 291.

You can run the BR*Tools CLI or a supported GUI to perform ProtectPoint backup, restore, or recovery operations with the database application agent.

The SAP and Oracle documentation provides details about the available options.

Performing ProtectPoint backups with BR*Tools

To perform database and archived redo log backups, you can run the `brbackup` and `brarchive` commands with the appropriate options. For example:

- Offline backup of the whole database:

```
brbackup -m all -t offline -d util_file
```

- Online backup of a single tablespace:

```
brbackup -m system -t online -d util_file_online
```

To back up the archived redo logs if you perform an online backup, you can run the `brarchive` command:

```
brarchive
```

Note

Without the redo logs, you can recover a database that was backed up online only to the time of its last full backup.

To back up nondatabase files or directories, you can run the following command:

```
brbackup -m {sap_dir | ora_dir | all_dir | <full_directory_path> | <nondatabase_file_path>}
```

To ensure that the backup is readable and complete after a backup is finished, you can run the `brbackup` command with the `-verify` option.

Performing ProtectPoint restore and recovery with BR*Tools

To ensure that the backup to be restored exists before you perform a restore or recovery, you can run the `brrestore` command with the `-verify` option. The BR*Tools documentation provides details about the `-verify` option.

For example, the following command inquires for the latest backup:

```
brrestore -b last -verify only_conf
```

To perform a restore or recovery, you can use the `brrestore` command or `brrecover` interface. The `brrestore` command restores the backed-up database files, and you must run Oracle SQL Plus to recover the database by applying the transaction logs to roll forward the database to a specific point-in-time. The `brrecover` interface automates the restore and recovery process by calling `brrestore` and SQL Plus to perform specific steps as needed.

For example, to restore the entire database, you can run the following command:

```
brrestore -u / -m all -d util_file -c force
```

To restore a single tablespace, you can run the following command:

```
brrestore -u / -m PSAPSR3DB -d util_file -c force
```

NOTICE

The `brrestore` program used for a restore operation, whether run directly from the command line or indirectly from the `brrecover` command, first deletes all the original files to be restored before the program runs `backint` to restore the backed-up files. You will lose the original files if `brrestore` or `backint` fails.

Perform one of the following actions to prevent these issues:

- Restore the files to a different location that does not include any files by using the `-m` option and specifying the restore destination. For example:

```
brrestore -m tablespace_name=<restore_directory>
```

- To prevent deletion of the original files by `brrestore`, use the `brrestore -NFD` option.

To restore data from a directory (nondatabase) backup, use the `brrestore -m non_db` command. When you restore the SAP directory data, prevent deletion of the BR*Tools and `backint` binaries during the restore by performing a relocated restore. For example, run the following command:

```
brrestore -m non_db=<restore_directory>
```

Performing a ProtectPoint restore to an alternate host

To optionally restore the data to a different SAP Oracle host (destination host) than the one that was backed up:

- Set the `CLIENT` parameter to the hostname of the host where the data was backed up.
- Follow the disaster recovery steps to re-create the environment and restore the configuration files and BR*Tools logs. [Preparing for SAP with Oracle disaster recovery](#) on page 302 provides details.

Preparing for SAP with Oracle disaster recovery

For a comprehensive disaster recovery plan, you must ensure that you can reconstruct the computing environment and all the SAP Oracle server files associated with maintaining data on the application host.

Use the following guidelines to prepare for a disaster recovery of the SAP with Oracle server host:

- Maintain accurate and complete records of the network and system configurations. Keep all the original software media and the following items in a safe location:
 - Original operating system media and patches
 - Device drivers and device names
 - File system configuration
 - IP addresses and hostnames
- To ensure the recovery of customer-specific settings of the database, back up the system configuration files by following the SAP with Oracle documentation. These files are not backed up during database backups.
- Ensure that you have a current full backup of the database and all the archived redo logs as required for a recovery of the database to the current time.

To recover from a disaster, you must first restore any lost Oracle and SAP configuration files and lost BR*Tools backup log files according to the following instructions. After you restore these files, you can perform a database point-in-time recovery or a whole database reset according to the instructions in [Performing ProtectPoint restore and recovery with BR*Tools](#) on page 301.

Restoring the required Oracle and SAP BR*Tools files

The following procedure is a concise version of the disaster recovery steps described in the SAP documentation, modified for the specific requirements of the database application agent. You can perform this procedure on the original host or a new host, both referred to as the destination host.

To restore to a different host than the one backed up, follow the guidelines in [Performing a ProtectPoint restore to an alternate host](#) on page 302.

Procedure

1. If the entire SAP Oracle system is lost:
 - a. Reinstall all the required SAP and Oracle software components according to the SAP and Oracle documentation.

- b. Reconfigure the SAP data layout, such as the `SAPDATA_HOME` directory and its subdirectories, to the same state as before the disaster.
2. Configure the BR*Tools configuration file, `init<DBSID>-dr.sap`, and the configuration file of the database application agent, `init<DBSID>-dr.utl`, for recovery on the destination host.
3. Re-create the lockbox on the destination host. [Configuring the lockbox](#) on page 103 provides details about the lockbox.
4. Ensure that the parameters are correctly set for a redirected restore according to [Configuring the SAP with Oracle parameters](#) on page 294.
5. To perform a disaster recovery of profiles and logs on the destination host, start BR*Tools on the host, and then follow the onscreen instructions.

For example, the following steps show how to use BR*Tools for disaster recovery to restore profiles and log files from a `BRBACKUP` backup.

- a. In the `brtools` menu, select **Restore and Recovery**, and then **Disaster recovery**.
- b. On the page **BRRECOVER options for disaster recovery**, set the location of the BR*Tools configuration file and the SAP with Oracle configuration file to the file names configured in step 2.

Note

If you are restoring the original BR*Tools configuration file or SAP with Oracle configuration file, ensure that the files used for the disaster recovery have different names or are stored in a different location than the original files to be restored.

- c. On the page **Device type for restoring profiles and log files from BRBACKUP backup**, select **Backup utility**.
- d. On the page **Parameters for restoring profiles and log files from BRBACKUP backup utility backup**, specify the files to restore.
- e. On the page **Restore of profiles and log files from BRBACKUP backup**, select the components that you want to restore.

The SAP documentation provides more details about disaster recovery.

Recovering an SAP Oracle database after disaster

To perform database point-in-time recovery or database reset, after you restore the correct SAP BR*Tools configuration file and logs on the application host, follow the normal BR*Tools recovery procedure.

[Performing ProtectPoint restore and recovery with BR*Tools](#) on page 301 provides the procedure to recover an SAP with Oracle database.

The SAP documentation provides more details about database recovery.

SAP with Oracle RAC and cluster requirements for ProtectPoint operations

You can perform ProtectPoint backups and restores in an Oracle RAC or active-passive cluster environment.

Active-passive cluster requirements

You must meet the following configuration requirements in an active-passive cluster environment:

- You have set up the database in the cluster according to the appropriate database server documentation.
- You have installed the database application agent on each node of the cluster that will participate in backups or restores.
- You have set the `CLIENT` parameter in the configuration file to the hostname of the virtual node. [Common parameters](#) on page 80 provides details on the parameter.

Oracle RAC requirements

The SAP documentation provides details about the Oracle RAC setup. You must meet the following configuration requirements in an Oracle RAC environment:

- One of the Oracle RAC instances, set as the dedicated database (DDB) instance, is used to perform all the database administration tasks. You must be able to administer all the RAC instances from the DDB instance.
- You have installed BR*Tools and the database application agent on the DDB instance host.
- You have created the BR*Tools log directories, such as `SAPBACKUP` and `SAPARCH`, on a shared file system.
- You have correctly configured Oracle SQL Net.
- You have set the required parameters, such as `parallel_instances` and `db_services`, in the SAP initialization file, `init<DBSID>.sap`, located on the DDB instance.

Oracle RAC requirements for the backint interface

All the Oracle RAC backups performed with BR*Tools and the `backint` interface run on a single RAC instance, the DDB instance. The configuration of the DDB instance is the same as for a stand-alone Oracle system.

If you change the DDB instance after some backups were performed with the original DDB instance, set the parameter `CLIENT=<original_DDB_instance_hostname>` to ensure that all the backups are stored in the same location on the Data Domain system. [Common parameters](#) on page 80 provides details on the parameter.

ProtectPoint restore and rollback for VCS on Solaris

Use the procedures in the following topics to perform ProtectPoint restore and rollback operations for a VCS system on Solaris.

Performing a ProtectPoint VCS restore

Procedure

1. On the primary VCS node, perform the following steps as the root user.

a. List the VCS Service Groups:

```

root:/# hastatus -sum

-- SYSTEM STATE
-- System          State          Frozen
A ledma054         RUNNING        0
A ledma056         RUNNING        0

-- GROUP STATE
-- Group           System          Probed
AutoDisabled      State
B ClusterService ledma054         Y
N                 ONLINE
B ClusterService ledma056         Y
N                 OFFLINE
B oracle_ctl_sg   ledma054         Y
N                 ONLINE
B oracle_ctl_sg   ledma056         Y
N                 OFFLINE
B oracle_sg       ledma054         Y
N                 ONLINE
B oracle_sg       ledma056         Y
N                 OFFLINE
B vxfen           ledma054         Y
N                 ONLINE
B vxfen           ledma056         Y
N                 ONLINE

```

b. Enable the VCS configuration as Read/Write:

```

root:/# haconf -makerw

```

c. Freeze the VCS service groups by disabling On line/Off line. Type the following command:

Note

This is an example of a VCS and Oracle configuration.

```

root:/# hagrps -freeze <oracle_sg> -persistent

```

d. Confirm the VCS status by typing the following command:

```

root:/# hastatus -sum

-- SYSTEM STATE
-- System          State          Frozen
A ledma054         RUNNING        0
A ledma056         RUNNING        0

-- GROUP STATE
-- Group           System          Probed
AutoDisabled      State
B ClusterService ledma054         Y
N                 ONLINE

```

```

B ClusterService ledma056 Y
N OFFLINE
B oracle_ctl_sg ledma054 Y
N ONLINE
B oracle_ctl_sg ledma056 Y
N OFFLINE
B oracle_sg ledma054 Y
N ONLINE
B oracle_sg ledma056 Y
N OFFLINE
B vxfen ledma054 Y
N ONLINE
B vxfen ledma056 Y
N ONLINE

-- GROUPS FROZEN
-- Group
C oracle_ctl_sg
C oracle_sg

-- RESOURCES DISABLED
-- Group Type Resource
H oracle_ctl_sg DiskGroup oracle_ctl_dg_DG_res1
H oracle_ctl_sg Mount oracle_ctl_dg_MNT_res1
H oracle_ctl_sg Volume oracle_ctl_dg_VOL_res1
H oracle_sg DiskGroup oracle_dg_DG_res1
H oracle_sg Mount oracle_dg_MNT_res1
H oracle_sg Volume oracle_dg_VOL_res1

```

e. Make the VCS configuration as Read Only. Type the following command:

```
root:/# haconf -dump -makero:
```

2. On the primary VCS node, perform the following steps as the Oracle user.

a. Run the `shutdown` and `startup mount` commands on the Oracle database:

a. `oracle:/# sqlplus / as sysdba`

b. `SQL > shutdown immediate`

c. `SQL > startup mount`

d. `SQL > exit`

b. Perform the RMAN restore and recovery.

3. On the primary VCS node, perform the following steps as the root user.

a. Make the VCS configuration Read/Write. Type the following command:

```
root:/# haconf -makerw
```

- b. Unfreeze the service groups, and allow On line/Off line. Type the following command:

```
root:/# hagr -unfreeze <oracle_ctl_sg> -persistent
root:/# hagr -unfreeze <oracle_sg> -persistent
```

- c. Confirm the VCS status. Type the following command:

```
root:/# hastatus -sum
```

```
-- SYSTEM STATE
-- System          State          Frozen
A ledma054         RUNNING       0
A ledma056         RUNNING       0

-- GROUP STATE
-- Group           System          Probed
AutoDisabled     State
B ClusterService ledma054       Y
N                 ONLINE
B ClusterService ledma056       Y
N                 OFFLINE
B oracle_ctl_sg  ledma054       Y
N                 ONLINE
B oracle_ctl_sg  ledma056       Y
N                 OFFLINE
B oracle_sg      ledma054       Y
N                 ONLINE
B oracle_sg      ledma056       Y
N                 OFFLINE
B vxfen          ledma054       Y
N                 ONLINE
B vxfen          ledma056       Y
N                 ONLINE
```

Performing a ProtectPoint VCS rollback

Note

A rollback fails if you change the style of the mpio device name. The rollback to the source LUN is successful. However, the fsck and mount fails. In this scenario, manually mount the FS.

Procedure

1. On the primary VCS node, perform the following steps as the root user.
 - a. List the VCS Service Groups:

```
root:/# hastatus -sum
```

```
-- SYSTEM STATE
-- System          State          Frozen
A ledma054         RUNNING       0
A ledma056         RUNNING       0

-- GROUP STATE
```

```

-- Group          System          Probed
AutoDisabled     State
B ClusterService ledma054      Y
N                ONLINE
B ClusterService ledma056      Y
N                OFFLINE
B oracle_ctl_sg  ledma054      Y
N                ONLINE
B oracle_ctl_sg  ledma056      Y
N                OFFLINE
B oracle_sg      ledma054      Y
N                ONLINE
B oracle_sg      ledma056      Y
N                OFFLINE
B vxfen          ledma054      Y
N                ONLINE
B vxfen          ledma056      Y
N                ONLINE

```

b. Enable the VCS configuration as Read/Write:

```
root:/# haconf -makerw
```

c. Freeze the VCS service groups by disabling On line/Off line. Type the following command:

Note

This is an example of a VCS and Oracle configuration.

```

root:/# hagrps -freeze <oracle_sg> -persistent
root:/# hagrps -freeze <oracle_ctl_sg> -persistent

```

d. Confirm the VCS status, by typing the following command:

```
root:/# hastatus -sum
```

```

-- SYSTEM STATE
-- System          State          Frozen
A ledma054         RUNNING       0
A ledma056         RUNNING       0
-- GROUP STATE
-- Group          System          Probed
AutoDisabled     State
B ClusterService ledma054      Y
N                ONLINE
B ClusterService ledma056      Y
N                OFFLINE
B oracle_ctl_sg  ledma054      Y
N                ONLINE
B oracle_ctl_sg  ledma056      Y
N                OFFLINE
B oracle_sg      ledma054      Y
N                ONLINE
B oracle_sg      ledma056      Y
N                OFFLINE
B vxfen          ledma054      Y
N                ONLINE

```

```

B vxfen          ledma056          Y
N                ONLINE

-- GROUPS FROZEN
-- Group
C oracle_ctl_sg
C oracle_sg

-- RESOURCES DISABLED
-- Group          Type          Resource
H oracle_ctl_sg  DiskGroup  oracle_ctl_dg_DG_res1
H oracle_ctl_sg  Mount      oracle_ctl_dg_MNT_res1
H oracle_ctl_sg  Volume     oracle_ctl_dg_VOL_res1
H oracle_sg      DiskGroup  oracle_dg_DG_res1
H oracle_sg      Mount      oracle_dg_MNT_res1
H oracle_sg      Volume     oracle_dg_VOL_res1

```

e. Make the VCS configuration as Read Only. Type the following command:

```
root:/# haconf -dump -makero:
```

2. On the primary VCS node, perform the following steps as the Oracle user.

a. Run the shutdown and startup mount commands on the Oracle database:

```
a. oracle:/# sqlplus / as sysdba
```

```
b. SQL > shutdown immediate
```

```
c. SQL > startup mount
```

```
d. SQL > exit
```

b. Perform the RMAN rollback and recovery.

3. On the primary VCS node, perform the following steps as the root user.

a. Make the VCS configuration Read/Write. Type the following command:

```
root:/# haconf -makerw
```

b. Unfreeze the service groups, and allow On line or Off line. Type the following command:

```
root:/# hagrps -unfreeze <oracle_ctl_sg> -persistent
root:/# hagrps -unfreeze <oracle_sg> -persistent
```

c. Confirm the VCS status. Type the following command:

```
root:/# hastatus -sum
```

```

-- SYSTEM STATE
-- System          State          Frozen

```

```

A ledma054          RUNNING          0
A ledma056          RUNNING          0

-- GROUP STATE
-- Group           System           Probed
AutoDisabled      State
B ClusterService  ledma054          Y
N                  ONLINE
B ClusterService  ledma056          Y
N                  OFFLINE
B oracle_ctl_sg   ledma054          Y
N                  ONLINE
B oracle_ctl_sg   ledma056          Y
N                  OFFLINE
B oracle_sg       ledma054          Y
N                  ONLINE
B oracle_sg       ledma056          Y
N                  OFFLINE
B vxfen           ledma054          Y
N                  ONLINE
B vxfen           ledma056          Y
N                  ONLINE

```

Note

The service groups will be faulted, but will come back online in a short time.

SAP with Oracle troubleshooting tips for ProtectPoint operations

[General troubleshooting tips](#) on page 132 provides common troubleshooting information that applies to the database application agent operations with all the supported databases and applications.

The database application agent maintains the operational and debugging logs in the standard directories:

- On UNIX or Linux: /opt/dpsapps/dbappagent/logs, linked to /var/opt/ddbda/logs
- On Windows: C:\Program Files\DPSAPPS\DBAPPAGENT\logs

ProtectPoint operations maintain a separate operational log named `ddbsm.log`, which has details about operations and errors.

You must set the `DEBUG_LEVEL` parameter to enable debugging.

APPENDIX A

Performance Optimization

This appendix includes the following topics:

- [Backup and recovery performance optimization](#)..... 312
- [Hardware component 70 percent rule](#).....312
- [Impact of software components on performance](#)..... 312
- [Performance optimization in DB2 systems](#).....313
- [Performance optimization in Oracle systems](#)..... 314
- [Performance optimization in SAP HANA systems](#)..... 314
- [Performance optimization in SAP with Oracle systems](#)..... 315

Backup and recovery performance optimization

Every backup environment has a bottleneck. The bottleneck determines the maximum throughput of the system. Backup and restore operations are only as fast as the slowest component in the chain.

When you set the backup and recovery performance expectations, consider the sizing requirements of the backup environment.

Consider the sizing requirements of the backup environment:

- Review the network infrastructure and the Data Domain storage before you set the performance expectations.
- Review and set the Recovery Time Objective (RTO) for the application.
- Determine the backup window.
- Determine the amount of data to be backed up during full backups, incremental backups, and log backups.
- Determine the data growth rate.
- Determine backup retention requirements.

Hardware component 70 percent rule

Manufacturer throughput and performance specifications that are based on theoretical environments are rarely, if ever, achieved in real backup and recovery environments. As a best practice, never exceed 70 percent of the rated capacity of any component.

Consider the following hardware components:

- CPU
- Storage
- Network
- Internal bus
- Memory
- Fibre Channel

Performance and response time significantly decrease when the 70 percent utilization threshold is exceeded.

Impact of software components on performance

The applications that run on the client host are the primary users of CPU, network, and I/O resources. The applications typically use a significant amount of these resources, which affects the backups. Backups can also be resource-intensive and can impact the performance of the primary applications.

Several components can impact the performance in system configurations:

- The application backups and restores are object-based, where an object can be a file or a stream. A backup or restore of many small objects usually takes longer than a backup or restore of a small number of large objects, even when the same amount of data is processed:

- Configure the transaction logging to generate larger log files when possible.
- For applications that support the explicit running of log backups, schedule the best log backup frequency that is based on the RTO.
- Usually the number of backups (save sets) that are stored on the Data Domain system does not affect the backup performance. However, for some applications, such as SAP HANA or DB2, the number of stored backups might affect the restore performance:
 - Use application tools to set the proper retention for backups when possible and ensure that obsolete backups are cleaned up on a regular basis.
 - Follow the instructions in [Configuring the display and deletion of save set information](#) on page 117 to find out the list and number of backup objects (save sets) that are stored on the Data Domain system.
 - Use different device paths for different database instances to increase the speed of recovery. A client system with four or five database instances can have five times the number of backups in a directory, which might lead to a slower restore for some applications.
- Backup encryption and compression are resource-intensive operations on the client, which can significantly affect the backup performance and the deduplication ratio of data that is stored on the Data Domain system:
 - Do not use application-based backup compression because the DD Boost software stores data in a compressed format.
 - Use Data Domain in-flight encryption to protect the data in transit. [Enabling encryption over a WAN connection](#) on page 52 provides details.
- Running parallel (multistream) backups and restores increases the speed of operations:
 - Ensure that the number of total or concurrent streams to the Data Domain system does not exceed the maximum number of supported streams.

The following topics describe the performance impact of numerous backups for each type of database application and provide information on how to improve performance.

Performance optimization in DB2 systems

You can optimize the performance in a DB2 system through the recommended practices for DB2 transaction log archiving, multistream backups, and backup deletions.

Configuring DB2 transaction log archiving

It is recommended that you keep the database active so that the archived logs are backed up in time. This practice enables you to avoid the backup of many archived logs, which slows the concurrent database backups.

Note

The number of backed-up logs in the DB2 backup storage might affect the performance of rollforward operations.

For example, when you restore to a nonexistent database with a rollforward, the following rollforward operation might require a complete scan of all the backed-up log entries. You might restore to a nonexistent database during a disaster recovery when all the database data has been lost. The scan addresses a DB2 request to query for the highest log chain available.

Configuring multistream backups

The restore time of a DB2 multistream backup includes the time to search for the streams. How the data is streamed in the backup affects the backup deduplication ratio and the performance of both the backup and restore. Ensure that you complete an optimal configuration for a DB2 multistream backup. The following IBM article provides more details:

<http://www.ibm.com/developerworks/data/library/techarticle/dm-1302db2deduplication>

Configuring DB2 backup deletion for DD Boost operations

Use the DB2 automatic deletion of recovery objects to remove or prune the backup objects that become obsolete. When a database image is pruned, all the associated archived log backups are also removed.

If the configuration of database and log pruning is incorrect, the pruning fails silently. In that case, the DB2 recovery history is not cleaned up and obsolete backups are not removed. This situation can also interfere with subsequent backups and restores. To prevent the problem, ensure that the backup deletion is configured correctly, then periodically monitor the db2 diagnostic log for pruning issues and fix the issues in a timely manner. [DB2 troubleshooting tips for DD Boost operations](#) on page 174 provides details about possible error cases, such as DB2 pruning issues, and resolutions.

Managing and deleting ProtectPoint DB2 backups

The DB2 automatic deletion of recovery objects does not apply to DB2 snapshot backups. Run the `db2acsutil` command manually to clean up the obsolete snapshots periodically.

Performance optimization in Oracle systems

The number of backups in the Oracle backup storage does not affect the performance of restore operations. It is recommended that you keep only the backups that are required for RTO, which saves the backup storage space.

Set up the Oracle retention policy, which is based on the recovery window or redundancy, to make backups obsolete. Delete the obsolete backups regularly by running the `rman delete obsolete` command. Follow the instructions on performing Oracle backup deletion and maintenance operations in the Oracle chapters of this guide.

Performance optimization in SAP HANA systems

The number of backups in the SAP HANA backup storage can decrease the performance of restore operations. Evidence has shown that the time required for restore-related operations increases with the number of backups that are stored on the Data Domain system.

It is recommended that you prevent the creation of large numbers of redo log backups, to optimize the restore time.

Configuring automatic backups of SAP HANA redo logs

By default, SAP HANA backs up the redo logs for databases every 15 minutes, which results in many small backups if the database is not busy. If allowed by the RTO, increase the log backup interval to generate a smaller number of larger backups.

Consider backing up the SAP HANA data and logs into separate device paths on the Data Domain system. You can achieve this by specifying separate utility files for the data and log backups. For example, specify that the utility file for the data backups

uses the `DEVICE_PATH=/ device_path/DATA` setting, and the utility file for the log backups uses the `DEVICE_PATH=/ device_path/LOGS` setting. As a result, SAP HANA scans fewer records during the restore.

Deleting DD Boost backups with SAP HANA Studio

Deleting the old backups regularly from the Data Domain system reduces the number of backup entries, which enables faster restores. Follow the instructions in [Deleting DD Boost backups by using SAP HANA Studio](#) on page 263.

The deletion of SAP HANA backups using Backint is a process that runs in the background after SAP HANA reports that the entries have been deleted from its own catalog. The deletion process time can be affected by a large number of backups.

Performance optimization in SAP with Oracle systems

The number of backups in the SAP with Oracle backup storage does not affect the performance of restore operations. It is recommended that you keep only the backups that are required for RTO, which saves the backup storage space.

SAP BR*Tools does not provide an interface for the deletion of backups that were performed with the `backint` program. To delete these backups, follow the instructions on backup deletion in [Configuring the display and deletion of save set information](#) on page 117.

Follow the Oracle RMAN instructions for the deletion and maintenance of backups that are performed with BR*Tools and the RMAN interface.

GLOSSARY

This glossary contains the definitions of terms found in this manual. Most of the terms are specific to the database application agent.

A

active-passive cluster Type of cluster configuration where the data server runs on the active node, and other nodes are passive nodes that maintain data updates and wait to take over if the active node fails.

B

backup 1. Duplicate of database data or application data or an entire computer system that is stored separately from the original, which you can use to recover the original if it is destroyed or damaged.
2. Operation that saves data to a volume for use during a recovery.

backup device Encapsulated LUN (eLUN) or FAST.X LUN on a VMAX system, which is created by encapsulating a DD vdisk LUN during a ProtectPoint backup. [See restore device](#)

backup level [See level](#)

C

client Database or application server whose data can be backed up and restored with the database application agent software.

cluster nodes A group of linked virtual or physical hosts with shared storage in a cluster, which work together and represent themselves as a single host called a virtual cluster host.

cold backup [See offline backup](#)

consistency group RecoverPoint group that protects a set of source LUNs (volumes). Two data sets that are dependent on each other, such as a database and a database log, should be part of the same consistency group. Logical components of a consistency group include copies, replication sets, and journals.

D

database application agent Software that enables the DD Boost and ProtectPoint operations that are performed through the Data Domain Boost for Enterprise Applications and ProtectPoint workflows. This agent was formerly known as DD Boost for Databases and Applications (DDBDA).

DD Boost	An optimized library and communication framework with a special Data Domain API that enables the backup software to define and interact with storage devices on the Data Domain system.
deduplication backup	Type of backup in which redundant data blocks are identified and only unique blocks of data are stored. When the deduplicated data is restored, the restore returns the data to its original native format.
deprecated feature	Feature that is supported in the current release of the software but will be unsupported and removed in a future release.
destination client	Computer to which a directed recovery restores the database data.
directed recovery	Method that recovers data that originated on one client host and re-creates it on a different client host, known as the destination client.
disaster recovery	Restore and recovery of business operations and data if a hardware failure or software corruption occurs.
distributed segment processing (DSP)	Part of the DD Boost interface, which enables data deduplication on a host before the data is sent to the Data Domain system for storage.

E

eLUN	Encapsulated LUN in a VMAX system, which is created by using the FAST.X software.
-------------	---

F

FAST.X	FAST.X software on the VMAX system that encapsulates LUNs for Data Domain storage, preserves existing data on the LUNs, and enables access to the external LUNs through the VMAX system.
firewall	A system designed to prevent unauthorized access to or from a private network.
full backup	See level

G

group set	User-defined set of RecoverPoint consistency groups that is used to perform operational and recovery activities. For ProtectPoint with RecoverPoint operations: <ul style="list-style-type: none"> You cannot enable parallel bookmarking for a group set. In a consistency group, the local copy exists on the Data Domain system and there is no journal volume for that local copy.
------------------	--

H

high-availability system	System of multiple computers configured as interconnected nodes on a network that ensures the application services continue despite a hardware failure or a software failure.
---------------------------------	---

host	Computer on a network.
hot backup	See online backup
I	
incremental backup	See level
L	
level	Backup configuration option that specifies how much data is saved during a backup: <ul style="list-style-type: none"> • A full backup backs up all data objects, regardless of when they last changed. • An incremental backup backs up only data objects that have changed since the previous backup.
O	
offline backup	Backup of database objects that is performed while the corresponding database or instance is shut down and unavailable to users.
online backup	Backup of database objects that is performed while the corresponding database or instance is running and available to users.
Oracle Recovery Manager (RMAN)	Oracle utility that acts as an intelligent interface to Oracle databases for the backup and restore of Oracle database objects.
P	
parallelism	Method that backs up or recovers data through multiple concurrent streams.
pathname	Set of instructions to the operating system for accessing a file: <ul style="list-style-type: none"> • An absolute pathname indicates how to find a file starting from the root directory and working down the directory tree. • A relative pathname indicates how to find a file starting from the current location.
primary Data Domain system	Data Domain system that stores the backups that are performed from a database or application host. The database application agent can back up data to a primary Data Domain system only.
ProtectPoint	An alternative workflow that provides block-based data protection from the primary storage to the protection storage. The database application agent supports the use of the ProtectPoint workflow for backups of application data on a VMAX or XtremIO system to a Data Domain system.

ProtectPoint backup Backup of DB2, Oracle, or SAP with Oracle data at the LUN level on a VMAX or XtremIO system to a Data Domain system. The Data Domain Boost for Enterprise Applications software and ProtectPoint software use the following technologies to jointly perform the backup:

- For a backup from VMAX, use the FAST.X and SnapVX technologies on the VMAX system and the vdisk and FastCopy technologies on the Data Domain system.
- For a backup from XtremIO, use the RecoverPoint splitter and consistency group technologies on the XtremIO system and the vdisk, FastCopy, and DD Boost technologies on the Data Domain system.

R

recover To restore data files from backup media to a client disk and apply transaction logs or redo logs to make the data consistent with a given point-in-time.

RecoverPoint Software system that includes the RecoverPoint splitters and RecoverPoint appliances (RPAs) that are used to replicate and protect data. The database application agent uses ProtectPoint and RecoverPoint software to perform the backups of database data on an XtremIO system to a Data Domain system.

RecoverPoint appliance (RPA) RecoverPoint's intelligent data protection appliance in the form of a physical or virtual machine that manages all aspects of reliable data replication. In a ProtectPoint with RecoverPoint backup, the RPA reads the snapshot data from an XtremIO system and uses DD Boost to transfer the data to working files on a Data Domain system.

restore To retrieve individual data files from backup media and then copy the files to disk, without applying transaction logs. [See recover](#)

restore device Encapsulated LUN (eLUN) or FAST.X LUN on a VMAX system, which is used during the restore of a ProtectPoint backup. [See backup device](#)

rollback restore Block-level restore that is provided by the primary storage array.
Rollback restore from a ProtectPoint for VMAX backup is the restore of the entire LUN. Rollback restore from a ProtectPoint with RecoverPoint backup is the restore of the entire RecoverPoint consistency group, volume group, or LUN, depending on the type of database server and the version of RecoverPoint.
Rollback restore always restores the data to the original source LUNs on the primary storage.

roll forward To apply transaction logs to a recovered database to restore the database to a state that is consistent with a given point-in-time.

rollforward recovery Type of DB2 database recovery that applies transaction logs to restore the database to a given point-in-time.

S

save file Operating system file or block of data, as the simplest object that you can back up or restore.

save set Collection of one or more save files created during the backup session.

secondary Data Domain system	Data Domain system from which you can restore replicated backups to a database or application host by using the database application agent. The Data Domain administrator replicates the backups from a primary Data Domain system to the secondary Data Domain system.
shared disk	Storage disk connected to multiple nodes in the cluster.
SnapVX snapshot	Snapshot created with VMAX SnapVX, which is a snapshot technology supported by specific VMAX arrays.
source LUN	LUN in the VMAX system where the original data resides.
static image	Copy of a VMAX SnapVX snapshot on the Data Domain system.
Symdev-ID	Device ID assigned by VMAX when a device or volume is created in the VMAX system.
T	
tablespace	Oracle database structure that consists of one or more data files.
transaction log	Record of named database transactions or a list of changed files in a database, which is stored in a log file to enable quick restore and rollback transactions.
V	
vdisk	Virtual Disk technology that is available in DD OS 5.5 and later.

