

HP Data Protector 6.20

Zero Downtime Backup Integration Guide

for Oracle, SAP R/3, Microsoft SQL Server,
Microsoft Exchange Server, and Microsoft
SharePoint Server



© Copyright 2004, 2011 Hewlett-Packard Development Company, L.P.

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Intel®, Itanium®, Pentium®, Intel Inside®, and the Intel Inside logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Microsoft®, Windows®, Windows XP®, and Windows NT® are U.S. registered trademarks of Microsoft Corporation.

Adobe and Acrobat are trademarks of Adobe Systems Incorporated.

Java is a registered trademark of Oracle and/or its affiliates.

Oracle® is a registered US trademark of Oracle Corporation, Redwood City, California.

UNIX® is a registered trademark of The Open Group.

Contents

Publication history.....	10
About this guide.....	11
Intended audience.....	11
Documentation set.....	11
Guides.....	11
Online Help.....	13
Documentation map.....	14
Abbreviations.....	14
Map.....	14
Integrations.....	15
Document conventions and symbols.....	16
Data Protector graphical user interface.....	16
General information.....	17
HP technical support.....	17
Subscription service.....	17
HP websites.....	17
Documentation feedback.....	18
1 Data Protector Oracle Server ZDB integration.....	19
Introduction.....	19
Backup and restore types.....	20
Integration concepts.....	21
Oracle backup set ZDB concepts.....	24
Backup process.....	26
Oracle proxy-copy ZDB concepts.....	28
Backup process.....	29
Configuring the integration.....	31
Prerequisites.....	31
Limitations.....	32
Before you begin.....	33
Backup set method.....	34
Enabling zero downtime backup on disk arrays of the HP P6000 EVA Disk Array Family for ASM configurations.....	35
Enabling zero downtime backup on disk arrays of the HP P9000 XP Disk Array Family for ASM configurations.....	36
Cluster-aware systems.....	36
Linking Oracle Server with the Data Protector MML.....	36
Configuring Oracle user accounts.....	36
Configuring Oracle operating system user accounts.....	37
Clusters.....	37
Configuring Oracle database users accounts.....	38
Configuring Oracle databases.....	38
Using the Data Protector GUI.....	39
Using the Data Protector CLI.....	41
Checking the configuration.....	44
Using the Data Protector GUI.....	44
Using the Data Protector CLI.....	44
Handling errors.....	44
Checking configuration for instant recovery.....	44
Setting environment variables.....	45
Using the Data Protector GUI.....	46

Using the Data Protector CLI.....	46
Switching between Oracle backup methods.....	47
Backup.....	47
Creating backup specifications.....	48
Examples of pre-exec and post-exec scripts on UNIX systems.....	56
Editing the Oracle RMAN script.....	56
Starting backup sessions.....	59
Considerations.....	59
Scheduling backup specifications.....	59
Running an interactive backup.....	60
Starting a backup using the GUI.....	60
Starting a backup using the CLI.....	61
Restore.....	61
Prerequisites.....	63
Restoring from backup media to the application system on LAN.....	63
Restoring Oracle using the Data Protector GUI.....	63
Restoring database items in a disaster recovery.....	63
Changing the database state.....	63
Restoring the recovery catalog database.....	64
Restoring the control file.....	65
Restoring Oracle database objects.....	66
Restoring tablespaces and datafiles.....	68
Duplicating an Oracle database.....	69
Restore, recovery, and duplicate options.....	71
Restore action options.....	71
General options.....	72
Duplicate options.....	73
Restore and recovery options.....	73
Restoring Oracle using RMAN.....	74
Preparing the Oracle database for restore.....	75
Connection strings used in the examples.....	76
SBT_LIBRARY parameter.....	76
Example of full database restore and recovery.....	76
Example of point-in-time restore.....	77
Example of tablespace restore and recovery.....	78
Example of datafile restore and recovery.....	80
Example of archive log restore.....	83
Example of database restore using a different device (with the automatic device selection functionality disabled).....	84
Restoring using another device.....	84
Instant recovery and database recovery.....	85
Instant recovery using the Data Protector GUI.....	85
Oracle database recovery after the instant recovery.....	87
Oracle in Veritas Cluster instant recovery.....	88
Aborting sessions.....	89
Troubleshooting.....	89
Before you begin.....	90
Checks and verifications.....	90
Problems.....	95
2 Data Protector SAP R/3 ZDB integration.....	104
Introduction.....	104
Integration concepts.....	105
ZDB flow.....	105
Data Protector SAP R/3 configuration file.....	107

Setting, retrieving, listing, and deleting Data Protector SAP R/3 configuration file parameters using the CLI.....	109
Configuring the integration.....	111
Prerequisites.....	111
Before you begin.....	112
Cluster-aware clients.....	112
Configuring user accounts.....	113
Configuring SQL*Net V2 or Net8 TNS listener.....	113
Checking the connection.....	114
Authentication password file.....	115
Enabling archived logging.....	115
Sharing directories on the application system.....	116
UNIX application system.....	116
Windows application system.....	117
Choosing authentication mode.....	117
Configuring SAP R/3 databases.....	117
Before you begin.....	118
Using the Data Protector GUI.....	118
Using the Data Protector CLI.....	120
Handling errors.....	121
Checking the configuration.....	121
Using the Data Protector GUI.....	122
Using the Data Protector CLI.....	122
Configuring the SAP R/3 parameter file.....	123
Backup.....	123
Considerations	124
Creating backup specifications.....	124
Modifying backup specifications.....	130
Scheduling backup specifications.....	130
Scheduling example.....	130
Previewing backup sessions.....	131
Using the Data Protector GUI.....	131
Using the Data Protector CLI.....	131
What happens during the preview?.....	132
Starting backup sessions.....	132
Backup methods.....	132
Using the Data Protector GUI.....	132
Using the Data Protector CLI.....	132
Using the SAP BRTOOLS.....	132
Configuring SAP compliant ZDB sessions.....	134
Using the Data Protector GUI.....	134
Using the Data Protector CLI.....	135
Manual balancing.....	135
Restore.....	136
Considerations.....	137
Standard restore.....	137
Instant recovery.....	140
Considerations.....	140
Instant recovery using the Data Protector GUI.....	140
Database recovery options.....	141
Instant recovery using the Data Protector CLI.....	142
Instant recovery from replicas containing the control file.....	143
Restoring using another device.....	144
Using the Data Protector GUI.....	144
Using the Data Protector CLI or SAP commands.....	144

Localized SAP R/3 objects.....	144
Monitoring sessions.....	144
Troubleshooting.....	144
Before you begin.....	145
General troubleshooting.....	145
Verifying the prerequisites (Oracle side).....	145
Verifying the prerequisites (SAP side).....	146
Verifying the configuration.....	147
Verifying the backup configuration.....	148
Verifying restore.....	148
Configuration and backup problems.....	150
Restore problems.....	152
3 Data Protector Microsoft SQL Server ZDB integration.....	153
Introduction.....	153
Integration concepts.....	153
Configuring the integration.....	154
Prerequisites.....	154
Before you begin.....	155
Data Protector SQL Server configuration file.....	155
Configuring users.....	156
Configuring an SQL Server cluster.....	156
Configuring SQL Server instances.....	156
Using the Data Protector GUI.....	156
Using the Data Protector CLI.....	158
Changing and checking configuration.....	158
Using the Data Protector GUI.....	158
Using the Data Protector CLI.....	159
Backup.....	159
Creating ZDB specifications.....	160
SQL Server-specific backup options.....	164
Scheduling backups.....	165
Scheduling example.....	165
Starting backup sessions.....	165
Using the Data Protector GUI.....	165
Using the Data Protector CLI.....	166
Restore.....	166
Before you begin.....	166
.....	166
Restore options.....	168
Restoring to another SQL Server instance or/and another SQL Server.....	170
Instant recovery.....	170
Monitoring sessions.....	172
Troubleshooting.....	172
Before you begin.....	172
Checks and verifications.....	173
Problems.....	173
4 Data Protector Microsoft Exchange Server 2003 ZDB integration.....	178
Introduction.....	178
Integration concepts.....	178
Configuring the integration.....	179
Prerequisites.....	179
Before you begin.....	179
Backup.....	180
Configuring Exchange Server ZDB.....	180

Creating ZDB specifications.....	181
Examples of using omnicreatedl.....	185
Modifying ZDB specifications.....	186
Checking Exchange files for consistency.....	187
Scheduling backups.....	188
Scheduling example.....	188
Starting backup sessions.....	188
Using the Data Protector GUI.....	188
Using the Data Protector CLI.....	189
Restore.....	189
Standard restore.....	189
Point-in-time recovery.....	190
Rollforward recovery.....	192
Instant recovery.....	195
Point-in-time recovery.....	195
Rollforward recovery.....	195
Troubleshooting.....	197
Before you begin.....	197
Checks and verifications.....	197
Problems.....	198
5 Data Protector Microsoft Exchange Server 2010 ZDB integration.....	200
Introduction.....	200
Integration concepts.....	201
Supported environments.....	201
Standalone environments.....	201
DAG environments.....	201
Configuring the integration.....	203
Prerequisites.....	203
Before you begin.....	204
Configuring user accounts.....	204
Backup.....	205
Backup types.....	206
Microsoft Exchange Server backup types.....	206
ZDB backup types.....	206
VSS backup types.....	206
Backup parallelism.....	207
Replica rotation in DAG environments.....	207
Backup considerations.....	208
Creating backup specifications.....	209
Modifying backup specifications.....	215
Scheduling backup specifications.....	215
Scheduling example.....	215
Previewing backup sessions.....	216
Using the Data Protector GUI.....	216
Using the Data Protector CLI.....	216
What happens during the preview?.....	216
Starting backup sessions.....	216
Using the Data Protector GUI.....	216
Using the Data Protector CLI.....	216
Backup objects.....	217
Restore.....	218
Restore methods.....	218
Repair all passive copies with failed status.....	218
Restore to the latest state.....	218

Restore to a point in time.....	219
Restore to a new mailbox database.....	219
Restore files to a temporary location.....	219
Restore destination.....	220
Restoring to a standalone database.....	220
Restoring to an active copy.....	220
Restoring to a passive copy.....	220
Restoring data to a new database.....	220
Restoring data to a temporary location.....	221
Instant recovery in DAG environments.....	221
Restore chain.....	221
Restore chain during instant recovery.....	222
Restore parallelism.....	222
Finding information for restore.....	222
Using the Data Protector GUI.....	222
Using the Data Protector CLI.....	223
Standard restore.....	223
Restoring using the Data Protector GUI.....	223
Restoring using the Data Protector CLI.....	229
Restoring using another device.....	232
Instant recovery.....	232
Performing instant recovery using the Data Protector GUI.....	232
Performing instant recovery using the Data Protector CLI.....	237
Restore options.....	239
Monitoring sessions.....	242
Troubleshooting.....	243
Before you begin.....	243
Checks and verifications.....	243
Problems.....	243
6 Data Protector Microsoft SharePoint Server 2007/2010 VSS based solution...	246
Introduction.....	246
Backup.....	246
Limitations.....	247
Restore.....	247
Installation and configuration.....	247
ZDB prerequisites.....	247
Microsoft Office SharePoint Server 2007.....	247
Microsoft SharePoint Server 2010.....	247
Licensing.....	248
Installing the integration.....	248
Configuring the integration.....	249
Configuring user accounts.....	249
Backup.....	249
How the command works.....	250
Microsoft Office SharePoint Server 2007.....	250
Microsoft SharePoint Server 2010.....	252
Considerations.....	252
The command syntax.....	253
Option description.....	253
Starting Windows PowerShell.....	255
Creating backup specifications (examples).....	256
Modifying backup specifications.....	257
Source page.....	257
Destination page.....	257

Options page.....	257
Starting backup sessions (examples).....	258
Scheduling backup sessions.....	261
Restore.....	261
Before you begin.....	262
Restoring data.....	262
Considerations.....	262
Prerequisites.....	263
Restoring using the Data Protector GUI.....	263
Restoring using the Data Protector CLI.....	265
Limitations.....	265
After the restore.....	265
Restoring index files on the Query system.....	266
Troubleshooting.....	266
Before you begin.....	266
Checks and verifications.....	266
After restore, you cannot connect to the Central Administration web page.....	267
Backup fails with the error Failed to resume Service Windows SharePoint Services Help Search...	267
After restore, a quiesce operation fails.....	267
After restore, you cannot connect to the FAST Search Server.....	268
The SharePoint_VSS_backup.ps1 script stops responding and the farm stays in read only mode...	268
A Appendix.....	269
In this appendix.....	269
Reconfiguring an Oracle instance for instant recovery.....	269
Examples for moving the control files and redo logs to different locations	270
ZDB integrations omnirc variables.....	271
Glossary.....	276
Index.....	306

Publication history

Guide updates may be issued between editions to correct errors or document product changes. To ensure that you receive updated or new editions, subscribe to the appropriate product support service. See your HP sales representative for details.

Table 1 Edition history

Part number	Guide edition	Product
B6960-90114	October 2004	Data Protector Release A.05.50
B6960-96013	July 2006	Data Protector Release A.06.00
B6960-96047	November 2008	Data Protector Release A.06.10
B6960-90163	September 2009	Data Protector Release A.06.11
N/A	March 2011	Data Protector Release 6.20
N/A	December 2011	Data Protector Release 6.20 with any of the following patches: DPWIN_00551, PHSS_42652, DPSOL_00477, DPLNX_00183
N/A	December 2011 (third edition)	Data Protector Release 6.20 with any of the following patches: DPWIN_00551, PHSS_42652, DPSOL_00477, DPLNX_00183

About this guide

This guide describes how to configure and use Data Protector disk array integrations with other software products.

Intended audience

This guide is intended for backup administrators responsible for planning, setting up, and maintaining network backups. It assumes you are familiar with:

- Basic Data Protector functionality
- Database administration

Conceptual information can be found in the *HP Data Protector Concepts Guide*, which is recommended to fully understand the fundamentals and the model of Data Protector.

It is also recommended to read the *HP Data Protector Zero Downtime Backup Concepts Guide* for fundamentals of Data Protector integrations with disk arrays.

Documentation set

Other documents and online Help provide related information.

Guides

Data Protector guides are available in the electronic PDF format. Install the PDF files during the Data Protector setup procedure by selecting the *English Documentation (Guides, Help)* component on Windows or the *OB2-DOCS* component on UNIX. Once installed, the guides reside in the *Data_Protector_home\docs* directory on Windows and in the */opt/omni/doc/C* directory on UNIX.

You can find these documents from the *Manuals* page of the HP Information Management Digital Hub website:

<http://www.hp.com/go/imhub>

In the *Storage* section, click **Storage Software** and then select your product.

- *HP Data Protector Concepts Guide*
This guide describes Data Protector concepts and provides background information on how Data Protector works. It is intended to be used with the task-oriented online Help.
- *HP Data Protector Installation and Licensing Guide*
This guide describes how to install the Data Protector software, taking into account the operating system and architecture of your environment. This guide also gives details on how to upgrade Data Protector, as well as how to obtain the proper licenses for your environment.
- *HP Data Protector Troubleshooting Guide*
This guide describes how to troubleshoot problems you may encounter when using Data Protector.
- *HP Data Protector Disaster Recovery Guide*
This guide describes how to plan, prepare for, test, and perform a disaster recovery.

- *HP Data Protector Integration Guides*

These guides describe how to configure and use Data Protector to back up and restore various databases and applications. They are intended for backup administrators or operators. There are six guides:

- *HP Data Protector Integration Guide for Microsoft Applications: SQL Server, SharePoint Server, and Exchange Server*

This guide describes the integrations of Data Protector with the following Microsoft applications: Microsoft SQL Server, Microsoft SharePoint Server, and Microsoft Exchange Server.

- *HP Data Protector Integration Guide for Oracle and SAP*

This guide describes the integrations of Data Protector with Oracle Server, SAP R/3, and SAP MaxDB.

- *HP Data Protector Integration Guide for IBM Applications: Informix, DB2, and Lotus Notes/Domino*

This guide describes the integrations of Data Protector with the following IBM applications: Informix Server, IBM DB2 UDB, and Lotus Notes/Domino Server.

- *HP Data Protector Integration Guide for Sybase, Network Node Manager, and Network Data Management Protocol Server*

This guide describes the integrations of Data Protector with Sybase Server, HP Network Node Manager, and Network Data Management Protocol Server.

- *HP Data Protector Integration Guide for Virtualization Environments*

This guide describes the integrations of Data Protector with virtualization environments: VMware Virtual Infrastructure and VMware vSphere, Microsoft Hyper-V, and Citrix XenServer.

- *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*

This guide describes the integration of Data Protector with the Microsoft Volume Shadow Copy Service. This guide also documents application writer specifics.

- *HP Data Protector Integration Guide for HP Operations Manager for UNIX*

This guide describes how to monitor and manage the health and performance of the Data Protector environment with HP Operations Manager and HP Service Navigator on UNIX.

- *HP Data Protector Integration Guide for HP Operations Manager for Windows*

This guide describes how to monitor and manage the health and performance of the Data Protector environment with HP Operations Manager on Windows.

- *HP Data Protector Zero Downtime Backup Concepts Guide*

This guide describes Data Protector zero downtime backup and instant recovery concepts and provides background information on how Data Protector works in a zero downtime backup environment. It is intended to be used with the task-oriented *HP Data Protector Zero Downtime Backup Administrator's Guide* and the *HP Data Protector Zero Downtime Backup Integration Guide*.

- *HP Data Protector Zero Downtime Backup Administrator's Guide*

This guide describes how to configure and use the integration of Data Protector with HP P6000 EVA Disk Array Family, HP P9000 XP Disk Array Family, HP P4000 SAN Solutions, and EMC Symmetrix Remote Data Facility and TimeFinder. It is intended for backup administrators or operators. It covers the zero downtime backup, instant recovery, and the restore of filesystems and disk images.

- *HP Data Protector Zero Downtime Backup Integration Guide*
This guide describes how to configure and use Data Protector to perform zero downtime backup, instant recovery, and standard restore of Oracle Server, SAP R/3, Microsoft Exchange Server, and Microsoft SQL Server databases.
- *HP Data Protector Granular Recovery Extension User Guide for Microsoft SharePoint Server*
This guide describes how to configure and use the Data Protector Granular Recovery Extension for Microsoft SharePoint Server. The Data Protector Granular Recovery Extension is integrated into Microsoft SharePoint Server Central Administration and enables you to recover individual items. This guide is intended for Microsoft SharePoint Server administrators and Data Protector backup administrators.
- *HP Data Protector Granular Recovery Extension User Guide for VMware vSphere*
This guide describes how to configure and use the Data Protector Granular Recovery Extension for VMware vSphere. The Data Protector Granular Recovery Extension is integrated into VMware vCenter Server and enables you to recover individual items. This guide is intended for VMware vCenter Server users and Data Protector backup administrators.
- *HP Data Protector Media Operations User Guide*
This guide provides information for network administrators responsible for maintaining and backing up systems on the tracking and management of offline storage media. It describes the tasks of installing and configuring the application, performing daily media operations and producing reports.
- *HP Data Protector Product Announcements, Software Notes, and References*
This guide gives a description of new features of HP Data Protector 6.20. It also provides information on installation requirements, required patches, and limitations, as well as known issues and workarounds.
- *HP Data Protector Product Announcements, Software Notes, and References for Integrations to HP Operations Manager*
This guide fulfills a similar function for the HP Operations Manager integration.
- *HP Data Protector Media Operations Product Announcements, Software Notes, and References*
This guide fulfills a similar function for Media Operations.
- *HP Data Protector Command Line Interface Reference*
This guide describes the Data Protector command-line interface, command options and their usage as well as providing some basic command-line examples.

Online Help

Data Protector provides Help topics and context-sensitive (F1) Help for Windows and UNIX platforms. You can access the online Help from the top-level directory of any installation DVD-ROM without installing Data Protector:

- **Windows:** Open `DP_help.chm`.
- **UNIX:** Unpack the zipped tar file `DP_help.tar.gz`, and access the online Help system through `DP_help.htm`.

Documentation map

Abbreviations

Abbreviations in the documentation map that follows are explained below. The guide titles are all preceded by the words “HP Data Protector”.

Abbreviation	Guide
CLI	Command Line Interface Reference
Concepts	Concepts Guide
DR	Disaster Recovery Guide
GS	Getting Started Guide
GRE-SPS	Granular Recovery Extension User Guide for Microsoft SharePoint Server
GRE-VMware	Granular Recovery Extension User Guide for VMware vSphere
Help	Online Help
IG-IBM	Integration Guide for IBM Applications: Informix, DB2, and Lotus Notes/Domino
IG-MS	Integration Guide for Microsoft Applications: SQL Server, SharePoint Server, and Exchange Server
IG-O/S	Integration Guide for Oracle and SAP
IG-OMU	Integration Guide for HP Operations Manager for UNIX
IG-OMW	Integration Guide for HP Operations Manager for Windows
IG-Var	Integration Guide for Sybase, Network Node Manager, and Network Data Management Protocol Server
IG-VirtEnv	Integration Guide for Virtualization Environments
IG-VSS	Integration Guide for Microsoft Volume Shadow Copy Service
Install	Installation and Licensing Guide
MO GS	Media Operations Getting Started Guide
MO RN	Media Operations Product Announcements, Software Notes, and References
MO UG	Media Operations User Guide
PA	Product Announcements, Software Notes, and References
Trouble	Troubleshooting Guide
ZDB Admin	ZDB Administrator's Guide
ZDB Concept	ZDB Concepts Guide
ZDB IG	ZDB Integration Guide

Map

The following table shows where to find information of different kinds. Shaded squares are a good place to look first.

								Integration Guides							ZDB			GRE		MO				
	Help	GS	Concepts	Install	Trouble	DR	PA	MS	O/S	IBM	Var	VSS	VirtEnv	OMU	OMW	Concept	Admin	IG	SPS	VMware	GS	User	PA	CLI
Backup	X	X	X					X	X	X	X	X	X			X	X	X						
CLI																								X
Concepts/ techniques	X		X					X	X	X	X	X	X	X	X	X	X	X	X	X				
Disaster recovery	X		X			X																		
Installation/ upgrade	X	X		X			X							X	X						X	X		
Instant recovery	X		X													X	X	X						
Licensing	X			X			X															X		
Limitations	X				X		X	X	X	X	X	X	X					X					X	
New features	X						X																X	
Planning strategy	X		X													X								
Procedures/ tasks	X			X	X	X		X	X	X	X	X	X	X	X		X	X	X	X		X		
Recommendations			X				X									X							X	
Requirements				X			X	X	X	X	X	X	X	X	X						X	X	X	
Restore	X	X	X					X	X	X	X	X	X				X	X	X	X				
Supported configurations																X								
Troubleshooting	X			X	X			X	X	X	X	X	X	X	X		X	X	X	X				

Integrations

Look in these guides for details of the integrations with the following software applications:

Software application	Guides
HP Network Node Manager (NNM)	IG-Var
HP Operations Manager	IG-OMU, IG-OMW
IBM DB2 UDB	IG-IBM
Informix Server	IG-IBM
Lotus Notes/Domino Server	IG-IBM
Media Operations	MO User
Microsoft Exchange Server	IG-MS, ZDB IG
Microsoft Hyper-V	IG-VirtEnv
Microsoft SharePoint Server	IG-MS, ZDB IG, GRE-SPS
Microsoft SQL Server	IG-MS, ZDB IG
Microsoft Volume Shadow Copy Service (VSS)	IG-VSS
Network Data Management Protocol (NDMP) Server	IG-Var
Oracle Server	IG-O/S, ZDB IG
SAP MaxDB	IG-O/S
SAP R/3	IG-O/S, ZDB IG

Software application	Guides
Sybase Server	IG-Var
VMware vSphere	IG-VirtEnv, GRE-VMware

Look in these guides for details of the integrations with the following families of disk array systems:


Disk array family	Guides
EMC Symmetrix	all ZDB
HP P4000 SAN Solutions	ZDB Concept, ZDB Admin, IG-VSS
HP P6000 EVA Disk Array Family	all ZDB, IG-VSS
HP P9000 XP Disk Array Family	all ZDB, IG-VSS

Document conventions and symbols

Table 2 Document conventions

Convention	Element
Blue text: "Document conventions" (page 16)	Cross-reference links and e-mail addresses
Blue, underlined text: http://www.hp.com	Website addresses
Bold text	<ul style="list-style-type: none"> Keys that are pressed Text typed into a GUI element, such as a box GUI elements that are clicked or selected, such as menu and list items, buttons, tabs, and check boxes
<i>Italic</i> text	Text emphasis
Monospace text	<ul style="list-style-type: none"> File and directory names System output Code Commands, their arguments, and argument values
<i>Monospace, italic</i> text	<ul style="list-style-type: none"> Code variables Command variables
Monospace, bold text	Emphasized monospace text

 **CAUTION:** Indicates that failure to follow directions could result in damage to equipment or data.

 **IMPORTANT:** Provides clarifying information or specific instructions.


NOTE: Provides additional information.

 **TIP:** Provides helpful hints and shortcuts.

Data Protector graphical user interface

Data Protector provides a cross-platform (Windows and UNIX) graphical user interface. You can use the original Data Protector GUI (Windows only) or the Data Protector Java GUI. For information about the Data Protector graphical user interface, see the online Help.

Figure 1 Data Protector graphical user interface



General information

General information about Data Protector can be found at <http://www.hp.com/go/dataprotector>.

HP technical support

For worldwide technical support information, see the HP support website:

<http://www.hp.com/support>

Before contacting HP, collect the following information:

- Product model names and numbers
- Technical support registration number (if applicable)
- Product serial numbers
- Error messages
- Operating system type and revision level
- Detailed questions

Subscription service

HP recommends that you register your product at the Subscriber's Choice for Business website:

<http://www.hp.com/go/e-updates>

After registering, you will receive e-mail notification of product enhancements, new driver versions, firmware updates, and other product resources.

HP websites

For additional information, see the following HP websites:

- <http://www.hp.com>
- <http://www.hp.com/go/software>
- <http://www.hp.com/go/imhub>
- <http://support.openview.hp.com/selfsolve/manuals>
- <http://www.hp.com/support/downloads>

Documentation feedback

HP welcomes your feedback.

To make comments and suggestions about product documentation, please send a message to DP.DocFeedback@hp.com. All submissions become the property of HP.

1 Data Protector Oracle Server ZDB integration

Introduction

You can employ a variety of backup strategies to best meet your system priorities. If database availability is the highest priority, for instance, your backup strategy should include online backups that are performed frequently to minimize recovery time. This strategy limits downtime, but uses system resources more intensively. The Data Protector zero downtime backup (ZDB) functionality offers online backup capabilities with minimal degradation of the application system performance.

Supported disk arrays

The following disk arrays can be used for zero downtime backup (ZDB) of the Oracle Server data:

- HP P6000 EVA Disk Array Family (P6000 EVA Array)
- HP P9000 XP Disk Array Family (P9000 XP Array)
- EMC Symmetrix (EMC)

NOTE: With the Data Protector EMC integration, instant recovery is not supported, and ZDB to tape is the only supported ZDB form.

In Oracle Server configurations that use Automatic Storage Management (ASM), zero downtime backup is only supported with the Data Protector P6000 EVA Array and Data Protector P9000 XP Array integrations, and instant recovery is only supported with the Data Protector P9000 XP Array integration.

Advantages

The advantages of using Data Protector Oracle ZDB integration are:

- ZDB reduces the performance degradation of the application system.
- The tablespaces are in backup mode (online backup) or the database is shut down (offline backup) only during the short period required to create a **replica** (split the mirror disks or create snapshots).
- The load to the application system is significantly reduced. Following the replica creation, tape backup can be started on the copied data, at leisure, using a separate backup system.

The Data Protector Oracle ZDB integration offers online and offline backup of your Oracle Server System (application system).

The online backup concept is widely used since it enables high application availability. Offline backup requires shutting down the database while creating a replica, and therefore does not offer high availability.

ZDB methods and Oracle versions

The installation, upgrade, configuration, and parts of backup flow are different depending on the selected Oracle ZDB method. These differences are indicated where appropriate.

The procedures for configuration of backup specifications and starting or scheduling backups are the same, regardless of the Oracle ZDB method.

Backup and restore types

Backup

Using Data Protector, you can perform the following types of backup:

- Online ZDB to disk, ZDB to tape, and ZDB to disk+tape.
During the creation of a replica, the database on the application system is in hot backup mode. If a ZDB-to-tape or a ZDB-to-disk+tape session is being performed, the streaming of the data to tape media is subsequently performed on the backup system.
- Offline ZDB to disk, ZDB to tape, and ZDB to disk+tape.
During the creation of a replica, the database is shut down on the application system. Therefore, the database is not available during the short time that it takes to create the replica. If a ZDB-to-tape or a ZDB-to-disk+tape session is being performed, the streaming of the data to tape media is subsequently performed on the backup system.

With both online and offline ZDB to tape or ZDB to disk+tape, a standard Data Protector (non-ZDB) backup of the recovery catalog and the control file is started automatically, after the target database backup is finished on the backup system. However, you can disable this when creating a backup specification.

NOTE: Backup of the recovery catalog and control file is not performed with ZDB to disk. The Oracle Recovery Manager utility (RMAN) is not aware of ZDB-to-disk sessions.

- ❗ **IMPORTANT:** Backup of archived logs cannot be done with the Data Protector Oracle ZDB integration. Backup of archive logs has to be done following the standard Data Protector Oracle integration backup procedure. For more information on Oracle archive log backup with Data Protector see the *HP Data Protector Integration Guide*.
-

NOTE: On EMC, decision support, application testing, and similar tasks are possible only if the Oracle binaries are installed on the backup system as well. In most cases, however, the Data Protector EMC integration requirement is that application binaries are installed on the application system only.

Restore

Using Data Protector and the disk array integrations, you can perform the following types of restore:

- Restoring from backup media to the application system on LAN (standard Data Protector restore) and using RMAN on the application system, you can:
 - recover a whole database
 - recover a part of a database
 - recover a whole database as it was at a specific point in time
- Using the instant recovery functionality and RMAN on the application system, you can:
 - perform a full database restore and database recovery
 - perform recovery from incremental backup (for ZDB to tape or ZDB to disk+tape)
 - perform recovery from a chain of incremental backups (for ZDB to tape or ZDB to disk+tape)
 - restore a datafile to a location other than its original one

“Oracle recovery methods” (page 21) provides an overview of recovery methods, depending on the type of backup that was performed and type of recovery required.

Table 3 Oracle recovery methods

Disk array	Backup types	Recover the whole database until		Recover a part of database until now
		Now	A point in time, logseq/thread, or SCN number	
P9000 XP, P6000 EVA, EMC	ZDB to tape - online	Restore	Restore	Restore
	ZDB to tape - offline	Restore	Restore ¹	Restore
P9000 XP, P6000 EVA	ZDB to disk - online	Instant recovery+ database recovery	Instant recovery+ database recovery	N/A
	ZDB to disk - offline	Instant recovery	Instant recovery+ database recovery ¹	N/A
	ZDB to disk+tape - online	<ul style="list-style-type: none"> • Restore or • Instant recovery+ database recovery 	<ul style="list-style-type: none"> • Restore or • Instant recovery+ database recovery 	Restore
	ZDB to disk+tape - offline	<ul style="list-style-type: none"> • Restore or • Instant recovery 	<ul style="list-style-type: none"> • Restore or • Instant recovery+ database recovery¹ 	Restore

¹ The database must be put in archive mode

Legend

Restore

Use the Data Protector GUI or RMAN scripts to restore the database from backup media to the application system on LAN.

Instant recovery + database recovery

The following three options are possible:

- Perform instant recovery followed by database recovery from the Data Protector Instant Recovery GUI context or
- Perform instant recovery first and then perform database recovery from the Data Protector Restore GUI context or
- Perform instant recovery first and then use RMAN scripts to recover the database.

Instant recovery

Perform instant recovery without database recovery.

See the *HP Data Protector Zero Downtime Backup Concepts Guide* for an overview of ZDB concepts and terminology.

Integration concepts

The Data Protector Oracle integration links the Oracle database management software with Data Protector. From the Oracle point of view, Data Protector represents a media management software. On the other hand, the Oracle database management system can be seen as a data source for backup, using media controlled by Data Protector.

Components

The software components involved in backup and restore processes are:

- The Oracle Recovery Manager (RMAN)
- The Data Protector Oracle integration software

Integration functionality overview

The Data Protector Oracle Integration agent (`ob2rman.pl`) works with RMAN to manage all aspects of the following operations on the Oracle target database:

- Database startup and shutdown
- Backups (backup and copy)
- Recovery (restore, recovery, and duplication)

How does the integration work?

`ob2rman.pl` executes RMAN, which directs the Oracle server processes on the target database to perform backup, restore and recovery. RMAN maintains the required information about the target databases in the recovery catalog, the Oracle central repository of information, and in the control file of a particular target database.

The main information which `ob2rman.pl` provides to RMAN is:

- Number of allocated RMAN channels
- RMAN channel environment parameters
- Information on the database objects to be backed up or restored

For backup, `ob2rman.pl` uses the Oracle target database views to get information on which logical (tablespaces) and physical (datafiles) target database objects are available for backup.

For restore, `ob2rman.pl` uses current control file or recovery catalog (if used) to get information on which objects are available for restore.

Using the Data Protector integration with RMAN, you can back up and restore the Oracle control files, datafiles, and Archived Redo Logs.

The interface from the Oracle server processes to Data Protector is provided by the Data Protector Oracle integration Media Management Library (**MML**), which is a set of routines that allows the reading and writing of data to General Media Agents.

Besides handling direct interaction with the media devices, Data Protector provides scheduling, media management, network backups, monitoring, and interactive backup.

A backup that includes all datafiles and current control file that belong to an Oracle Server instance is known as a whole database backup.

These features can be used for online or offline backup of the Oracle target database. However, you must ensure that the backup objects (such as tablespaces) are switched into the appropriate state before and after a backup session. For online backup, the database instance must operate in the ARCHIVELOG mode; whereas for offline backup, objects need to be prepared for backup using the `Pre-exec` and `Post-exec` options in the backup specification.

The Data Protector backup specification contains information about backup options, commands for RMAN, Pre- and Post-exec commands, media, and devices.

The Data Protector backup specification allows you to configure a backup and then use the same specification several times. Furthermore, scheduled backups can only be performed using a backup specification.

Backup and restore of an Oracle target database can be performed using the Data Protector User Interface or the RMAN utility.

The heart of the Data Protector Oracle integration is MML, which enables an Oracle server process to issue commands to Data Protector for backing up or restoring parts or all of the Oracle target database files. The main purpose is to control direct interaction with media and devices.

Non-ZDB flow

A Data Protector scheduled or interactive backup is triggered by the Data Protector Backup Session Manager, which reads the backup specification and starts the `ob2rman.pl` command on the

Oracle Server under the operating system user account specified in the backup specification. Further on, `ob2rman.pl` prepares the environment to start the backup, and issues the RMAN backup command. RMAN instructs the Oracle Server processes to perform the specified command.

The Oracle Server processes initialize the backup through MML, which establishes a connection to the Data Protector Backup Session Manager. The Backup Session Manager starts the General Media Agent, sets up a connection between MML and the General Media Agent, and then monitors the backup process.

The Oracle Server processes read the data from the disks and send it to the backup devices through MML and the General Media Agent.

RMAN writes information regarding the backup either to the recovery catalog (if one is used) or to the control file of the Oracle target database.

Messages from the backup session are sent to the Backup Session Manager, which writes messages and information regarding the backup session to the IDB.

The Data Protector General Media Agent writes data to the backup devices.

Restore flow

A restore session can be started using:

- Data Protector GUI
- RMAN CLI

You must specify which objects are to be restored.

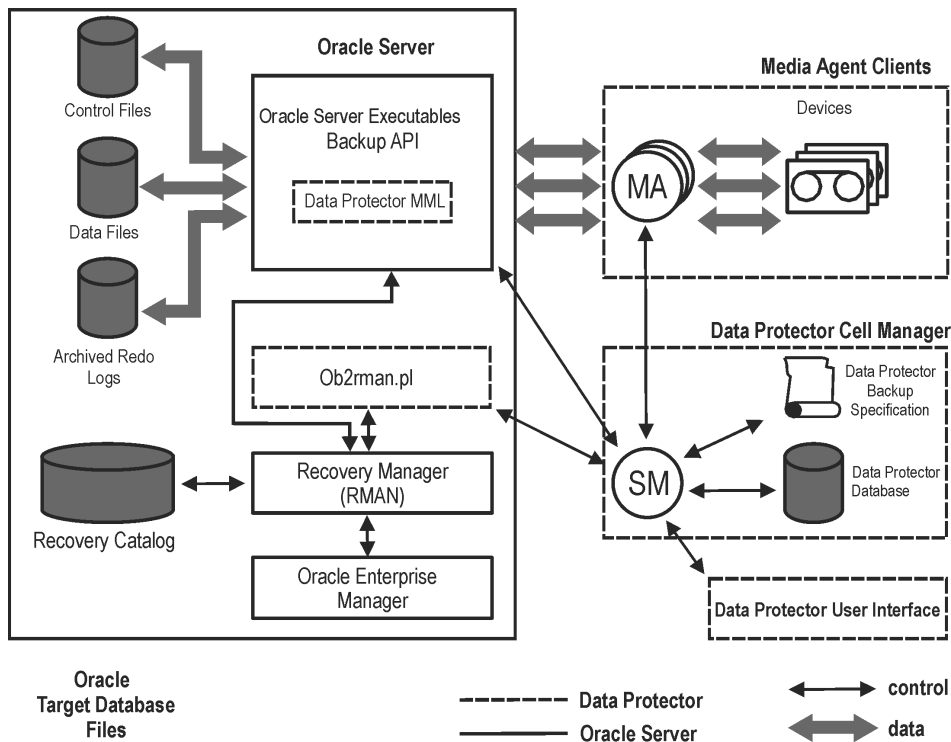
A restore from the Data Protector user interface is triggered by the Data Protector Restore Session Manager, which starts the `ob2rman.pl` command. `Ob2rman.pl` prepares the environment to start the restore, and issues the RMAN restore command. RMAN checks the recovery catalog (if one is used) or the control file to gather the information about the Oracle backup objects. It also contacts the Oracle Server processes, which initialize the restore through MML. MML establishes a connection with the Restore Session Manager and passes along the information about which objects and object versions are needed.

The Restore Session Manager checks the IDB to find the appropriate devices and media, starts the General Media Agent, establishes a connection between MML and the General Media Agent, and then monitors the restore and writes messages and information regarding the restore to the IDB.

The General Media Agent reads the data from the backup devices and sends it to the Oracle Server processes through MML. The Oracle Server Processes write the data to the disks.

The concept of Oracle integration, data and the control flow are shown in [“Data Protector Oracle integration concept” \(page 24\)](#), and the related terms are explained in the following table.

Figure 2 Data Protector Oracle integration concept



Database files can also be managed by **Automatic Storage Management (ASM)**.

Legend

SM	The Data Protector Session Manager, which can be the Data Protector Backup Session Manager during a backup session and the Data Protector Restore Session Manager during a restore session.
RMAN	The Oracle Recovery Manager.
Data Protector MML	The Data Protector Oracle integration Media Management Library, which is a set of routines that enables data transfer between the Oracle Server and Data Protector.
Backup API	The Oracle-defined application programming interface.
IDB	The IDB where all the information about Data Protector sessions, including session messages, objects, data, used devices, and media is written.
MA	The Data Protector General Media Agent, which reads and writes data from and to media devices.

Oracle backup set ZDB concepts

See the *HP Data Protector Zero Downtime Backup Concepts Guide* for a general description of ZDB-to-disk, ZDB-to-tape, and ZDB-to-disk+tape and instant recovery concepts.

With the Oracle backup set ZDB method, the entire data to be backed up is provided to Data Protector through the Oracle API—the data is streamed through the Data Protector Oracle integration MML.

Depending on the location of the Oracle control file, online redo log files, and SPFILE, the following two options are possible:

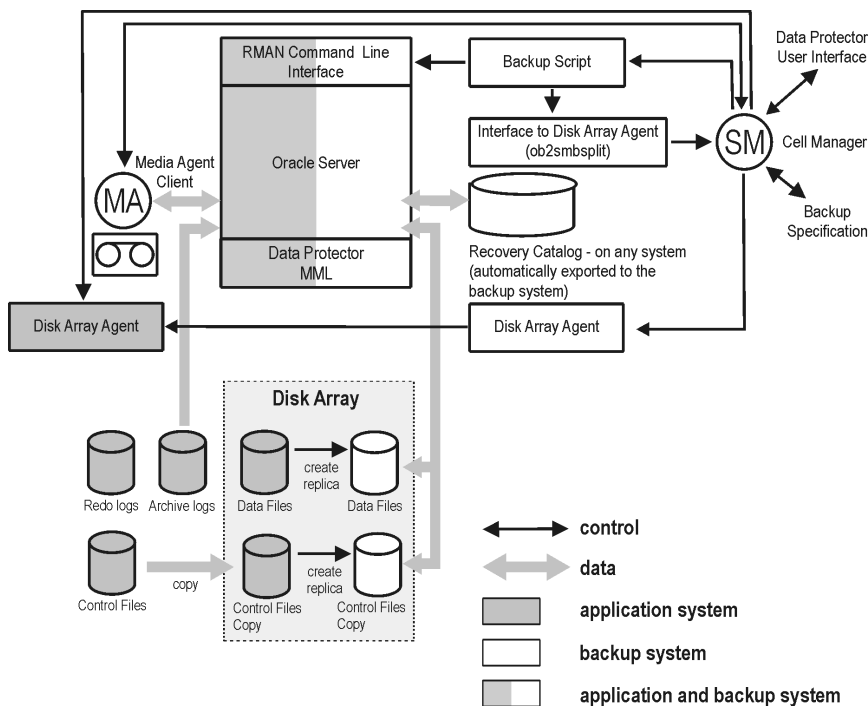
- Oracle control file, online redo log files, and SPFILE reside on a **different** volume group (if LVM is used) or source volume than Oracle datafiles.
By default, instant recovery for such a configuration is enabled.
- Oracle control file, online redo log files, SPFILE reside on the **same** volume group (if LVM is used) or source volume as Oracle datafiles.

By default, instant recovery for such a configuration is *not* enabled. You can enable instant recovery by setting the ZDB_ORa_INCLUDE_CF_OLF, ZDB_ORa_INCLUDE_SPF, and ZDB_ORa_NO_CHECKCONF_IR omnirc variables to 1. See [“ZDB integrations omnirc variables”](#) (page 271).

❗ **IMPORTANT:** If you enable instant recovery by setting the above mentioned variables, note that the control file, SPFILE, and online redo logs are overwritten during instant recovery.

The Oracle archived redo log files do not have to reside on source volumes.

Figure 3 Oracle backup set ZDB concept



“Oracle backup set ZDB concept” (page 25) presents only the default integration behavior, where Oracle control file, online redo log files, and SPFILE reside on a different volume group (if LVM is used) or source volume than Oracle datafiles. Oracle database files can also be managed by ASM, however specific limitations apply to Oracle ASM configurations. For details, see [“Limitations”](#) (page 32).

For more information on an alternative Oracle backup and restore concept, see [“ZDB integrations omnirc variables”](#) (page 271).


Legend

- MA The General Media Agent writes data from a replica to backup media. The General Media Agent typically resides on the backup system.
- SM The session manager controls backup and restore sessions and writes session information to the IDB.

<i>Disk Array Agent</i>	The disk array agents (ZDB agents) are SYMA (on EMC), SSEA (on P9000 XP Array), and SMISA (on P6000 EVA Array).
<i>Data Protector MML</i>	The Data Protector Oracle integration Media Management Library, which is a set of routines that enables data transfer between the Oracle Server and Data Protector. This is a Data Protector software library that is linked to the Oracle software.

Backup process

Figure 4 Oracle backup set ZDB flow



NOTE: ZDB agents are SYMA on EMC, SSEA on P9000 XP Array, and SMISA on P6000 EVA Array.

See the *HP Data Protector Zero Downtime Backup Concepts Guide* for a general description of ZDB and instant recovery concepts.

See the *HP Data Protector Zero Downtime Backup Administrator's Guide* for a general description of the ZDB-to-disk, ZDB-to-tape, and ZDB-to-disk+tape session flows and for the explanation of actions triggered by ZDB options.

This section provides only the information relevant to the Data Protector Oracle ZDB integration.

Operations on a replica (mounting, activating volume/disk groups,...) described below are dependent on or triggered by ZDB options. See the *HP Data Protector Zero Downtime Backup Administrator's Guide* for more information on these options.

- Data Protector executes the `ob2rman.pl` command on the backup system. This command retrieves a list of files or raw disks to be backed up from the Oracle database on the application system and starts the resolving process. The list is used only to determine the source volumes to be replicated. If the location for control file copy is specified during configuration, `ob2rman.pl` makes a copy of the control file to the specified directory on the application system. This directory has to reside on a disk array source volume.
- When performing an *online* ZDB session, `ob2rman.pl` then sets the Oracle target database into backup mode by issuing the `sqlplus` command "ALTER TABLESPACE BEGIN BACKUP", starts the procedure to create a replica of the source volumes on which the database is installed; and after the replica is created, takes the database out of backup mode by issuing the `sqlplus` command "ALTER TABLESPACE END BACKUP".

When performing an *offline* ZDB session, `ob2rman.pl` shuts down the Oracle database, starts the procedure to create a replica of the source volumes on which the database is installed; and after the replica is created, starts the Oracle database.

- `Ob2rman.pl` then starts the procedure to prepare the replica on the backup system. In this step, volume/disk groups on the backup system are enabled and, unless the database is installed on raw partitions, the mount points with the Oracle database files are mounted.
- A ZDB agent then mounts the database on the backup system to the mount points with the same names (created by Data Protector) as on the application system.

NOTE: There must be nothing already mounted on the mount point concerned on the backup system, or the resolving and backup will fail.

- If a ZDB-to-disk session is being performed, at this point the remaining ZDB options are processed and details of the session are written to the ZDB database. The session then finishes. The following steps in this description are not performed, therefore RMAN is not given any information about ZDB-to-disk session.
- If a ZDB-to-tape or a ZDB-to-disk+tape session is being performed, the processing continues as follows:
 - `Ob2rman.pl` starts the Oracle backup command RMAN on the backup system, and then sends the Oracle RMAN Backup Command Script to the RMAN cmdfile (input command file).
 - RMAN contacts the Oracle database instance on the backup system, which contacts Data Protector via SBT API and initiates a backup.
 - The Oracle database instance on the backup system reads data from the replica and sends it to the Data Protector General Media Agent for writing to the backup device.
 - At the end of data transfer, the backup system is disabled (filesystems are unmounted on all platforms and volume/disk groups deactivated on UNIX systems) and links are re-established.
 - The recovery catalog and the control file are backed up automatically after the target database backup is finished on the backup system. However, you can disable this when creating a backup specification.

NOTE: A replica of the archive logs is not created; therefore, the archive logs should be backed up from the application system, following the standard Data Protector Oracle archive logs backup procedure.

Oracle proxy-copy ZDB concepts

See the *HP Data Protector Zero Downtime Backup Concepts Guide* for a general description of ZDB-to-disk, ZDB-to-tape, ZDB-to-disk+tape, and instant recovery concepts.

The Data Protector Oracle integration MML supports the Proxy Copy functionality. This enables Data Protector to perform backup using filesystem backup methods.

Depending on the location of the Oracle control file, online redo log files, and SPFILE, the following two options are possible:

- Oracle control file, online redo log files, and SPFILE reside on a **different** volume group (if LVM is used) or source volume than Oracle datafiles.
By default, instant recovery is enabled if this option is selected in the GUI.
- Oracle control file, online redo log files, and SPFILE reside on the **same** volume group (if LVM is used) or source volume as Oracle datafiles.


By default, instant recovery is *not* enabled, even if this option is selected in the GUI. You can enable instant recovery by setting the `ZDB_ORA_INCLUDE_CF_OLF`, `ZDB_ORA_INCLUDE_SPF`, and `ZDB_ORA_NO_CHECKCONF_IR` omnirc variables to 1. See [“ZDB integrations omnirc variables”](#) (page 271).

❗ **IMPORTANT:** If you enable instant recovery by setting the above mentioned variables, note that the control file, SPFILE, and online redo logs are overwritten during instant recovery.

The Oracle archived redo log files do not have to reside on source volumes.

“Oracle proxy-copy ZDB concept” (page 28) shows the architecture of the Data Protector Oracle ZDB integration. The figure illustrates the configuration, in which the backup is performed on the backup system. It presents the default integration behavior, where Oracle control file, online redo log files, and SPFILE reside on different disk array source volumes than the Oracle data files. For more information on alternative Oracle backup and restore concepts, see [“ZDB integrations omnirc variables”](#) (page 271).

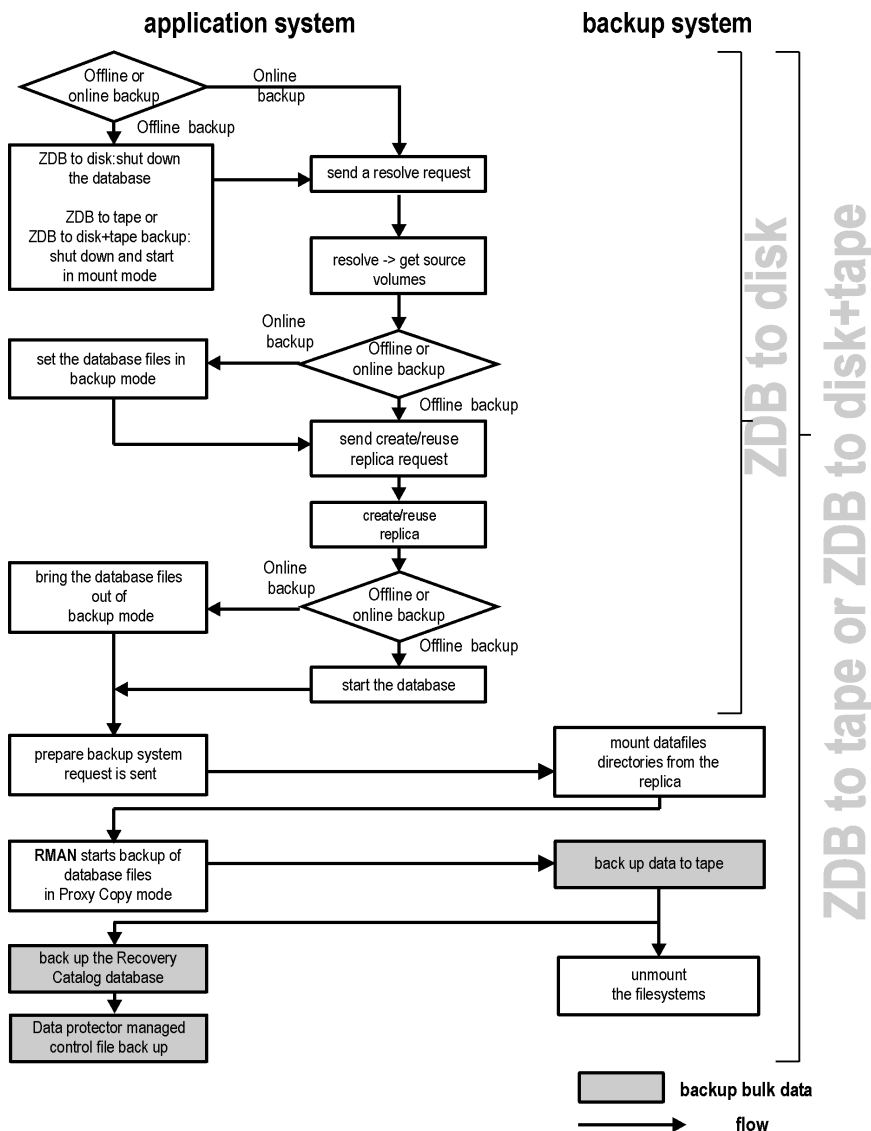
Figure 5 Oracle proxy-copy ZDB concept



MA	The General Media Agent writes data from a replica to backup media. The General Media Agent typically resides on the backup system.
SM	The session manager controls backup and restore sessions and writes the session information to the IDB.
Disk Array Agent	The disk array agents (ZDB agents) are SYMA (on EMC), SSEA (on P9000 XP Array), and SMISA (on P6000 EVA Array).
MML	The Data Protector Oracle integration Media Management Library, which is a set of routines that enables data transfer between the Oracle Server and Data Protector. This is a Data Protector software library that is linked to the Oracle software.

Backup process

Figure 6 Oracle proxy-copy ZDB flow



See the *HP Data Protector Zero Downtime Backup Administrator's Guide* for a general description of the ZDB-to-disk, ZDB-to-tape, and ZDB-to-disk+tape sessions flows and for an explanation of actions triggered by ZDB options.

This section provides only the information relevant to the Data Protector Oracle ZDB integration.

Operations on a replica (mounting, activating volume/disk groups...) described below are dependent on or triggered by ZDB options. See the *HP Data Protector Zero Downtime Backup Administrator's Guide* for more information on these options.

- In the case of an *offline* ZDB-to-disk+tape or ZDB-to-tape session, `ob2rman.pl` shuts down and opens the database instance in mount state. For both, offline and online ZDB-to-disk+tape or ZDB-to-tape sessions, Data Protector starts RMAN in proxy-copy mode.

In the case of an *offline* ZDB-to-disk session, the database is shut down.

- Data Protector retrieves a list of files or raw disks to be included in the replica creation from the Oracle database and starts the resolving process. The list is used only to determine the source volumes to be replicated.

In the case of a *ZDB-to-disk* session, if the location for control file copy is specified during configuration, Data Protector makes a copy of the control file to the specified directory on the application system. This directory has to reside on a disk array source volume.

- In the case of an *online* backup, the Oracle target database is switched into the backup mode.
- `ob2smbsplit` or MML starts the procedure to create a replica of the source volumes on which the database is installed.
- In the case of an *online* backup, the database files are taken out of the backup mode after the replica is created.

In the case of an *offline* backup, the Oracle database is started by `ob2rman.pl` after the replica is created.

- Data Protector (for ZDB to disk) or MML (for ZDB to tape or ZDB to disk+tape) starts the procedure to prepare the replica on the backup system. In this step, volume/disk groups on the backup system are enabled (UNIX systems) and, unless the database is installed on raw disks, the mount points containing the Oracle database files are mounted.
- A ZDB agent then mounts the database on the backup system to the mount points with the same names (created by Data Protector) as on the application system.
- If a ZDB-to-disk session is being performed, at this point the remaining ZDB options are processed and details of the session are written to the ZDB database. The session then finishes. The following steps in this description are not performed; therefore, RMAN is not given any information about the ZDB-to-disk session.
- MML on the application system sends a request to the Data Protector **data movement agent** (DMA) on the backup system to back up the datafiles to tape.
- The DMA reads data from the backup system and sends it to the General Media Agent to write the actual data to the backup device.

DMA's role is also to disable the General Media Agent requests from accessing the application system. Thus, the database runs on the application system with greatly reduced performance degradation since the backup is performed on the backup system.

- At the end of data transfer, the backup system is disabled (filesystems are unmounted on all platforms and volume/disk groups deactivated on UNIX systems) and links are re-established.
- The recovery catalog and the control file are backed up automatically after the target database backup is finished on the backup system. However, you can disable this when creating a backup specification.

NOTE: A replica of the archive logs is not created; therefore, the archive logs should be backed up from the application system, following the standard Data Protector Oracle archive logs backup procedure.

Configuring the integration

Prerequisites

- It is assumed that you are familiar with the Oracle database administration and the basic Data Protector functionality.
- You need a license to use the Data Protector ZDB integration with Oracle. Additional licenses are required for instant recovery and for the online extension. For information on licensing, see the *HP Data Protector Installation and Licensing Guide*.
- Before you begin, ensure that you have correctly installed and configured the Oracle Server and Data Protector client systems. See the:
 - Latest support matrices at <http://www.hp.com/support/manuals> for an up-to-date list of supported versions, platforms, devices, and other information.
 - *HP Data Protector Installation and Licensing Guide* for instructions on how to install Data Protector on various architectures and how to install a Data Protector disk array integration (EMC, P9000 XP Array, or P6000 EVA Array) with Oracle.
 - *Oracle Recovery Manager User's Guide and References* for Oracle concepts and backup/recovery strategies.
 - *Oracle Backup and Recovery Guide* for the configuration and use of Recovery Manager, as well as for Oracle backup terminology and concepts.
 - *Oracle Enterprise Manager User's Guide* for information on backup and recovery with the Oracle Enterprise Manager, as well as information about SQL*Plus.
- A Data Protector disk array integration (EMC, P9000 XP Array, or P6000 EVA Array) must be correctly installed and configured. For installation, see the *HP Data Protector Installation and Licensing Guide*. For configuration, see the *HP Data Protector Zero Downtime Backup Administrator's Guide*.
- For Oracle Server configurations in which Automatic Storage Management (ASM) is used, the disk array that will be used in zero downtime backup sessions must support creation of replicas with cross-volume data consistency:
 - If the HP P6000 EVA Disk Array Family will be used, it must support multisnapping. For configuration details, see “Enabling zero downtime backup on disk arrays of the HP P6000 EVA Disk Array Family for ASM configurations” (page 35).
 - If the HP P9000 XP Disk Array Family will be used, it must support the atomic split operation. For configuration details, see “Enabling zero downtime backup on disk arrays of the HP P9000 XP Disk Array Family for ASM configurations” (page 36).

For details about which disk array models and disk array firmware revisions support creation of replicas with cross-volume data consistency, see the latest support matrices at <http://www.hp.com/support/manuals>.

- The Oracle Server software must be installed on the application system and the Oracle target database must be open or mounted there.
- The Oracle recovery catalog database must be properly configured and open.
- Oracle net services must be properly configured and running (on the application system) for the Oracle target database and the recovery catalog. The net services are needed for the Data Protector Oracle agent to be connected to the Oracle database on the application system through Oracle.

For more information about different connection options, see the *Oracle Recovery Manager User's Guide and References*.

For details about how to check the prerequisites listed above, see [“Troubleshooting” \(page 89\)](#).

Note that the Data Protector Oracle integration uses RMAN for backup and restore. RMAN connection to a target database requires a dedicated server process. To ensure that RMAN does not connect to a dispatcher when the target database is configured for a shared server, the net service name used by RMAN must include (SERVER_DEDICATED) in the CONNECT_DATA attribute of the connection string.

- On Windows systems, if the Oracle target database and the Oracle recovery catalog are installed on two different systems, configure a *domain* user account that is a member of the Administrators group on both systems. After that, on the system with the Oracle target database installed, if the system is not Windows Server 2008, restart the Data Protector Inet service under this account.

For information on how to change the Data Protector Inet service account, see the online Help index: “changing Data Protector Inet account”.

- In case of Real Application Cluster (RAC), each node must have a dedicated disk for storing archive logs. Such disks must be NFS mounted on all other RAC nodes.
However, if the archive logs are not on a NFS mounted disk, you must modify the archive log backup specification. See [“Problem” \(page 99\)](#).
- In a RAC environment, for Oracle 11.2.0.2 or subsequent versions, the control file should be created on a shared disk and the OB2_DPMCTL_SHRLOC variable must point to this location, from this location the control file is backed up.

Limitations

- The MAXPIECESIZE RMAN parameter option is not supported because the restore of multiple backup pieces created during a backup is not possible using the Data Protector Oracle integration.
- The Oracle recovery catalog database must be used as RMAN repository for backup and restore operations. ZDB using the Oracle control file are not supported. This is set when configuring the database. See [“Configuring Oracle databases” \(page 38\)](#).
- The Oracle database identifier (DBID) must be a unique in a Data Protector cell. If you clone a database you must change the DBID.
- Preview of zero downtime backup and instant recovery sessions is not available.
- Using the Oracle proxy-copy ZDB method, individual tablespaces or datafiles cannot be backed up during a ZDB-to-disk or ZDB-to-disk+tape session (instant recovery enabled), only the whole database can be backed up.
- The Oracle backup set ZDB method is not supported on Windows.
- The Oracle backup set ZDB method is supported on UNIX raw logical volumes only if these were created with LVM or VxVM.
- When using the Oracle backup set ZDB method, you must reconfigure the Oracle integration if the initialization parameter file has been changed since the last configuration execution. See [“Configuring Oracle databases” \(page 38\)](#).
- The single-host configuration (BC1, TF/1) is not supported for Oracle backup set ZDB sessions.
- Object copying and object mirroring is not supported for ZDB to disk.
- Recovery files residing in the **flash recovery area** cannot be backed up using ZDB.

- For zero downtime backup and instant recovery sessions involving ASM-managed files, the following limitations apply:
 - Only the backup set ZDB method is supported.
 - In order to enable instant recovery of the corresponding ASM-managed files, names of the ASM disk groups involved in a zero downtime backup session must not be changed after the session.
 - With disk arrays of the HP P6000 EVA Disk Array Family, the ASM-managed files can only be backed up in ZDB-to-tape sessions involving a single P6000 EVA Array.
- **Oracle Data Guard:** Standby database is not supported for ZDB.
- The Data Protector Oracle integration does not support non-ASCII characters in backup specification names.

Before you begin

- Test whether the Oracle Server system and the Cell Manager communicate properly: Configure and run a Data Protector filesystem backup and restore on the Oracle Server system.
- Identify the Oracle database *user* that will be used by Data Protector for backup. This user must have the `SYSDBA` privilege granted. For example, it could be the Oracle user `sys`, which is created during database creation.

See the Oracle documentation for more information on user privileges in Oracle.

- In case of the backup set method, if the Oracle database is installed on symbolic links, create these symbolic links on the backup system, too.
- If an Oracle ASM instance manages files of more than one database, to enable instant recovery, you must reconfigure Oracle Server to use a separate ASM disk group for each database.
- From the application system, using SQL*Plus, connect to the target database and recovery catalog by specifying the user, password, and net connect identifier. Connect to the target database as the database administrator and to the recovery catalog database as the recovery catalog owner.

Example

If the user name for the target database is `system`, password `manager`, net service name `PROD`, and the user name and password for the recovery catalog is `rman` and the net service name `RMANCAT`, then the commands will look like:

```
sqlplus /nolog
```

```
SQL> connect system/manager@PROD as sysdba;
Connected.
```

```
SQL> connect rman/rman@RMANCAT;
Connected.
```

- For *online backup* only, enable the Oracle automatic log archiving:
 1. Shut down the Oracle target database instance on the application system.
 2. Back up the entire database using a filesystem backup.
 3. Select the location for archive logs:
 - If SPFILE is used:

Run:

```
alter system set log_archive_dest=path_to_archive_logs
SCOPE=SPFILE;
```
 - If the `init.ora` file is used:

Run:

```
log_archive_start=true
log_archive_dest=path_to_archive_logs
```

The default path of the file is:

Windows systems: *ORACLE_HOME\database\initDB_NAME.ora*

UNIX systems: *ORACLE_HOME/dbs/initDB_NAME.ora*

where *DB_NAME* is the name of the Oracle database instance.

4. Mount the target database and to enable the archive log mode, start SQL*Plus and type:

```
startup mount
alter database archivelog;
alter database open;
```

Example

If the user name for the target database is *system*, password manager, instance name *PROD*, and the user name and password for the recovery catalog is *rman*, then the commands will look like:

```
sqlplus /nolog
SQL> connect system/manager@PROD as sysdba;
Connected.
SQL> startup mount;
SQL> alter database archivelog;
Statement processed.
SQL> archive log start;
Statement processed.
SQL> alter database open;
```

5. Back up the entire database.

Backup set method

For backup set method:

- Ensure that the Oracle software on the backup system and application system have the same directory structure. That means that *ORACLE_HOME* for both Oracle installations has to be identical.
- Ensure that the following files are the same on the application system and the backup system. Check also that the permissions are identical as on the application system:
 - *names.ora*
Default path: *ORACLE_HOME/network/admin/names.ora*
 - *initDB_NAME.ora*
Default path: *ORACLE_HOME/dbs/initDB_NAME.ora*.
 - *orapwDB_NAME*
Default path: *ORACLE_HOME/dbs/orapwDB_NAME*
 - *admin/DB_NAME*
Default path: *ORACLE_BASE/admin/DB_NAME*

Ensure that the Oracle net services on the application system and the backup system have the same directory structure. This can be accomplished by either NFS sharing of the files, manually copying the files from the application system to the backup system, or by using the UNIX *rdist* or *tar* commands to distribute the files from the application system.

- Test whether the Oracle user can log in to the Oracle target database as the Oracle database administrator and to the Oracle recovery catalog database as the Oracle recovery catalog owner from the backup system:

1. Export `ORACLE_HOME`, `DB_NAME`, and on UNIX systems also `SHLIB_PATH` variables.
2. Using SQL*Plus, connect to the Oracle recovery catalog database by specifying the user (recovery catalog owner), password, and net connect identifier.
3. Connect to the Oracle target database locally using the Oracle Net software as the Oracle database administrator with the `SYSDBA` role.

Example

If the `DB_NAME` of the target database is `PROD`, the `DB_NAME` of the Oracle recovery catalog database is `RMANCAT`, and `ORACLE_HOME` is `/oracle/PROD`, then the commands will look like:

```
su - ora
id
uid=101(ora) gid=101(dba)

export DB_NAME=PROD
oracle/PROD/bin/sqlplus
SQL> connect rman/rman@RMANCAT
Connected.

SQL> connect system/manager as sysdba
SQL> connect system/manager@PROD as sysdba;
Connected.
```

- Test whether the user `root` and the Oracle administrator (for example, the user `oracle`) can connect to the target database and the recovery catalog database using the `RMAN` command on the backup system:
 1. Log in as the Oracle database administrator to the backup system (for example, the user `oracle`).
 2. Execute the `RMAN` command and connect to the target database and the recovery catalog database.

Example

If the `DB_NAME` of the target database is `PROD`, the `DB_NAME` of the Oracle recovery catalog database is `RMANCAT`, and `ORACLE_HOME` is `/oracle/PROD`, then the commands will look like:

```
su - ora
id
uid=101(ora) gid=101(dba)
export DB_NAME=PROD

rman target system/manager catalog rman/rman
Recovery Manager: Release 10.1.0.2.0 - Production
RMAN-06005: connected to target database: PROD
RMAN-06008: connected to recovery catalog database
RMAN> exit
Recovery Manager completed.
```

Enabling zero downtime backup on disk arrays of the HP P6000 EVA Disk Array Family for ASM configurations

Zero downtime backup with the HP P6000 EVA Disk Array Family is supported for ASM configurations, provided that the P6000 EVA Array supports multisnapping. Additionally, to enable zero downtime backup, the following prerequisites must be fulfilled:

- The ASM-managed files that will be backed up must reside on raw disks, not on raw logical volumes.

Note that the maximum number of source disks that can be involved in multisnapping depends on the firmware revision of the P6000 EVA Array that will be used and the installed Command View

(CV) version. If the number of source disks selected for a zero downtime backup session exceeds this limitation, the session is aborted. For limitation details, see the HP P6000 EVA Disk Array Family documentation.

Enabling zero downtime backup on disk arrays of the HP P9000 XP Disk Array Family for ASM configurations

Zero downtime backup with the oHP P9000 XP Disk Array Family is supported for ASM configurations, provided that the P9000 XP Array's atomic split configuration is enabled. Additionally, to enable zero downtime backup, the following prerequisites must be fulfilled:

- The ASM-managed files that will be backed up must reside on raw disks, not on raw logical volumes.
- The autoextend feature of the Oracle Server ASM must be disabled.
- The P9000 XP Array mirror groups must be created in the group mode with a specific CT number assigned. The CT number 0 is not supported.
- The Data Protector `omnirc` variable `SSEA_ATOMIC_SPLIT` must be enabled (set to 1). By default, the variable is disabled (set to 0).

Cluster-aware systems

In cluster environment, if you intend to use the Data Protector CLI, set the Data Protector environment variable `OB2BARHOSTNAME` to the virtual server name. Set the variable on the Oracle Server system as follows:

Windows systems: `set OB2BARHOSTNAME=virtual_server_name`

UNIX systems: `export OB2BARHOSTNAME=virtual_server_name`

HP-UX with RAC: To enable instant recovery, create an MC/ServiceGuard package containing only the virtual IP and the virtual hostname parameters and distribute it among the RAC nodes.

Linking Oracle Server with the Data Protector MML

To use the Data Protector Oracle integration, the Oracle Server software needs to be linked with the Data Protector Oracle integration **Media Management Library (MML)** on every system on which an Oracle instance is running.

You do not need to link Oracle Server with the Data Protector MML manually. When you start backups or restores using the Data Protector GUI or CLI, Data Protector automatically links Oracle Server with the correct platform-specific Data Protector MML. However, for testing purposes, you can override this automatic selection. You can manually specify which platform-specific Data Protector MML should be used by setting the Data Protector `SBT_LIBRARY` parameter. On how to set the parameter, see the `util_cmd` man page. The parameter is saved in the Data Protector Oracle instance configuration file.

MML is invoked by the Oracle server when it needs to write to or read from devices using Data Protector.

Configuring Oracle user accounts

Decide under which user accounts you want backups to run. Data Protector requires the following user accounts:

- Oracle operating system user account
For details, see [“Configuring Oracle operating system user accounts” \(page 37\)](#).
- Oracle database user accounts
For details, see [“Configuring Oracle database users accounts” \(page 38\)](#).

Configuring Oracle operating system user accounts


For each Oracle database, Data Protector requires an operating system user account that has Oracle rights to back up the database. This user account usually belongs to the DBA user group (**OSDBA user**). The user account under which the Oracle database is running has these rights. For example, to find such a user on UNIX systems, run:

```
ps -ef | grep ora_pmon_DB_NAME
```

or

```
ps -ef | grep ora_lgwr_DB_NAME
```

Figure 7 Finding the Oracle user



The following table explains how to configure users on different operating systems:

Client system	Description
UNIX system	<p>Ensure that the Oracle user <code>oracle</code> from the Oracle Inventory group (<code>oinstall</code>) has been added to the Data Protector <code>admin</code> user group. For details on adding users, see the online Help index: "adding users".</p> <p>Add the OSDBA user account and <code>root</code> user account from both the application system and backup system to the Data Protector <code>admin</code> or <code>operator</code> user group. The OSDBA user on the backup system must have the same numerical user ID and group ID as the OSDBA user on the application system (for example, <code>uid=101 (ora)</code> <code>gid=101 (dba)</code>).</p> <p>TIP: To find the user ID, connect to a system under this user account and run:</p> <pre>#id</pre>
Windows system	<p>On Windows systems, Data Protector connects to the Oracle database using the Data Protector <code>Inet</code> service on the related system. By default, the service runs under the <code>Local System</code> account, which is automatically added to the Data Protector <code>admin</code> user group. However, if you have restarted the Data Protector <code>Inet</code> service on the application system and backup system under OSDBA user accounts, you need to add the new users to the Data Protector <code>admin</code> or <code>operator</code> user group.</p>

For information on adding users to Data Protector user groups, see the online Help index: "adding users".

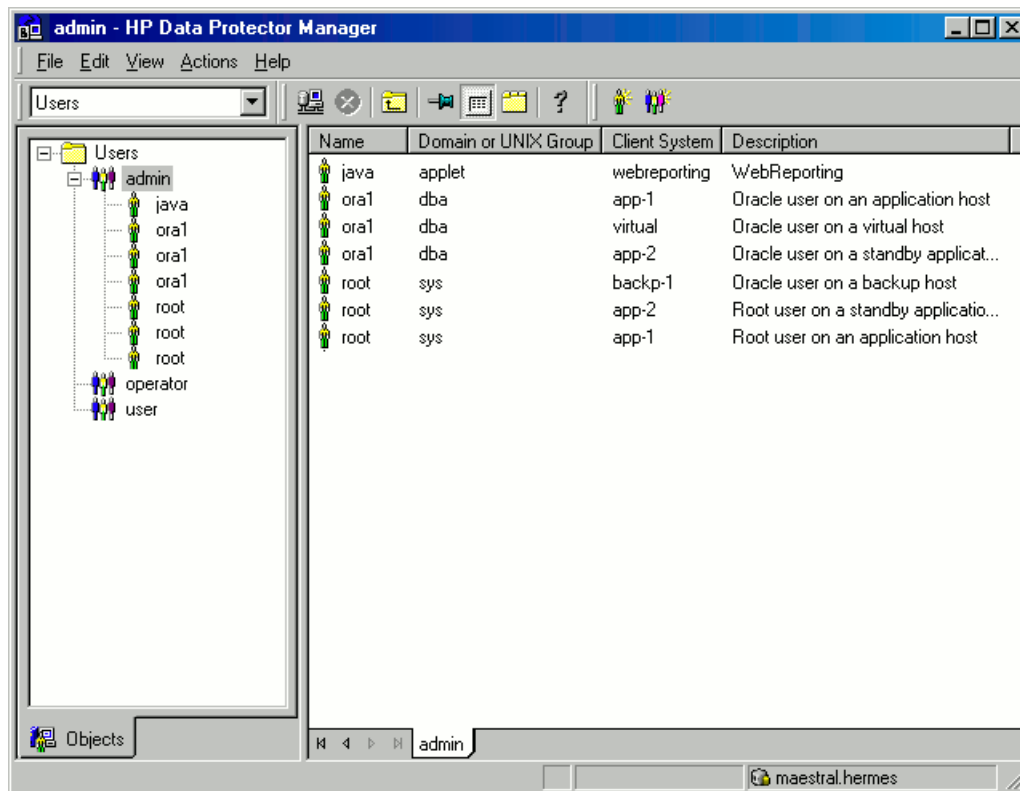
NOTE: The OSDBA user account for the backup system needs to be added to a Data Protector user group only if you plan to use the Oracle backup set ZDB method.

Clusters

In cluster environments, ensure to add the following users to the Data Protector `admin` or `operator` user group:

- OSDBA user for all physical nodes
- OSDBA user for the virtual server (applicable for MC/ServiceGuard clusters)
- **UNIX systems only:** `root` user for all physical nodes

Figure 8 Example user configuration in a cluster environment



Configuring Oracle database users accounts

Identify or create the following Oracle database user accounts. You need to provide these user accounts when you configure the Oracle database as described in [“Configuring Oracle databases”](#) (page 38).

Table 4 Oracle database user accounts

User	Description
Primary database user	Required to log in to the primary database.
Recovery catalog user	<p>The owner of the recovery catalog (for example, <code>rman</code>). Required to log in to the catalog database. Needed if you use the recovery catalog.</p> <p>If you are using Oracle 11g R2 or later, ensure that the owner of the Oracle recovery catalog:</p> <ul style="list-style-type: none">is granted the <code>CREATE ANY DIRECTORY</code> and the <code>DROP ANY DIRECTORY</code> system privileges, which are required to use the Data Pump Export (<code>expdp</code>) and the Data Pump Import (<code>impdp</code>) utilities.has <code>SELECT</code> permissions on <code>sys.v\$instance</code> view. Start SQL*Plus and type: <code>grant select on v_\$instance to recovery_catatalog_user;</code>
Standby database user	Required to log in to the standby database. Applicable only in Oracle Data Guard environments. Needed to back up the standby database.

Configuring Oracle databases

Configuration of an Oracle database consists of providing Data Protector with the following data:

- Oracle Server home directory
- Login information to the target database
- Optionally, login information to the recovery catalog database
- Optionally, login information to the standby database

- Optionally, ASM-related information.
- Backup method to be used and the related options

During the configuration, the `util_oracle8.pl` command, which is started on the application system, saves the specified parameters in the Data Protector Oracle database specific configuration file on the Cell Manager.

Ensure that the database is open during the configuration procedure and that you are able to connect to the database.

To configure an Oracle database, you can use the Data Protector GUI or the Data Protector CLI. However, to prepare an Oracle environment which uses Automatic Storage Management (ASM), and the ASM database uses a different home directory or the Data Protector Oracle integration agent must connect to it through a corresponding net service, you must use the Data Protector CLI.

Using the Data Protector GUI

Configure an Oracle database when you create the first ZDB backup specification for the database. Start with the procedure described in “[Creating backup specifications](#)” (page 48) and at [Step 6](#) proceed as follows:

1. In the **Configure Oracle** dialog box and in the **General** page, specify the pathname of the Oracle Server home directory.

Figure 9 Configuring Oracle - General (Windows)





Figure 10 Configuring Oracle - General (UNIX)



2. In the **Primary** page, specify the login information to the primary database.

Note that the user must have the `SYSDBA` privilege granted.

In **Services**, type the net service name for the primary database instance. The backup will be performed on the system where this database instance resides.

RAC: List all net services names for the primary database separated by a comma.

Figure 11 Configuring Oracle - Primary

The screenshot shows the 'Configure Oracle' dialog box with the 'Primary' tab selected. The 'General' tab is also visible. The 'Oracle login information to primary database' section contains three text boxes: 'Username' with 'system', 'Password' with 'XXXXXXXX', and 'Services' with 'NETSERVICE1, NETSERVICE2'.

3. In the **Catalog** page, select **Use target database control file instead of recovery catalog** to use the primary database control file.

To use the recovery database catalog as an RMAN repository for backup history, select **Use recovery catalog** and specify the login information to the recovery catalog.

Note that for ZDB, you must use the recovery catalog.

The user specified must be the owner of the recovery catalog.

In **Services**, type the net service name for the recovery catalog.

Figure 12 Configuring Oracle - Catalog

The screenshot shows the 'Configure Oracle' dialog box with the 'Catalog' tab selected. The 'General' and 'Primary' tabs are also visible. The 'Use recovery catalog' radio button is selected. The 'Oracle login information to recovery database' section contains three text boxes: 'User name' with 'rman', 'Password' with 'XXXXXXXX', and 'Services' with 'CATSERVICE'.

4. If you have Oracle Data Guard configuration for *non-ZDB sessions* and if you intend to back up a standby database, configure also the standby database:

In the **Standby** page, select **Configure standby database** and specify the login information to the standby database.

In **Services**, type the net service name for the standby database instance.

RAC: List all net services names for the standby database separated by a comma.

Figure 13 Configuring Oracle - Standby

The screenshot shows the 'Configure Oracle' dialog box with the 'Standby' tab selected. The 'General', 'Primary', and 'Catalog' tabs are also visible. The 'Configure standby database' checkbox is checked. The 'Oracle login information to standby database' section contains three text boxes: 'Username' with 'system', 'Password' with 'XXXXXXXX', and 'Services' with 'NETSERVICESB1, NETSERVICESB2'.


5. In the **ZDB** page, select **Backup method** and then select **PROXY** or **BACKUP SET** in the drop-down list.

In **Backup control file copy location**, you can specify the location on the source volumes where a backup copy of the current control file will be made during ZDB to disk.

If you do not specify the location, `ob2rman.pl` will copy the copy of the control file from the application system to the backup system when it is needed. Thus, you do not need to create an additional disk for this location if you do not need the control file copy on a replica.

If your backup method is *backup set* and if your database instance uses PFILE (and not SPFILE), select the **Parameter file (PFILE)** option and specify the pathname of PFILE residing on the application system.

Figure 14 Configuring Oracle - ZDB



Click **OK**.

The Oracle database is configured. Exit the GUI or proceed with creating the backup specification at [Step 7](#).

Using the Data Protector CLI

1. **UNIX systems only:** Log in to the Oracle Server system with an OSDBA user account.
2. On the Oracle Server system, from the directory:

Windows systems: `Data_Protector_home\bin`

UNIX systems: `/opt/omni/lbin`

run:

Windows systems:

```
perl -I..\lib\perl util_oracle8.pl -config -dbname DB_NAME -orahome ORACLE_HOME PRIMARY_DB_LOGIN [CATALOG_DB_LOGIN] [STANDBY_DB_LOGIN] [ZDB_OPTIONS] [ASM_OPTIONS] [-client CLIENT_NAME]
```

UNIX systems:

```
util_oracle8.pl -config -dbname DB_NAME -orahome ORACLE_HOME PRIMARY_DB_LOGIN [CATALOG_DB_LOGIN] [STANDBY_DB_LOGIN] [ZDB_OPTIONS] [ASM_OPTIONS] [-client CLIENT_NAME]
```

where:

`PRIMARY_DB_LOGIN` is:

```
-prouser PRIMARY_USERNAME
```

```
-prpasswd PRIMARY_PASSWORD
```

```
-prmservice PRIMARY_NET_SERVICE_NAME_1[, PRIMARY_NET_SERVICE_NAME_2 ...]
```

`CATALOG_DB_LOGIN` is:

```

-rcuser CATALOG_USERNAME
-rcpasswd CATALOG_PASSWORD
-rcservice CATALOG_NET_SERVICE_NAME
STANDBY_DB_LOGIN is:
-stbuser STANDBY_USERNAME
-stbpasswd STANDBY_PASSWORD
-stbservice STANDBY_NET_SERVICE_NAME_1[,STANDBY_NET_SERVICE_NAME_2 ...]
ZDB_OPTIONS are:
-zdb_method {PROXY | BACKUP_SET}
[-ctlcp_location BACKUP_CONTROL_FILE_COPY_LOCATION]
[-pfile PARAMETER_FILE]
[-bkphost BACKUP_SYSTEM]
ASM_OPTIONS are:
[-asmhome ASM_HOME]
[-asmuser ASM_USER -asmpasswd ASM_PASSWORD -asmervice
ASM_NET_SERVICE_NAME_1[,ASM_NET_SERVICE_NAME_2 ...]]

```

If you have Oracle Data Guard configuration for *non-ZDB sessions* and if you intend to back up a standby database, you must provide the *STANDBY_DB_LOGIN* information.

To configure an Oracle database for ZDB, you must provide the *ZDB_OPTIONS* information. If your ZDB method is *backup set*, you must also provide the *BACKUP_SYSTEM* information.

To prepare an Oracle environment which uses Automatic Storage Management (ASM), and the ASM database uses a different home directory or the Data Protector Oracle integration agent must connect to it through a corresponding net service, you must provide the *ASM_OPTIONS* information.

Parameter description

CLIENT_NAME

Name of the Oracle Server system with the database to be configured. It must be specified in a cluster environment or if the ZDB configuration is run on the backup system.

RAC: The virtual server of the Oracle resource group.

Oracle Data Guard: Name of either a primary system or secondary (standby) system.

DB_NAME

Name of the database to be configured.

ORACLE_HOME

Pathname of the Oracle Server home directory.

PRIMARY_USERNAME PRIMARY_PASSWORD

Username and password for login to the target or primary database. Note that the user must have the SYSDBA privilege granted.

PRIMARY_NET_SERVICE_NAME_1 [, *PRIMARY_NET_SERVICE_NAME_2*, ...]

Net services names for the primary database.

RAC: Each net service name must resolve into a specific database instance.

CATALOG_USERNAME CATALOG_PASSWORD

Username and password for login to the recovery catalog. This is optional and is used only if you use the recovery catalog database as an RMAN repository for backup history.

CATALOG_NET_SERVICE_NAME

Net service name for the recovery catalog.

STANDBY_USERNAME STANDBY_PASSWORD

This is used in Oracle Data Guard environment for backing up a standby database. Username and password for login to the standby database.

STANDBY_NET_SERVICE_NAME_1 [,STANDBY_NET_SERVICE_NAME_2, ...]

Net services names for the standby database.

BACKUP_CONTROL_FILE_COPY_LOCATION

A location on a source volume where a copy of the current control file is made before a ZDB to disk. This is optional and if not specified, `ob2rman.pl` will copy the copy of the control file from the application system to the backup system when it is needed. Thus, you do not need to create an additional disk for this location if you do not need the control file copy on a replica.

PARAMETER_FILE

Full pathname of the PFILE residing on the application system. This is optional and used if backup method is backup set and the database instance uses PFILE (and not SPFILE).

BACKUP_SYSTEM

Name of the backup system. It must be specified for a ZDB backup set configuration.

ASM_HOME

Home directory of the ASM database in an Oracle ASM configuration.

ASM_USERNAME ASM_PASSWORD

User name and password (authentication credentials) used by the Data Protector Oracle integration agent to connect to the ASM database.

ASM_NET_SERVICE_NAME_1 [,ASM_NET_SERVICE_NAME_2 ...]

Name of the net service to be used to access the ASM database. For Oracle environments involving multiple net services, multiple names can be specified.

The message `*RETVAL*0` indicates successful configuration, even if followed by additional messages.

Example

The following example represents configuration on a UNIX system of an Oracle database and its recovery catalog with the backup set method used and the parameter file location specified.

The following names are used in the example:

- database name: `oracle`
- Oracle Server home directory: `/app10g/oracle10g/product/10.1.0`
- primary user name: `system`
- primary password: `manager`
- primary net service name 1: `netservice1`
- primary net service name 2: `netservice2`
- recovery catalog user name: `rman`
- recovery catalog password: `manager`
- recovery catalog net service name: `catSERVICE`
- backup system name: `bcksys`
- ASM user name: `asm`
- ASM password: `asmmanager`
- ASM net service name: `netSERVICEasm`

Syntax

```
/opt/omni/sbin/util_oracle8.pl -config -dbname oracle -orahome
/app10g/oracle10g/product/10.1.0 -prouser system -prmpasswd manager
-prmservice netservice1,netSERVICE2 -rcuser rman -rcpasswd manager
-rcservice catSERVICE -zdb_method BACKUP_SET -pfile
/app10g/oracle10g/product/10.1.0/dbs/pfile.ora -bkphost bcksys -asmuser
asm -asmpasswd asmmanger -asmSERVICE netSERVICEasm
```

If you need to export some variables before starting SQL*Plus, listener, or RMAN, these variables must be defined in the **Environment** section of the Data Protector Oracle global configuration file or using the Data Protector GUI.

Checking the configuration

You can check the configuration of an Oracle database after you have created at least one backup specification for the database. If you use the Data Protector CLI, a backup specification is not needed.

Using the Data Protector GUI

1. In the Context List, select **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **Oracle Server**. Click the backup specification to display the server with the database to be checked.
3. Right-click the server and click **Check configuration**.

❗ **IMPORTANT:** Data Protector does not check if the specified user has appropriate Oracle backup permissions.

Using the Data Protector CLI

1. **UNIX systems only:** Log in to the application system with an OSDBA user account.
2. From the directory:

Windows systems: *Data_Protector_home\bin*

UNIX systems: */opt/omni/sbin*

run:

Windows systems:

perl -I..\lib\perl util_oracle8.pl -chkconf_smb -dbname DB_NAME

UNIX systems:

util_oracle8.pl -chkconf_smb -dbname DB_NAME

Handling errors

If an error occurs, the error number is displayed in the form **RETVAL*error_number*.

To get the error description, on the Cell Manager, run:

Windows systems: *Data_Protector_home\bin\omnigetmsg 12 error_number*

UNIX systems: */opt/omni/sbin/omnigetmsg 12 error_number*

❗ **IMPORTANT:** On UNIX systems, it is possible that although you receive **RETVAL*0*, backup still fails because Data Protector does not check if the specified user has appropriate Oracle backup permissions.

Checking configuration for instant recovery

Check if the Oracle configuration is suitable for instant recovery.

On the application system, from the directory:

Windows systems: `Data_Protector_home\bin`

UNIX systems: `/opt/omni/lbin`

run:

Windows systems:

```
perl util_oracle8.pl -chkconf_ir -dbname DB_NAME
```

UNIX systems:

```
util_oracle8.pl -chkconf_ir -dbname DB_NAME
```

If the control files, SPFILE, and online redo logs are on the same volume group (if LVM is used) or source volume as datafiles, a warning is displayed stating that instant recovery is not possible. You can either:

- Reconfigure the Oracle database instance. See [“Reconfiguring an Oracle instance for instant recovery” \(page 269\)](#) on how to move the control files and redo logs to source volumes that are not replicated.
Or:
- Set the `ZDB_ORA_INCLUDE_CF_OLF`, `ZDB_ORA_INCLUDE_SPF`, and `ZDB_ORA_NO_CHECKCONF_IR` omnirc variables and ignore the warning. However, note that the control file, SPFILE, and online redo logs are overwritten during instant recovery. See [“ZDB integrations omnirc variables” \(page 271\)](#) on how to set the omnirc variables.

Setting environment variables

Use environment variables to modify backup environment to suit your needs. Environment variables are Oracle database specific. It means that they can be set differently for different Oracle databases. Once specified, they are saved to related Data Protector Oracle database configuration files.

For details of how environment variables affect your environment, see [Table 5 \(page 45\)](#).

NOTE: Environment variables are not supported on HP OpenVMS systems.

Table 5 Environment variables

Environment variable	Default value	Description
OB2_RMAN_COMMAND_TIMEOUT	300 s	This variable is applicable when Data Protector tries to connect to a target or catalog database. It specifies how long (in seconds) Data Protector waits for RMAN to respond that the connection succeeded. If RMAN does not respond within the specified time, Data Protector aborts the current session.
OB2_SQLP_SCRIPT_TIMEOUT	300 s	This variable is applicable when Data Protector issues an SQL*Plus query. It specifies how long Data Protector waits for SQL*Plus to respond that the query completed successfully. If SQL*Plus does not respond within the specified time, Data Protector aborts the current session.
OB2_DPMCTL_SHRLOC	N/A	Defines the path for the control file, where it is created and backed up as a part of the Data Protector managed control file backup. By default Data Protector copies the control file to the following location: <code>/var/opt/omni/tmp</code> (UNIX systems) or <code>Data_Protector_home\tmp</code> (Windows systems) directory. This variable is used to redirect this file to a specified location. With Oracle 11.2.0.2 or subsequent versions in a RAC environment, ensure this location is on a

Table 5 Environment variables (continued)

Environment variable	Default value	Description
		shared disk. This way, all the nodes can access the control file and the session finishes successfully.
ORA_ASM_LCL_INSTANCE	N/A	Specifies the Oracle Server instance name of the ASM database to be used when Data Protector Oracle integration agent connects directly to the database, without using a corresponding net service. If this variable is not set, the agent uses the instance name <i>SID</i> +ASM where <i>SID</i> is the instance name of the primary database.

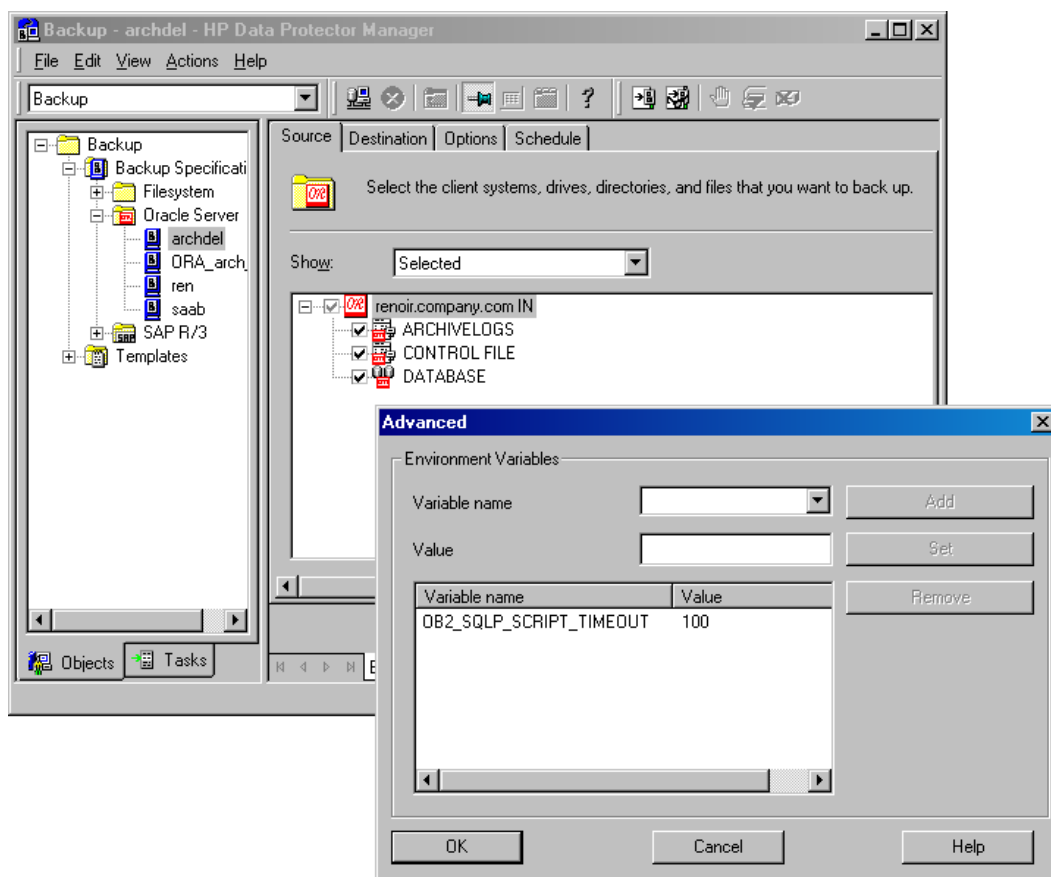
To set environment variables, use the Data Protector GUI or CLI.

Using the Data Protector GUI

You can set a variable when you create a backup specification or modify an existing one:

1. In the Source page of the backup specification, right-click the Oracle database at the top and click **Set Environment Variables**.
2. In the Advanced dialog box, specify the variable name, its value, and click **Add**. See Figure 15 (page 46).

Figure 15 Setting environment variables



Click **OK**.

Using the Data Protector CLI

From the directory:

Windows systems: *Data_Protector_home\bin*

UNIX systems: /opt/omni/sbin/

run:

```
util_cmd -putopt Oracle8 DatabaseName Variable Value -sublist Environment
```

For details, see the `util_cmd` man page or the *HP Data Protector Command Line Interface Reference*.

Example

To set the environment variable `OB2_RMAN_COMMAND_TIMEOUT` to 100 seconds for the Oracle database `INST2`, run:

```
util_cmd -putopt Oracle8 INST2 OB2_RMAN_COMMAND_TIMEOUT 100 -sublist Environment
```

Switching between Oracle backup methods

You can switch between the Oracle backup methods by reconfiguring the Data Protector Oracle integration for each database. It is *not* possible to select the method during the backup specification creation.

-
- ❗ **IMPORTANT:** When switching between the Oracle backup set and proxy-copy methods, you must carefully follow the instructions given below to ensure a successful switch between both methods and to ensure that during a restore or recovery RMAN does not select backup objects backed up using different methods in one restore session. If such a mixed set is used, the restore procedure will fail.
-

To switch between the backup methods:

1. Successfully back up the entire database using the *currently* selected method.
2. To avoid selecting backup specifications with a backup method different than the current backup method, you may remove or move all ZDB backup specifications belonging to the selected database instance. The backup specifications are located on the Cell Manager in:

Windows systems: `Data_Protector_home\Config\Server\BarLists\Oracle8`

UNIX systems: `/etc/opt/omni/server/barlists/oracle8`

3. Re-configure the database with the new *method* selected while creating a new Oracle ZDB specification.
4. Optionally, if you switch *from backup set to proxy-copy*, you may:
 - a. On the Cell Manager, remove the file:

Windows systems: `Data_Protector_home\Config\Server\Integ\Config\Oracle8\client_name%initDB_NAME_bckp.ora`

UNIX systems: `/etc/opt/omni/server/integ/config/Oracle8/client_name%initDB_NAME_bckp.ora`

- b. Remove the Oracle software from the backup system.
5. Perform ZDB of the entire database.

-
- ❗ **IMPORTANT:** If you need to perform a restore from a time between the start and the end of the first backup of the entire database using the new backup method, RMAN may try to use backup files from old method through a channel allocated for the files from the old method and the restore will fail. See [“Problem” \(page 97\)](#) on how to restore such a backup.
-

Backup

To configure an Oracle ZDB, perform the following steps:

1. Configure the devices you plan to use for a backup. For instructions, see the online Help index: "configuring devices".
2. Configure media pools and media for a backup. For instructions, see the online Help index: "creating media pools".
3. Ensure you are able to connect to the database.
4. Configure a non-ZDB backup specification and run the backup of Oracle data on the application system to verify that you have properly configured the Oracle environment. On how to create a non-ZDB backup specification, see the *HP Data Protector Integration Guide for Oracle and SAP*.
5. Create a Data Protector Oracle ZDB backup specification. See "Creating backup specifications" (page 48).

Creating backup specifications

Online ZDB

To perform an online ZDB of an Oracle database, the database has to run in the ARCHIVELOG mode.

Offline ZDB

To perform an offline ZDB, create only a ZDB backup specification.

Cluster-aware systems

Before you perform an *offline* ZDB in a cluster environment, take the Oracle Database resource offline and bring it back online after the replica is created. This can be done using the Oracle `fscommand` command line interface commands in the `Pre-exec` and `Post-exec` commands for the client system in a particular backup specification, or by using the Cluster Administrator.

You cannot perform a ZDB of the archived redo log files. Therefore, you need to create two backup specifications:

- ZDB backup specification for backing up database files
- standard Data Protector Oracle integration backup specification for backing up the application system archived log files

Procedure

To create an Oracle ZDB backup specification:

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **Oracle Server**, and click **Add Backup**.
3. In the Create New Backup dialog box, select the following:

Backup set method

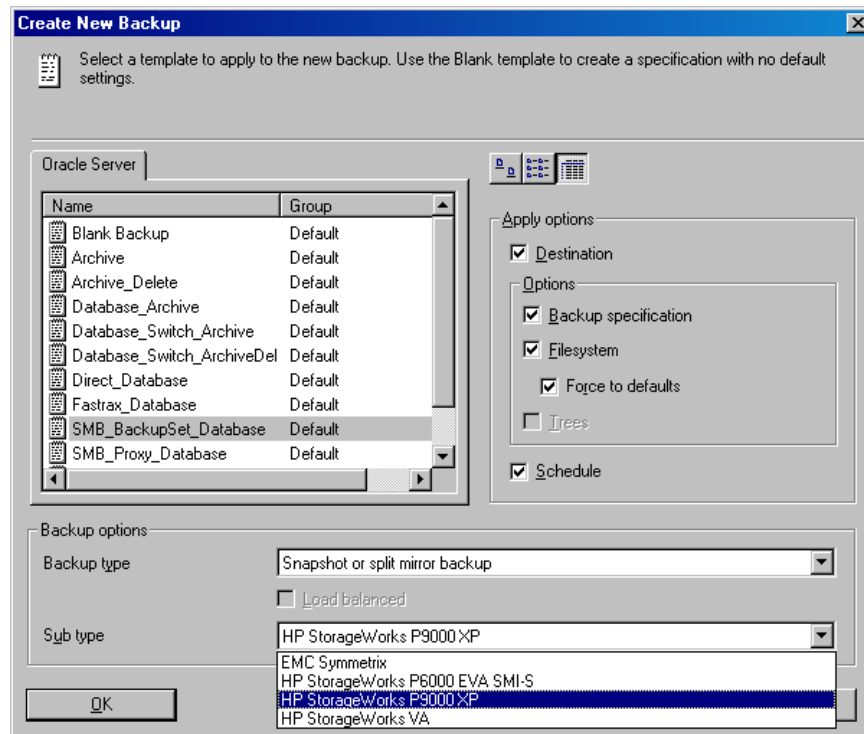
To perform a ZDB of the entire database using the backup set method, select the **SMB_BackupSet_Database** template.

Proxy-copy method

To perform a ZDB of the entire database using the proxy-copy method, select the **SMB_Proxy_Database** template.

From the **Backup type** drop-down list, select **Snapshot or split mirror backup**, and from the **Sub type** drop-down list, select the appropriate disk array agent. The agent must be installed on the application system and the backup system. See "Selecting an Oracle backup template and the snapshot or split mirror backup" (page 49).

Figure 16 Selecting an Oracle backup template and the snapshot or split mirror backup



Click **OK**.

4. In **Application system**, select the Data Protector Oracle integration client. In a non-RAC cluster environment, select the virtual server.

RAC: Select the virtual server of the Oracle resource group.

In **Backup system**, select the backup system.

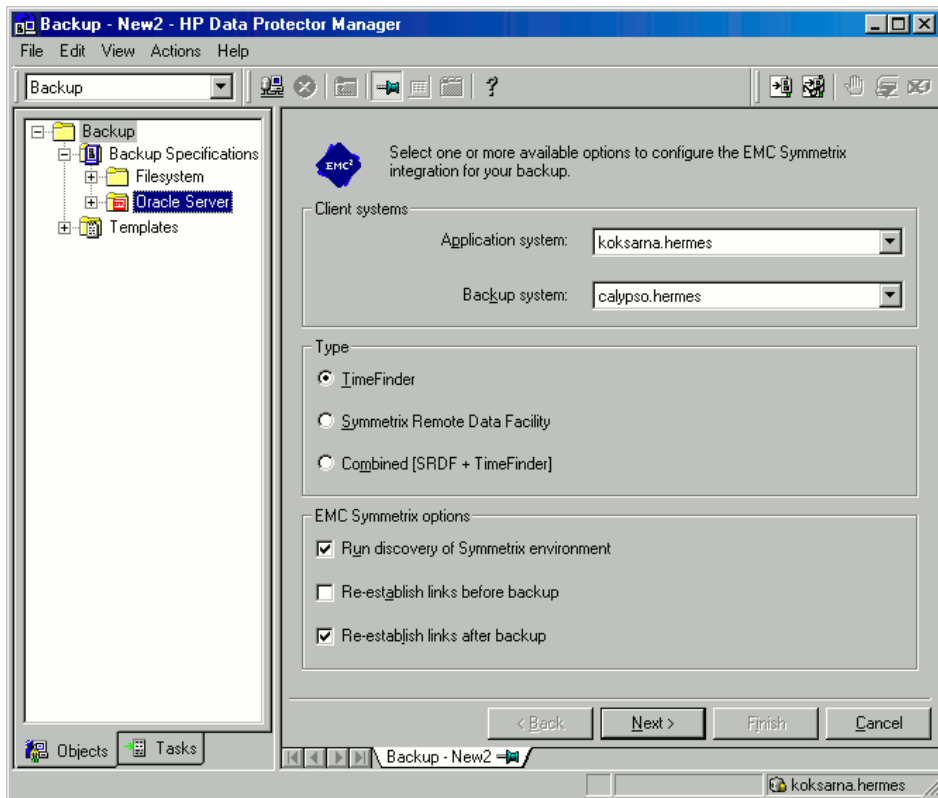
Select other disk array-specific backup options (see [“EMC backup options”](#) (page 50) for EMC, [“P9000 XP Array backup options”](#) (page 50) for P9000 XP Array, or [“P6000 EVA Array backup options”](#) (page 51) for P6000 EVA Array. For detailed information on the backup options, press **F1**.

EMC GeoSpan specifics

In the EMC GeoSpan for Microsoft Cluster Service environment, select the backup system for the active node and specify the TimeFinder configuration.

After a failover in EMC GeoSpan for MSCS, select the backup system for the currently active node and save the backup specification.

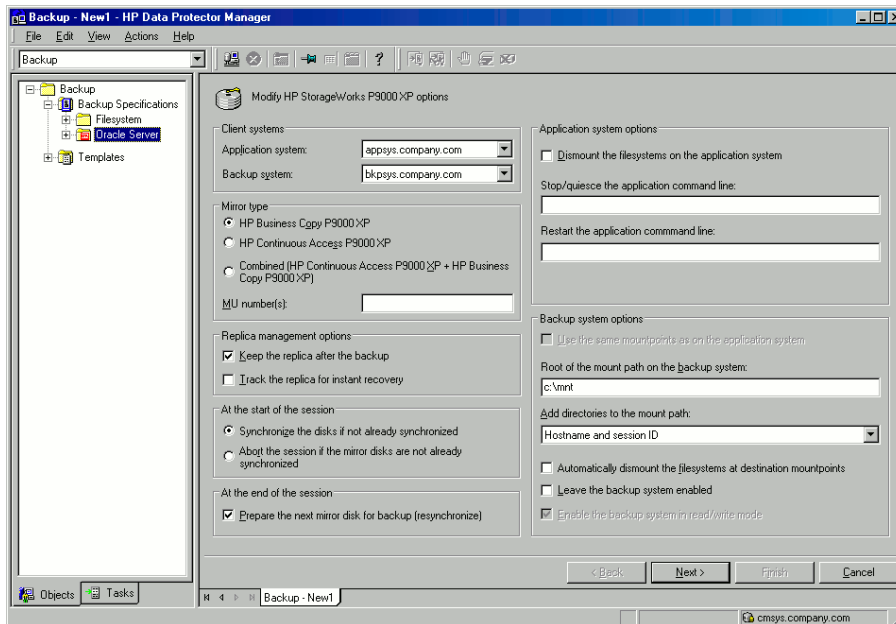
Figure 17 EMC backup options



P9000 XP Array specifics

To enable instant recovery, leave the **Track the replica for instant recovery** option selected. It is not possible to run instant recovery with Data Protector if this option is cleared.

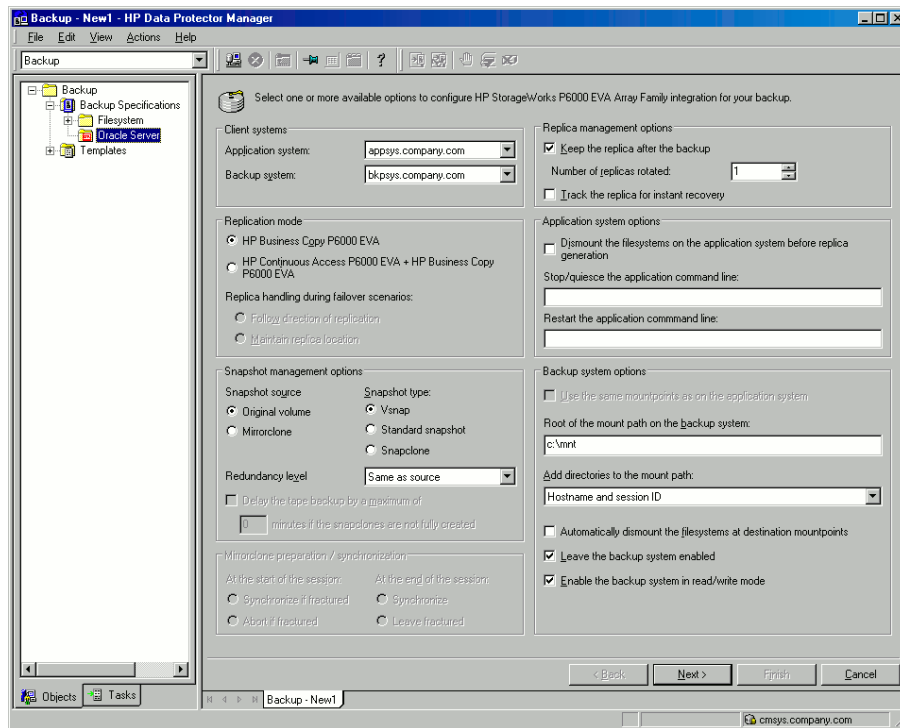
Figure 18 P9000 XP Array backup options



P6000 EVA Array specifics

To enable instant recovery, select the **Track the replica for instant recovery** option.

Figure 19 P6000 EVA Array backup options



Click **Next**.

5. In **Application database**, type the name of the database to be backed up.

The database name can be obtained using SQL*Plus:

```
SQL>select name from v$database;
```

NOTE: In a single-instance configuration, the database name is usually the same as its instance name. In this case, the instance name can be also used. The instance name can be obtained as follows:

```
SQL>select instance_name from v$instance;
```

Specify the **User and group/domain** options, which are available on UNIX and Windows Server 2008 systems, as follows:

- **UNIX systems:** In **Username** and **Group/Domain name**, specify the OSDBA user account under which you want the backup to start (for example, the user name ora, group DBA). This user must be configured as described in [“Configuring Oracle user accounts”](#) (page 36).
- **Windows Server 2008 systems:** It is not mandatory to specify these options and if they are not specified, the backup runs under the Local System Account.

In **Username** and **Group/Domain name**, specify the operating system user account under which you want the backup session to run (for example, the user name Administrator, domain DP).

Ensure that this user has been added to the Data Protector admin or operator user group and has the Oracle database backup rights. This user becomes the backup owner.

NOTE: If this is not your first backup specification, Data Protector fills in **Username** and **Group/Domain name** for you, providing the values of the last configured Oracle database.

Figure 20 Specifying an Oracle Server system (Windows)


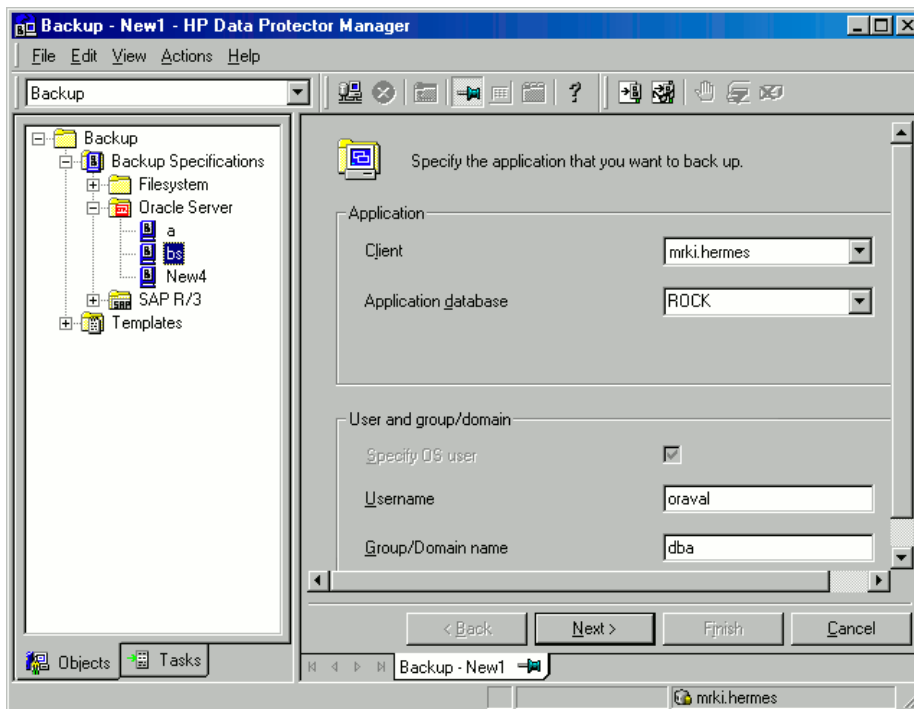


Figure 21 Specifying an Oracle Server system (UNIX)



Click **Next**.

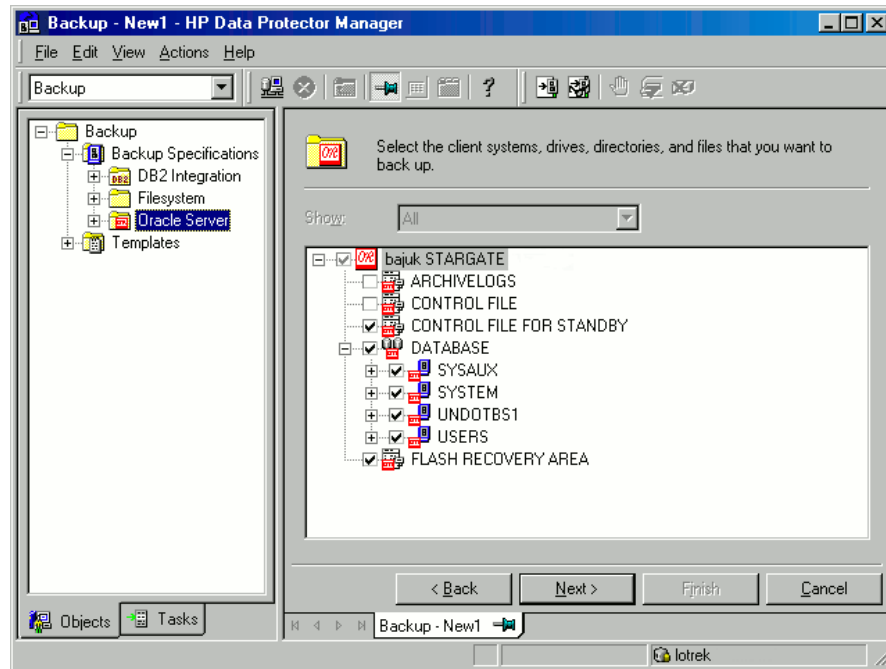
NOTE: When you click **Next**, Data Protector performs a configuration check.

UNIX systems only: The check is started under the specified OSDBA user account. If it completes successfully, the OSDBA user and group are also saved in both the Oracle database specific configuration file and Oracle system global configuration file, overriding previous values if they exist.

6. If the Oracle database is not configured yet for use with Data Protector, the Configure Oracle dialog box is displayed. Configure the Oracle database for use with Data Protector as described in [“Configuring Oracle databases” \(page 38\)](#).
7. Select the Oracle database objects to be backed up.

NOTE: Since temporary tablespaces do not contain permanent database objects, RMAN and Data Protector do not back them up. For more information, see Oracle documentation.

Figure 22 Selecting backup objects



Click **Next**.

If the backup method configured for this instance does not correspond to the method in the backup specification, Data Protector will display a warning and abort the configuration.

8. Select the device(s) you want to use for the backup. Click **Properties** to set the device concurrency, media pool, and preallocation policy. For more information on these options, click **Help**.

You can also specify whether you want to create additional copies (mirrors) of the backup during the backup session. Specify the desired number of mirrors by clicking the **Add mirror** and **Remove mirror** buttons. Select separate devices for the backup and for each mirror.

For detailed information on the object mirror functionality, see the online Help index: “object mirroring”.

NOTE: Object mirroring is not supported for ZDB to disk.

Click **Next** to proceed.

9. Set the backup options.


For information on other the Backup Specification Options and Common Application Options, press **F1**.

Offline ZDB

To perform an offline ZDB, select the **Backup offline** option in the Application Specific Options dialog box. This option stops the database before creating a replica, and restarts it after the replica is created. Note that if a ZDB-to-tape or ZDB-to-disk+tape session is being performed,

the database is not offline during the actual backup to tape. See [“Backup offline option ” \(page 54\)](#).

Figure 23 Backup offline option



For information on other Application Specific Options, see [“Oracle backup options” \(page 55\)](#) or press **F1**.

Click **Next**.

10. Optionally, schedule the backup. For more details, see [“Scheduling backup specifications” \(page 59\)](#).

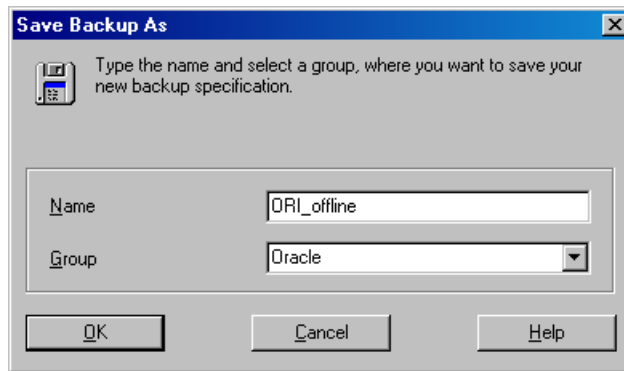
Note that only the **Full** backup type is supported.

Click **Next**.

11. Save the backup specification. It is recommended that you save all Oracle backup specifications in the **Oracle** group.

❗ **IMPORTANT:** The word `DEFAULT` is a reserved word and therefore must not be used for backup specification names or labels of any kind. Therefore, do not use a punctuation in the names of backup specifications, since the Oracle channel format is created from the backup specification name.

Figure 24 Saving the backup specification



Click **OK**.

To start the backup, see [“Starting backup sessions” \(page 59\)](#).

12. For **online backup**, create also a standard Data Protector Oracle integration backup specification for backing up the application system archived log files. See the *HP Data Protector Integration Guide*.

TIP: The backup specification for the backup of archived log files can be either triggered by the **Post-Exec** command defined in the ZDB backup specification for the backup of database files (recommended), or started manually after the ZDB backup specification has been started. See online Help index: “pre- and post-exec commands” for more information on configuring the **Pre-Exec** and **Post-Exec** commands.

Table 6 Oracle backup options

Disable recovery catalog auto backup	By default, Data Protector backs up the recovery catalog after every ZDB to tape or ZDB to disk+tape. Select this option to disable backup of the recovery catalog.
Disable Data Protector managed control file backup	By default, Data Protector backs up the Data Protector managed control file after every ZDB to tape or ZDB to disk+tape. Select this option to disable backup of the Data Protector managed control file.
Back up standby database	This option is ignored for ZDB.
RMAN Script	You can edit the Oracle RMAN script section of the Data Protector Oracle backup specification. The script is created by Data Protector during the creation of a backup specification and reflects the backup specification’s selections and settings. You can edit the script only after the backup specification has been saved. For information on how to edit the RMAN script section, see “Editing the Oracle RMAN script” (page 56) .
Pre-exec, Post-exec	Specify a command or RMAN script that will be started by <code>ob2rman.pl</code> on the Oracle Server system before the backup (<code>pre-exec</code>) or after it (<code>post-exec</code>). RMAN scripts must have the <code>.rman</code> extension. Do not use double quotes. For example, you can provide scripts to shut down and start an Oracle instance. For examples of shut-downing and starting an Oracle instance on a UNIX system, see “Examples of pre-exec and post-exec scripts on UNIX systems” (page 56) . Provide the pathname of the command or RMAN script.
Backup offline	Select this option to perform an offline ZDB session. This option stops the database before creating a replica, and restarts it after the replica is created. See “Backup offline option ” (page 54) .

Examples of pre-exec and post-exec scripts on UNIX systems

Pre-exec example

The following is an example of a script that *shuts down* an Oracle instance:

```
#!/bin/sh
export ORACLE_HOME=$2
export ORACLE_SQLNET_NAME=$1
if [ -f $ORACLE_HOME/bin/sqlplus ]; then
$ORACLE_HOME/bin/sqlplus << EOF
connect sys/manager@$ORACLE_SQLNET_NAME as sysdba
shutdown
EOF
echo "Oracle database \"$DB_NAME\" shut down."
exit 0
else
echo "Cannot find Oracle SQLPLUS ($ORACLE_HOME/bin/sqlplus)."
exit 1
fi
```

Post-exec example

The following is an example of a script that *starts* an Oracle instance:

```
#!/bin/sh
export ORACLE_HOME=$2
export ORACLE_SQLNET_NAME=$1
if [ -f $ORACLE_HOME/bin/sqlplus ]; then
$ORACLE_HOME/bin/sqlplus << EOF
connect sys/manager@$ORACLE_SQLNET_NAME as sysdba
startup
EOF
echo "Oracle database \"$DB_NAME\" started."
exit 0
else
echo "Cannot find Oracle SQLPLUS ($ORACLE_HOME/bin/sqlplus)."
exit 1
fi
```

Editing the Oracle RMAN script

The RMAN script is used when the Data Protector backup specification is started to perform a backup of the Oracle objects.

The RMAN script section is not written to the backup specification until the backup specification is either saved or manually edited by clicking the **Edit** button.

You can edit the RMAN script section of only after the Data Protector Oracle backup specification has been saved.

Limitations

When editing the RMAN script sections of the Data Protector backup specifications, consider the following limitations:

- The Oracle manual configuration convention must be used and not the Oracle automatic configuration convention.
- Double quotes (") must not be used - single quotes should be used instead.
- By default, RMAN scripts created by Data Protector contain instructions for backing up one or more of the following objects:
 - Databases, tablespaces, or datafiles (the first backup command)
 - Archive logs (the second backup command)
 - Control files (the last backup command)

The RMAN scripts with all combinations of the above listed backup objects are recognized by Data Protector as its own scripts and it is possible to modify the selection of objects that will be backed up in the **Source** tab of the Results Area.

If the RMAN script contains *additional* manually entered backup commands, for example a second backup command for backing up a database that is already listed in the first backup command, the object selection is disabled and it is only possible to browse the **Source** tab.

To edit an Oracle RMAN script, click **Edit** in the **Application Specific Options** window (see [“Recovery catalog settings dialog” \(page 64\)](#)), edit the script, and then click **Save** to save the changes to the script.

See the *Oracle Recovery Manager User’s Guide and References* for more information on Oracle RMAN commands.

Data Protector RMAN script structure

The RMAN script created by Data Protector consists of the following parts:

- **The Oracle channel allocation** together with the Oracle environment parameters’ definition for every allocated channel.

For all backup specifications except for Oracle proxy-copy ZDB backup specifications, the number of allocated channels is the same as the sum of concurrency numbers for all devices selected for backup.

NOTE: Once the backup specification has been saved, changing the concurrency number does not change the number of allocated channels in the RMAN script. This has to be done manually by editing the RMAN script.

- ❗ **IMPORTANT:** On Windows systems, a maximum of 32 or 64 (if device is local) channels can be allocated. If the calculated number exceeds this limitation, you have to manually edit the RMAN script and reduce the number of allocated channels.
-

When an Oracle channel is manually defined by editing the RMAN script, the environment parameters must be added in the following format:

```
parms 'ENV=(OB2BARTYPE=Oracle8, OB2APPNAME=DB_NAME,
OB2BARLIST=Backup_Specification_Name) ';
```

Proxy-copy

For Oracle proxy-copy ZDB backup sessions, Data Protector allocates *one* channel.

For Oracle proxy-copy ZDB, the OB2SMB parameter must be set to 1. If you use the Blank Oracle Backup template, the number of concurrently running DMA (OB2DMAP) is automatically calculated as the sum of all device concurrences; for example, if there are 4 devices with concurrency set to 3 then OB2DMAP will be set to 12.

If you use the Oracle_SMB template, the OB2DMAP parameter is set to 1. To improve the backup and restore performance, you may want to increase the value of this parameter. The environment parameters must be added in the following format:

```
parms 'ENV=(OB2BARTYPE=Oracle8, OB2APPNAME=DB_NAME,
OB2BARLIST=Backup_Specification_Name, OB2SMB=1,
OB2DMAP=Concurrent_DMAs) ';
```

NOTE: The OB2DMAP parameter does not change after it has been calculated, even if you adjust the device concurrency. To change OB2DMAP, you have to manually edit the RMAN script.

- Depending on the backup objects selection, **an RMAN backup statement for the backup of the whole database instance, and/or for any combination of RMAN commands to back up tablespaces and datafile**. The backup statement consists of the following:
 - The Oracle format of the backup file in the following format:

```
format 'Backup_Specification_Name<DB_NAME_%s:%t:%p>.dbf' database;
```
- In case of an Oracle proxy-copy ZDB-to-disk+tape or ZDB-to-tape session, the PROXY ONLY option is required. Only one BACKUP command with the proxy only option is permitted and only one additional backup command for backing up the control file is permitted.
- The RMAN datafile `tablespace_name*datafile_name` command.
- If the Archived Redo Logs were selected for a backup, **an RMAN backup statement for the backup of Oracle archive logs**.

The backup statement consists of the Oracle format of the backup file:

```
format 'Backup_Specification_Name<DB_NAME_%s:%t:%p>.dbf'
```

NOTE: When an Oracle format of the backup file is manually defined or changed by editing the RMAN script, any user-defined combination of the Oracle substitution variables can be added to the obligatory `%s:%t:%p` substitution variables and `DB_NAME`.

- If the control file was selected for a backup, **an RMAN backup statement for the backup of Oracle control files**. The backup statement consists of the following:
 - The Oracle format of the backup file in the following format:

```
format 'Backup_Specification_Name<DB_NAME_%s:%t:%p>.dbf' current controlfile;
```
- The RMAN `current controlfile` command.

For Oracle proxy-copy ZDB to disk or disk+tape, it is not possible to select only the control file. You must also select either a DATABASE, TABLESPACE, or DATAFILE object.

Example of the Oracle proxy-copy ZDB to disk+tape RMAN script

The following is an example of the RMAN script section as created by Data Protector based on the Oracle SMB_Proxy_Database template, after the whole database selection:

```
run {
allocate channel 'dev_0' type 'sbt_tape' parms
'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DIPSI,OB2BARLIST=New1,
OB2SMB=1,OB2DMAP=1)';
backup incremental level <incr_level>
format 'New1<DIPSI_%s:%t:%p>.dbf'
proxy only
```

```

database
;
backup format 'New1<DIPSI_%s:%t:%p>.dbf' controlfile;
}

```

Starting backup sessions

To run a ZDB-to-disk, ZDB-to-tape, or ZDB-to-disk+tape backup session of an Oracle database, use any of the following methods:

Backup methods

- Schedule a backup of an existing Oracle ZDB backup specification using the Data Protector Scheduler. See [“Scheduling backup specifications”](#) (page 59).
- Start an interactive backup of an existing Oracle ZDB backup specification using the Data Protector GUI or the Data Protector CLI. See [“Running an interactive backup”](#) (page 60).

Considerations

Before running an Oracle ZDB session, consider the following:

- It is not possible to start a ZDB, restore, or instant recovery sessions using the same source volume on the application system at the same time. A ZDB, restore, or instant recovery session must be started only after the preceding session that is using the same source volume on the application system has finished the ZDB session or restore; otherwise, the session will fail.
- For the backup set method, if the Oracle database is installed on symbolic links, then these symbolic links have to be also created on the backup system.
- On P9000 XP Array, if the LVM mirroring configuration is used, Data Protector displays a warning during a backup because the volume group source volumes on the application system do not have their HP BC P9000 XP pairs assigned. This message should be ignored.
- If the control file, SPFILE, or online redo logs are on the same source volumes as the datafiles and the **Track the replica for instant recovery** option is selected, the backup session will be aborted. In this case, you need to either reconfigure the database or set the `ZDB_ORA_INCLUDE_CF_OLF`, `ZDB_ORA_INCLUDE_SPF`, and `ZDB_ORA_NO_CHECKCONF_IR` omnirc variables. See [“Reconfiguring an Oracle instance for instant recovery”](#) (page 269) or [“ZDB integrations omnirc variables”](#) (page 271).

Scheduling backup specifications

Scheduling a backup specification means setting the time, date, and type of a backup that starts unattended once the scheduling options are defined and saved in the backup specification.

For more information on scheduling, see the online Help index: “scheduled backups”.

To schedule an Oracle ZDB backup specification, proceed as follows:

1. In the Data Protector Manager, switch to the **Backup** context.
2. In the **Scoping Pane**, expand **Backup Specifications** and then **Oracle Server**.
3. Double-click the backup specification you want to schedule and click the **Schedule** tab.
4. In the **Schedule** page, select a date in the calendar and click **Add** to open the **Schedule Backup** dialog box.

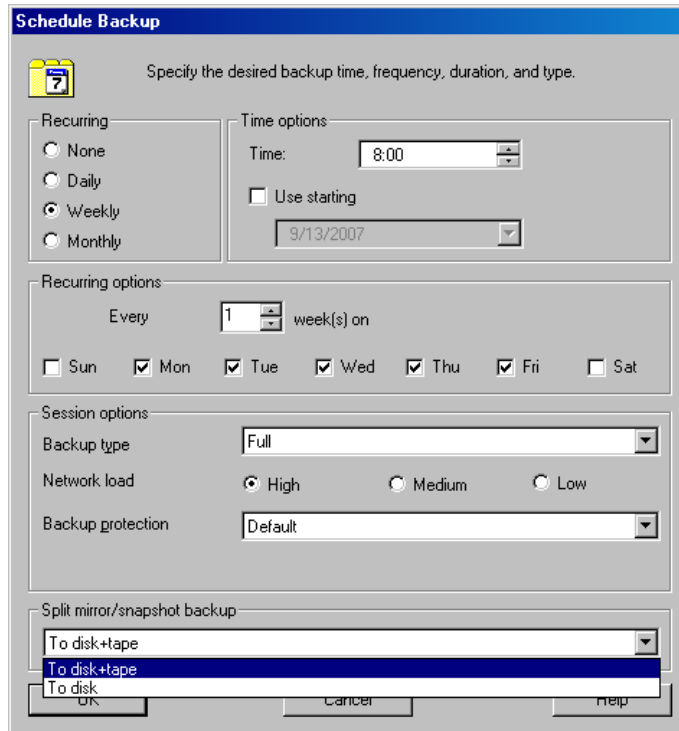
5. Specify **Recurring**, **Time options**, **Recurring options**, and **Session options**..

Note that only the `Full` backup type is supported.

In the case of a ZDB-to-disk or a ZDB-to-disk+tape session, specify the **Split mirror/snapshot backup** option. See “[Selecting ZDB to disk or ZDB to disk+tape session using the Data Protector scheduler](#)” (page 60).

NOTE: You can run a ZDB-to-disk or a ZDB-to-disk+tape session only if the **Track the replica for instant recovery** option is selected in the backup specification.

Figure 25 Selecting ZDB to disk or ZDB to disk+tape session using the Data Protector scheduler



Click **OK** and then **Apply** to save the changes.

Running an interactive backup

An interactive backup can be performed any time after a backup specification has been created and saved. You can use the Data Protector GUI or CLI.

Starting a backup using the GUI

To start an interactive ZDB session of an Oracle database using the Data Protector GUI, proceed as follows:

1. In the Context List, click **Backup** context.
2. In the Scoping Pane, expand **Backup Specifications** and then **Oracle Server**. Right-click the backup specification you want to use and click **Start Backup**.

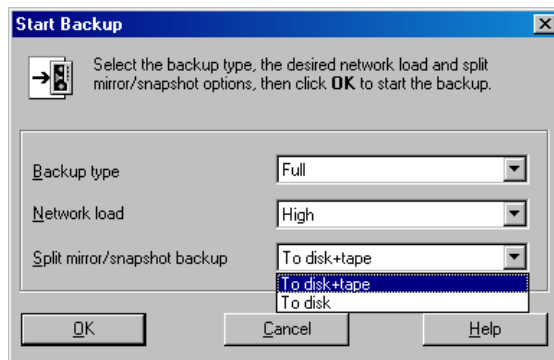
3. In the **Start Backup** dialog box, select the **Network load** option. For information on network load, click **Help**.

Note that only the **Full** backup type is supported.

In the case of a ZDB-to-disk or a ZDB-to-disk+tape session, specify the **Split mirror/snapshot backup** option. See [“Selecting ZDB to disk or ZDB to disk+tape session when starting an interactive backup”](#) (page 61).

NOTE: You can run a ZDB-to-disk or a ZDB-to-disk+tape session only if the **Track the replica for instant recovery** option is selected in the backup specification.

Figure 26 Selecting ZDB to disk or ZDB to disk+tape session when starting an interactive backup



Click **OK**.

Starting a backup using the CLI

To start an Oracle **ZDB-to-tape** or **ZDB-to-disk+tape** session using the Data Protector CLI, run:

```
omnib -oracle8_list Name
```

To start an Oracle **ZDB-to-disk** session using the Data Protector CLI, run:

```
omnib -oracle8_list Name -disk_only
```

where *Name* is the name of the backup specification. For more information on the `omnib` command, see its man page or the *HP Data Protector Command Line Interface Reference*.

NOTE: It is not possible to run a ZDB-to-disk or a ZDB-to-disk+tape session if the **Track the replica for instant recovery** backup option is not selected in the backup specification.

Restore

You can restore the following database objects using both the Data Protector GUI or RMAN:

- Control files
- Datafiles
- Tablespace
- Databases
- Recovery Catalog Databases

Using the Data Protector GUI, you can also **duplicate** a production database. See [“Duplicating an Oracle database”](#) (page 69).

The following are the available methods in Data Protector for restoring database objects:

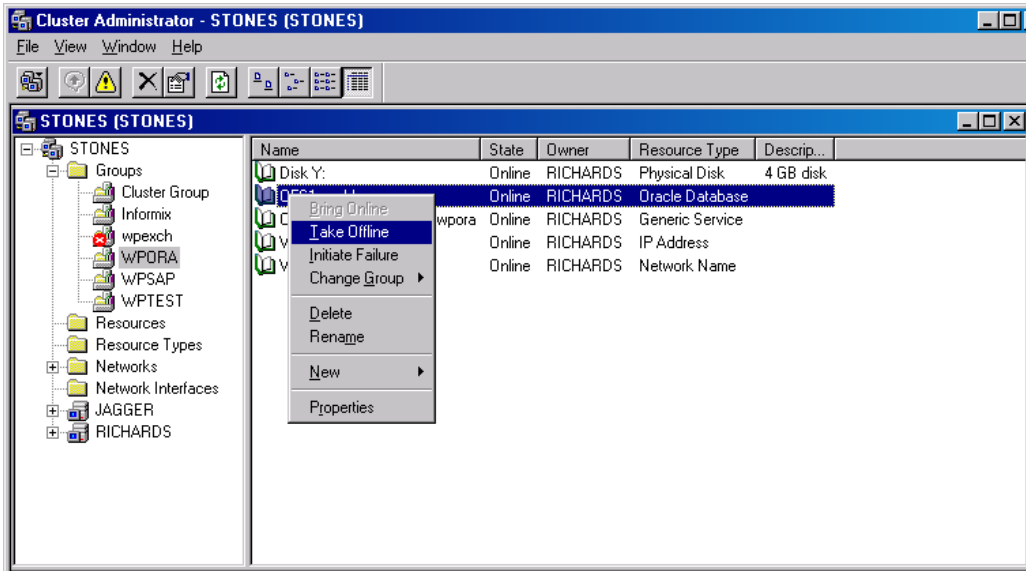
- Standard restore from backup media to the application system on LAN. See [“Restoring from backup media to the application system on LAN”](#) (page 63).
- Instant recovery. See [“Instant recovery and database recovery”](#) (page 85).

See also “Oracle recovery methods” (page 21) for an overview of recovery methods depending on the backup type and type of recovery.

Microsoft Cluster Server systems

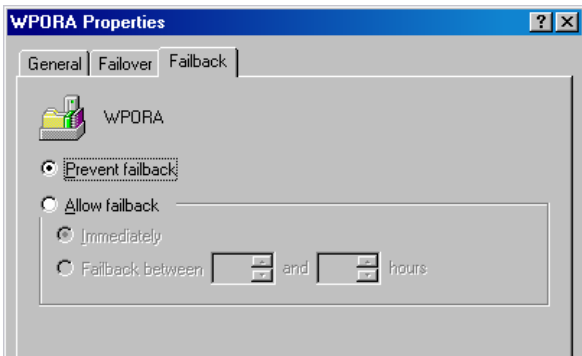
Before you start restoring a cluster-aware Oracle server, take the Oracle Database resource offline using, for example, the **Cluster Administrator** utility. See “Taking the Oracle resource group offline” (page 62).

Figure 27 Taking the Oracle resource group offline



Verify that you have set the **Prevent Fallback** option for the Oracle resource group and **Do not restart** for the `DB_NAME.world` resource, which is an Oracle Database resource.

Figure 28 Checking properties



MC/ServiceGuard systems

When restoring the database from a backup performed on a virtual host, you should set `OB2BARHOSTNAME` environment variable in the RMAN script. For example:

```
run {
allocate channel dev1 type 'sbt_tape' parms
  'SBT_LIBRARY=Path_to_Data_Protector_MML,
  ENV=(OB2BARHOSTNAME=virtual.domain.com)';
restore datafile '/opt/ora10g/oradata/MAKI/example02.dbf';
release channel dev1;
}
```

Prerequisites

- An instance of Oracle must be created on the system to which you want to restore or duplicate the database.
- The database must be in the `Mount` state if the whole database is being restored, or in the `NoMount` state if the control file is being restored or a database duplication is performed.
- You must be able to connect to the database.
One way of achieving this is by configuring static service information for your Oracle listener. For details, see the Oracle documentation. You can find an example of static service information configuration in the troubleshooting [“Instant recovery of an Oracle database fails”](#) (page 101).

Restoring from backup media to the application system on LAN

You can restore the database objects using one of the following tools within Data Protector:

- Data Protector GUI. See [“Restoring Oracle using the Data Protector GUI”](#) (page 63).
- RMAN. See [“Restoring Oracle using RMAN”](#) (page 74).

Restoring Oracle using the Data Protector GUI

For restore, RMAN scripts are generated with necessary commands, depending on selections made in the GUI. To use additional commands, use them manually from RMAN itself. You can also use the workaround described in [“How to modify the RMAN restore script”](#) (page 100).

Restoring database items in a disaster recovery

In a disaster recovery situation, database objects must be restored in a certain order. The following list shows you in which order database items must be restored. Under normal conditions it is possible to restore database items in any order.

1. Restore the recovery catalog database (if it was lost)
2. Restore the control file
3. Restore the entire database or data items

Changing the database state

Before you restore any database item or you perform a duplication of a database, ensure that the database is in the correct state:

Table 7 Required database states

Item to restore	Database state
Control file, duplicating a database	NoMount (started)
All other items ¹	Mount

¹ When restoring only a few tablespaces or datafiles, then the database can be open with the tablespaces or datafiles to be restored offline.

To put the database into the correct state, run:

```
sqlplus /nolog
```

```
SQL>connect user/password@service as sysdba;
```

```
SQL>shutdown immediate;
```

To put the database into `NoMount` state, run:

```
SQL>startup nomount;
```

To put the database into `Mount` state, run:

```
SQL>startup mount;
```

Restoring the recovery catalog database

The Oracle recovery catalog database is exported using the Oracle export utility to a binary file and backed up by Data Protector. This file has to be restored back to the disk and then imported into the Oracle database using the Oracle import utility. Data Protector provides a facility to do this automatically using the Oracle integration.

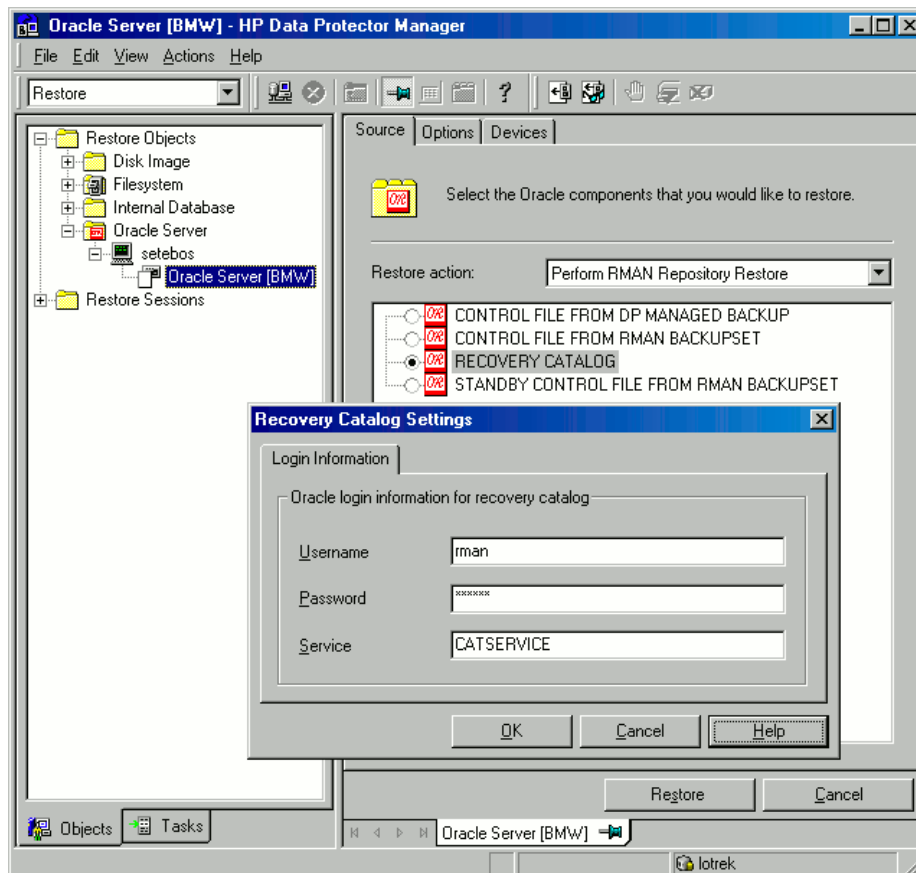
To restore the recovery catalog database:

1. Ensure that the recovery catalog database is in the **Open** state.
2. Remove the recovery catalog from the database (if it exists), using the RMAN command `DROP CATALOG`.
3. In the Data Protector GUI, switch to the **Restore** context.
4. Under **Restore Objects**, expand **Oracle Server**, expand the system on which the database, for which you want to restore the recovery catalog, resides, and then click the database.
5. In the **Restore action** drop-down list, select **Perform RMAN Repository Restore**.

In the Results Area, select **RECOVERY CATALOG**.

If you want to change the recovery catalog login information, right-click **RECOVERY CATALOG** and click **Properties**. In **Recovery Catalog Settings**, specify the login information for recovery catalog.

Figure 29 Recovery catalog settings dialog



6. In the **Options** page:
In **User name** and **User group**, specify the user name and password to the recovery catalog database.
From the **Session ID** drop-down list, select the Session ID.
For further information, see ["Restore, recovery, and duplicate options"](#) (page 71).
7. Click **Restore**.

Proceed to restore the control file.

Restoring the control file

The control file contains all the information about the database structure. If the control file has been lost, you must restore it before you restore any other part of the database. The database should be in the `NoMount` state.

Depending on the type of the control file backup, the following types of restore are possible when restoring the control file:

- Restoring from Data Protector managed control file backup (`CONTROLFILE FROM DP MANAGED BACKUP`)

The control file was backed up automatically by `ob2rman.pl` at the end of a backup session, unless the option `Disable Data Protector managed control file backup` was selected.

The recovery catalog is *not* required for this restore option.

The control files (`ctrlDB_NAME.dbf`) are restored to:

Windows systems: `Data_Protector_home\tmp`

UNIX systems: `/var/opt/omni/tmp`

NOTE: With Oracle 11.2.0.2 in a RAC environment, the control files are created, backed up and restored to the location specified with the `OB2_DPMCTL_SHRLOC` variable. This directory must be located on a shared disk and accessible from all nodes for a successful restore session.

After the restore, run the following script:

```
run {
  allocate channel 'dev0' type disk;
  restore controlfile from 'TMP_FILENAME';
  release channel 'dev0';
}
```

Where `TMP_FILENAME` is the location to which the file was restored.

- Restoring from RMAN backup set (`CONTROLFILE FROM RMAN BACKUPSET`)
The recovery catalog *is* required.

A backup session can contain more than one type of the control file backup.

To restore the control file:

1. Open the `sqlplus` window and put the database in the `nomount` state. See [“Changing the database state” \(page 63\)](#).
2. In the Data Protector GUI, switch to the **Restore** context.
3. Under **Restore Objects**, expand **Oracle Server**, expand the system on which the database, for which you want to restore the control file, resides, and then click the database.
4. In the **Restore Action** drop-down list, select **Perform RMAN Repository Restore**.
In the Results area, select the control file for restore.
5. In the **Options** page, from the **Client** drop-down list, select the system on which the Data Protector Oracle integration agent (`ob2rman.pl`) will be started. To restore the control file to a different database than it is selected, click **Settings** and specify the login information for the target database.
Set the other restore options. For information, see [“Restore, recovery, and duplicate options” \(page 71\)](#).
6. Click **Restore**.

Proceed with restoring the Oracle database objects.

Restoring Oracle database objects

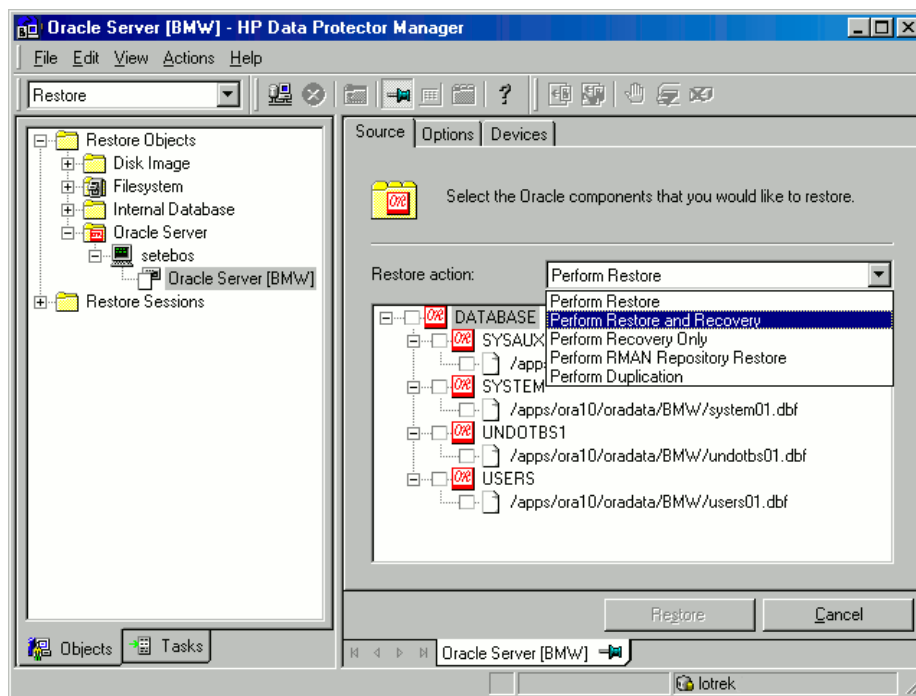
Before you restore Oracle database objects, ensure that you have an up-to-date version of the recovery catalog database and the control file. They contain the database structure information. If you do not have up-to-date versions of these files, restore them as described in “Restoring the recovery catalog database” (page 64) and “Restoring the control file” (page 65).

To restore Oracle database objects:

1. Put the database in the mount state. See “Changing the database state” (page 63).
2. In the Data Protector GUI, switch to the **Restore** context.
3. Under **Restore Objects**, expand **Oracle Server**, expand the system on which the database, for which you restore the database objects, resides, and then click the database.
4. In the **Restore action** drop-down list, select the type of restore you wish to perform. For information on the options, see “Restore, recovery, and duplicate options” (page 71).

❗ **IMPORTANT:** If you do not select **Perform Restore and Recovery** or **Perform Recovery Only**, you will have to recover the database objects manually using RMAN. For information, see “Restoring Oracle using RMAN” (page 74).

Figure 30 Source page



5. In the Results Area, select objects for restore.

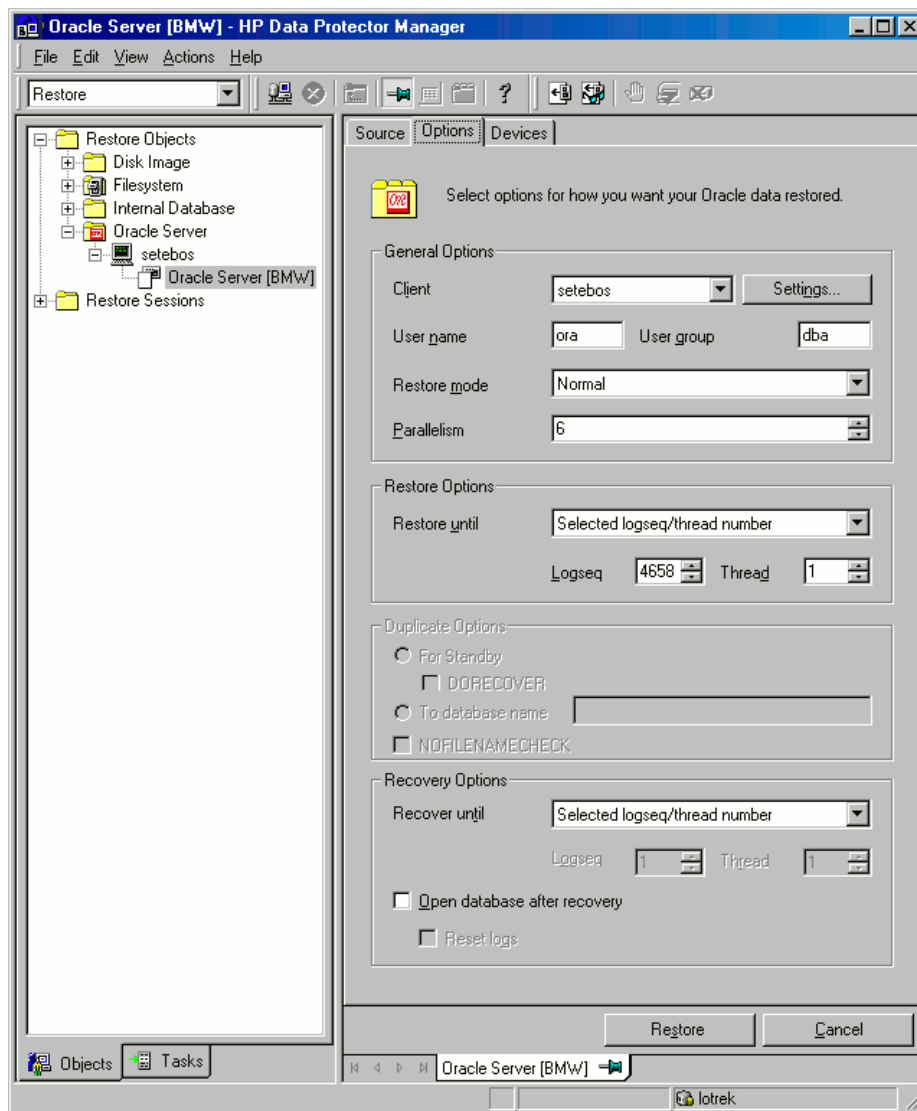
If you are restoring datafiles, you can restore the files to a new location. Right-click the database object, click **Restore As**, and in the **Restore As** dialog box, specify the new datafile location.

NOTE: When restoring to a new location, current datafiles will be switched to the restored datafile copies only if you have selected **Perform Restore and Recovery** from the **Restore action** drop-down list.

6. In the **Options** page, from the **Client** drop-down list, select the system on which the Data Protector Oracle integration agent will be started. To restore the database objects to a different database than it is selected, click **Settings** and specify the login information for the target database.


Set the other restore options. For information, see “Restore, recovery, and duplicate options” (page 71).

Figure 31 Options page



7. In the **Devices** page, select the devices to be used for the restore.
For more information on how to specify devices for a restore, see the online Help index: "restore, selecting devices for".

Figure 32 Devices page



8. Click **Restore**.

After the restore:

1. Put the database in the correct state.

If you selected **Perform Restore and Recovery** or **Perform Recovery Only** in the **Source** page, then the database is automatically put into **Open** state by Data Protector.

2. If you performed an Oracle database restore and recovery until point in time, and the session has finished successfully, reset the database to register the new incarnation of database in the recovery catalog.

Connect to the target and recovery catalog database using RMAN and reset the database:

```
rman target Target_Database_Login catalog Recovery_Catalog_Login
RMAN> RESET DATABASE;
RMAN> exit
```

3. If you did not choose to use Data Protector to recover the database objects and if you have all archived redo logs on disk, perform the following after the database is restored:

Open a command line window and enter the following commands:

```
sqlplus /nolog
SQL>recover database;
SQL>connect user/password@service as sysdba;
SQL>alter database open;
```

Restoring tablespaces and datafiles

To restore tablespaces and datafiles:

1. Open a command line window and enter the following commands if you have the database in the Open state:

```
sqlplus /nolog  
SQL>connect user/password@service as sysdba;  
SQL>alter database datafile 'datafile name' offline;  
If you are restoring a tablespace enter:  
SQL>alter tablespace tablespace_name offline;
```
2. When the restore has been completed put the datafiles and tablespaces back online with the following procedures:
Open a command line window and enter the following commands:

```
sqlplus /nolog  
SQL>connect user/password@service as sysdba  
If you are restoring a datafile enter:  
SQL>alter database datafile 'datafile_name' online;  
If you are restoring a tablespace enter:  
SQL>alter tablespace tablespace_name online;
```

Duplicating an Oracle database

Perform a production database duplication to create:

- A standby database which has the same DBID as the production (primary) database. With this, you can:
 - Create a new standby database.
 - Re-create a standby database after:
 - Loss of entire standby database
 - Primary database control file was restored or recreated
 - Database point-in-time recovery was performed on the primary database
 - Switchover or failover of database roles occurred
- An independent copy, with a unique DBID, which can be used for data mining or testing purposes.

Prerequisites

- The whole primary database with the archived logs must be backed up.
- Archive logs, which have not been backed up to tape since the last full backup and are required for duplication must be available on the duplicate system with the same path names as on the target system (system with the production database to be duplicated).
- Net service name for the auxiliary instance must be configured.
- When duplicating a database on the same system on which the target database resides, set all `*_PATH`, `*_DEST`, `DB_FILE_NAME_CONVERT`, and `LOG_FILE_NAME_CONVERT` initialization parameters appropriately. Thus, the target database files will not be overwritten by the duplicate database files.

Limitations

- Database duplication is not supported using proxy copy backups of the primary database.
- If you perform duplication of a database (not for standby) on the same system on which the target or production database resides, note that you cannot use the same database name for the target and duplicate databases when the duplicate database resides in the same Oracle home directory as the target database. Note also that if the duplicate database resides in a different Oracle home directory than the target database, then the duplicate database name has to differ from other database names in that same Oracle home directory.

To duplicate a production database:

1. On the system where the selected database will be duplicated, put the Oracle auxiliary database instance in the nomount state. See [“Changing the database state” \(page 63\)](#).
2. In the Context List of the Data Protector GUI, click **Restore**.
3. Under **Restore Objects**, expand **Oracle Server**, expand the system on which the production database resides, and then click the production database which you want to duplicate. If there are several such systems, select the system on which you want the Data Protector Oracle integration agent (`ob2rman.pl`) to be started.
4. In the **Restore Action** drop-down list, select **Perform Duplication**.
5. In the **Options** page, from the **Client** drop-down list, select the system on which the Data Protector Oracle integration agent (`ob2rman.pl`) will be started.

Click **Settings** to specify the login information (a user name, password, and net services name) for the auxiliary database. If you do not provide the login information, the duplication session will fail.

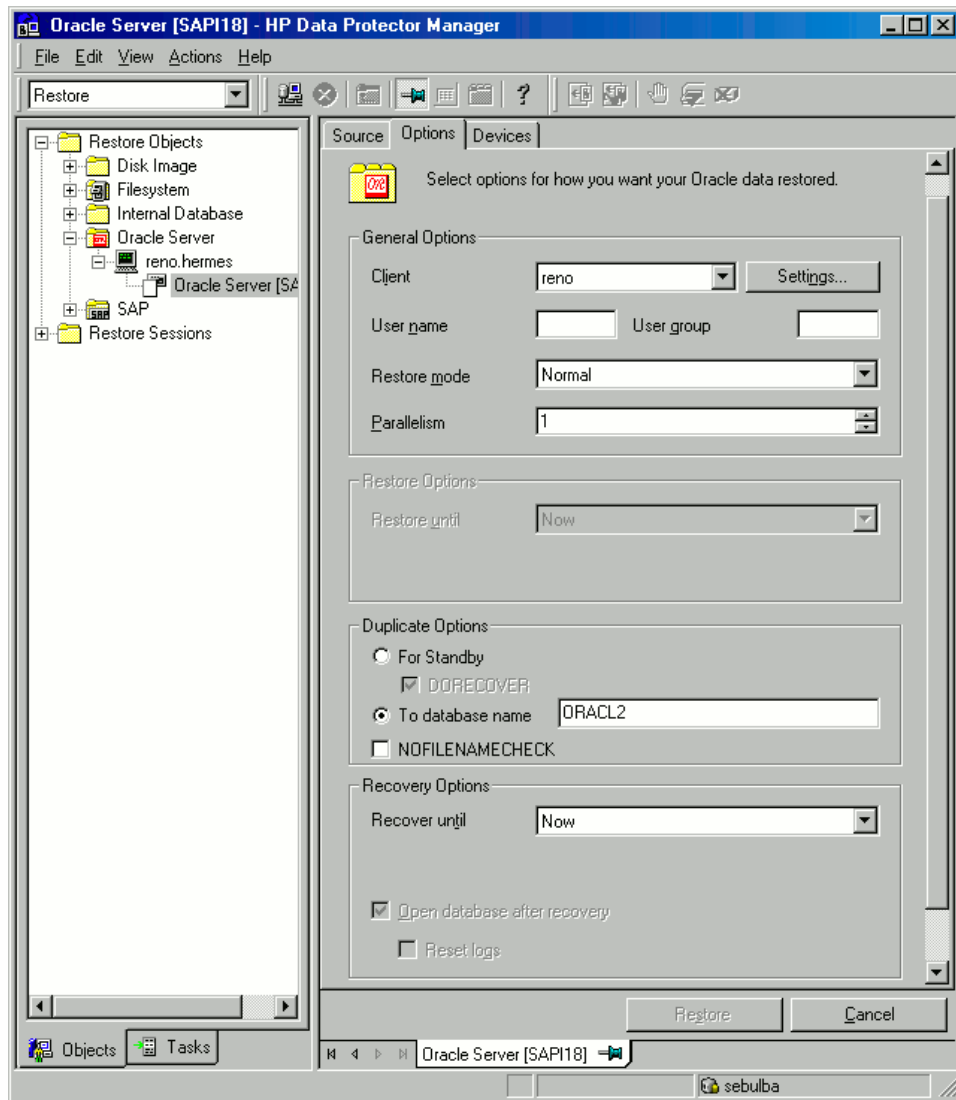
In **User name** and **User group**, specify the user name and group for the OSDBA account, which will be used by the Data Protector Oracle integration agent.

In **Parallelism**, specify the number of RMAN auxiliary channels to be allocated for database duplication.

Set duplicate options. For information, see [“Duplicate options” \(page 73\)](#) or press **F1**.

If you are creating a new database copy (not for standby), specify also the **Recover until** option to recover the duplicated database until a specified point in time.

Figure 33 Oracle duplicate options



6. Click **Restore**.

When the standby database is created, it is left mounted. Start the managed recovery process (log apply services) manually.

For information on how to use the RMAN commands to duplicate a database, see Oracle documentation.

Restore, recovery, and duplicate options

Restore action options

The following describes each of the options in the **Source** page. This page is used to define the combination of restore and recovery you would like to perform using the GUI.

In the context of Data Protector “restore” means to restore the datafiles. You can select which database, tablespace, or datafiles they would like to restore and up to which point in time they would like them to be restored. “Recover” means applying the redo logs. You can select which redo logs to apply according to SCN number, logseq, or you can apply all the redo logs to the time of the last backup.

Perform Restore

Use this option to only restore (but not recover) the database objects using Data Protector. After restore, recover the database manually using RMAN. For information on recovering the database using RMAN, see [“Restoring Oracle using RMAN” \(page 74\)](#).

Perform Restore and Recovery

Use this option to perform both the restore and recovery of the database objects using Data Protector.

Perform Recovery Only

Use this option to only recover the database objects using Data Protector.

Perform RMAN Repository Restore

Use this option to restore the recovery catalog or the control file when the database objects are not available in the **Source** page.

Perform Duplication

This option is used to perform duplication of a production database.

General options

Client

This option specifies the system on which the Data Protector Oracle integration agent (`ob2rman.pl`) will be started.

Settings

Click **Settings** to specify the login information (user name, password, and net service name) for the target database (in case of restore and recovery) or auxiliary database (in case of duplication) where you want the selected database objects to be restored or duplicated.

If this is not specified in the case of restore or recovery, the login information of the selected database that resides on the selected system will be used.

If this is not specified in the case of duplication, the duplication session will fail.

User name, User group (UNIX systems only)

Specify the operating system user account under which you want the restore to start.

Ensure that this user has Oracle rights to restore the database (for example, it is in the DBA user group). The user must also be in the Data Protector `admin` or `operator` user group (actually, the `Start restore` and `See private objects` user rights suffice).

Restore mode

This drop-down list allows you to specify which type of restore you would like perform. The options are:

- Normal
This option should be used when a conventional backup or ZDB using the backup set method was performed.
- Proxy copy
This option should be used when the original Oracle backup was made using the Oracle RMAN proxy-copy method.

This option is disabled when you perform recovery only.

Parallelism

This field is used to specify the number of concurrent data streams that can read from the backup device. The default value is one.

In case of `Normal` restore mode, to optimize restore performance, specify the same number of data streams as were used during the backup. For example, if you set the backup concurrency to 3, set the number of parallel data streams to 3 as well. Note that if a very high number of

parallel data streams is specified this may result in a resource problem because too much memory is being used.

For Oracle proxy-copy ZDB sessions, this option is disabled and Data Protector sets the number of concurrent data streams to the value that was used at backup. If you are restoring a backup created using a previous version of Data Protector, parallelism is set to the number of devices that were used for backup, regardless of the concurrency numbers for these devices.

Duplicate options

Available if **Perform Duplication** was selected.

For Standby

Select this option to create a standby database.

Default: selected.

DORECOVER

Available if **For Standby** was selected.

Select this option if you want RMAN to recover the database after creating it.

To database name

Select this option to create a new database copy. In the text box, specify its name. The name should match the name in the initialization parameter file that was used to start the auxiliary database instance. By default, the database name is set to the database name of the currently selected target database.

NOFILENAMECHECK

Select this option to disable RMAN to check whether the target datafiles share the same names with the duplicated datafiles.

Select this option when the target datafiles and duplicated datafiles have the same names, but reside on different systems.

Default: not selected.

Restore and recovery options

Restore until

The options in this drop-down list allow you to limit the selection to those backups that are suitable for an incomplete recovery to the specified time.

- **Now**

Use this option to restore the most recent full backup. By default, this option is selected.

- **Selected time**

Use this option to specify an exact time to which you wish the database to be restored. Data Protector restores the backup that can be used in recovery to the specified time.

- **Selected logseq/thread number**

A logseq number is a redo log sequence number. Use this option to specify a particular redo log sequence and a thread number which will act as an upper limit of redo logs to restore. Data Protector restores the backup that can be used in recovery to the specified log sequence number.

- **Selected SCN number**

Use this option to specify the SCN number to which you wish the database to be restored. Data Protector restores the backup that can be used in recovery to the specified SCN number.

Recover until

The options in this drop-down list allow you to specify to which point in time you would like the recovery to be performed.

- **Now**

Data Protector starts RMAN to recover the database to the most recent time possible by applying all archived redo logs. By default, this option is selected.

- **Selected time**

Use this option to specify an exact time to which the archive logs are applied.

- **Selected logseq/thread number**

A logseq number is a redo log sequence number. Use this option to specify a particular redo log sequence and a thread number which will act as an upper limit of redo logs to recover.

- **Selected SCN number**

Use this option to specify the SCN number to which you perform the recovery.

If you reset the logs, also reset the database; otherwise, Oracle will during the next backup try to use the logs that were already reset and the backup will fail. Login to the target and recovery catalog database and run:

```
rman target Target_Database_Login catalog Recovery_Catalog_Login
RMAN> RESET DATABASE;
RMAN> exit
```

Open database after recovery

Opens the database after a recovery is performed.

Reset logs

Resets the archive logs after the database is opened.

Always reset the logs:

- After an incomplete recovery (not **Recover until now**).
- If a backup of a control file is used in recovery or restore and recovery.

Do not reset the logs:

- After a complete recovery (**Recover until now**) when the backup of a control file was not used in recovery or restore and recovery.
- On the primary database, if the archive logs are used for a standby database. However, if you must reset the archive logs, you will need to recreate the standby database.

If you reset the logs when the **Recover until** option is set to **Now**, a warning is displayed, stating that you should reset the logs only if you use an older control file for restore.

NOTE: Oracle recommends that you perform a complete backup immediately after a database was opened with the **Reset Logs** option.

Restoring Oracle using RMAN

Data Protector acts as a media management software for the Oracle system, therefore RMAN can be used for a restore.

This section only describes *examples* of how you can perform a restore. The examples provided do not apply to all situations where a restore is needed.

See the *Oracle Recovery Manager User's Guide and References* for detailed information on how to perform:

- Restore and recovery of the database, tablespace, control file, and datafile.
- Duplication of a database.

The following examples of restore are given:

- [“Example of full database restore and recovery” \(page 76\)](#)
- [“Example of point-in-time restore” \(page 77\)](#)
- [“Example of tablespace restore and recovery” \(page 78\)](#)
- [“Example of datafile restore and recovery” \(page 80\)](#)
- [“Example of archive log restore” \(page 83\)](#)

The restore and recovery procedure of Oracle control files is a very delicate operation, which depends on the version of the Oracle database you are using. For detailed steps on how to perform the restore of control files, see the *Recovery Manager User's Guide and References*.

Preparing the Oracle database for restore

The restore of an Oracle database can be performed when the database is in mount mode. However, when you are performing the restore of tablespaces or datafiles, only a part of the Oracle database can be put offline.

Prerequisites

The following requirements must be met before you start a restore of an Oracle database:

- Make sure that the recovery catalog database is open. If the recovery catalog database cannot be brought online, you will probably need to restore the recovery catalog database. See [“Restore” \(page 61\)](#) for details of how to restore the recovery catalog database.
- Check which ZDB method (proxy-copy or backup set) was used for the backup session that you plan to restore.
- Control files must be available. If the control files are not available, you must restore them. See the *Oracle Recovery Manager User's Guide and References* for more details.

If you have to perform a restore of the recovery catalog database, you must perform this restore first. Only then can you perform a restore of other parts of the Oracle database.

When you are sure that the recovery catalog database files are in place, start the recovery catalog database.

- Make sure that the following environment variables are set:
 - ORACLE_BASE
 - ORACLE_HOME
 - ORACLE_TERM
 - DB_NAME
 - PATH
 - NLS_LANG
 - NLS_DATE_FORMAT

Windows systems example

```
ORACLE_BASE=Oracle_home
```

```
ORACLE_HOME=Oracle_home\product\10.1.0
```

```
ORACLE_TERM=HP
```

```
DB_NAME=PROD
PATH=$PATH:Oracle_home\product\10.1.0\bin
NLS_LANG=american
NLS_DATE_FORMAT='Mon DD YYYY HH24:MI:SS'
```

UNIX systems example

```
ORACLE_BASE=/opt/oracle
ORACLE_HOME=/opt/oracle/product/10.1.0
ORACLE_TERM=HP
DB_NAME=PROD
PATH=$PATH:/opt/oracle/product/10.1.0/bin
NLS_LANG=american
NLS_DATE_FORMAT='Mon DD YYYY HH24:MI:SS'
```

- Check that the `/etc/oratab` file has the following line:

Windows systems: `PROD:Oracle_home\product\10.1.0:N`

UNIX systems: `PROD:/opt/oracle/product/10.1.0:N`

The last letter determines whether the database will automatically start upon bootup (Y) or not (N).

Connection strings used in the examples

In the examples below, the following connection strings are used:

- Target connection string for target database:

`sys/manager@PROD`

where `sys` is the username, `manager` is the password and `PROD` is a net service name.

- Recovery catalog connection string for recovery catalog database:

`rman/rman@CATAL`

where `rman` is the username and password and `CATAL` is a net service name.

SBT_LIBRARY parameter

On Windows and UNIX systems, set the `SBT_LIBRARY` RMAN script parameter to point to the correct platform-specific Data Protector MML. The parameter must be specified for each RMAN channel separately. For details on the Data Protector MML location, see the *HP Data Protector Integration Guide for Oracle and SAP*.

In the following examples, the `SBT_LIBRARY` parameter is set to `/opt/omni/lib/libob2oracle8.so`, which is the correct path for 32-bit Solaris systems.

Example of full database restore and recovery

To perform a full database restore and recovery, you also need to restore and apply all the archive logs. To perform a full database restore and recovery:

1. Log in to the Oracle RMAN:

Windows systems: `ORACLE_HOME\bin\rman target sys/manager@PROD catalog rman/rman@CATAL`

UNIX systems: `ORACLE_HOME/bin/rman target sys/manager@PROD catalog rman/rman@CATAL`

2. Start the full database restore and recovery:

For a non-ZDB or ZDB backup set session:

```
run{
allocate channel 'dev1' type 'sbt_tape' parms
  'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
  ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME) ';
restore database;
recover database;
sql 'alter database open';
release channel 'dev1';
}
```

For a ZDB proxy-copy session:

```
run{
allocate channel 'dev1' type 'sbt_tape' parms
  'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
  ENV=(OB2BARTYPE=Oracle8,OB2PROXYCOPY=1, OB2APPNAME=DB_NAME) ';
restore database;
recover database;
sql 'alter database open';
release channel 'dev1';
}
```

You can also save the script into a file and perform a full database restore using the saved files. The procedure in such cases is as follows:

1. Create a file `restore_database` in the `/var/opt/omni/tmp` (UNIX systems) or `Data_Protector_home\tmp` (Windows systems) directory.
2. Start the full database restore:

Windows systems: `ORACLE_HOME\bin\rman target sys/manager@PROD catalog rman/rman@CATAL cmdfile=Data_Protector_home\tmp\restore_datafile`

UNIX systems: `ORACLE_HOME/bin/rman target sys/manager@PROD catalog rman/rman@CATAL cmdfile=/var/opt/omni/tmp/restore_datafile`

Example of point-in-time restore

To perform a point-in-time restore, you also need to restore and apply the archive logs to the specified point in time. To perform a point-in-time database restore and recovery:

1. Log in to the Oracle RMAN:

Windows systems: `ORACLE_HOME\bin\rman target sys/manager@PROD catalog rman/rman@CATAL`

UNIX systems: `ORACLE_HOME/bin/rman target sys/manager@PROD catalog rman/rman@CATAL`

2. Start the point-in-time restore:

For a non-ZDB or ZDB backup set session:

```
run{
allocate channel 'dev1' type 'sbt_tape' parms
  'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
  ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME) ';
set until time 'Mar 14 2004 11:40:00';
restore database;
recover database;
sql 'alter database open';
release channel 'dev1';
}
```

For a ZDB proxy-copy session, allocate one channel for restoring proxy-copy sessions and one channel for database recovery. Release the proxy-copy channel before the recovery:

```

run{
allocate channel 'dev1' type 'sbt_tape' parms
  'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
  ENV=(OB2BARTYPE=Oracle8,OB2PROXYCOPY=1, OB2APPNAME=DB_NAME)';
allocate channel 'dev2' type 'sbt_tape' parms
  'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
  ENV=(OB2BARTYPE=Oracle8, OB2APPNAME=DB_NAME)';
set until time 'Mar 14 2006 11:40:00';
restore database;
release channel 'dev1';
recover database;
sql 'alter database open';
release channel 'dev2';
}

```

3. After you have performed a point-in-time restore, reset the database in the Recovery Catalog.

You can also save the script into a file and perform a point-in-time restore using the saved files:

1. Create a file `restore_PIT` in the `/var/opt/omni/tmp` or `Data_Protector_home\tmp` directory.
2. Start the point-in-time restore:

Windows systems: `ORACLE_HOME\bin\rman target sys/manager@PROD catalog rman/rman@CATAL cmdfile=Data_Protector_home\tmp\restore_PIT`

UNIX systems: `ORACLE_HOME/bin/rman target sys/manager@PROD catalog rman/rman@CATAL cmdfile=/var/opt/omni/tmp/restore_PIT`

Example of tablespace restore and recovery

If a table is missing or corrupted, you need to perform a restore and recovery of the entire tablespace. To restore a tablespace, you may take only a part of the database offline, so that the database does not have to be in the mount mode. You can use either a recovery catalog database or control files to perform a tablespace restore and recovery. Follow the steps below:

1. Log in to the Oracle RMAN:

Windows systems: `ORACLE_HOME\bin\rman target sys/manager@PROD catalog rman/rman@CATAL`

UNIX systems: `ORACLE_HOME/bin/rman target sys/manager@PROD catalog rman/rman@CATAL`

2. Start the tablespace restore and recovery.

- If the database is in the open state, the script to restore and recover the tablespace should have the following format:

For a non-ZDB or ZDB backup set session:

```
run{
allocate channel dev1 type 'sbt_tape' parms
  'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
  ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME)';
sql 'alter tablespace TEMP offline immediate';
restore tablespace TEMP;
recover tablespace TEMP;
sql 'alter tablespace TEMP online';
release channel dev1;
}
```

For a ZDB proxy-copy session, allocate one channel for restoring proxy-copy sessions and one channel for database recovery. Release the proxy-copy channel before the recovery:

```
run{
allocate channel 'dev1' type 'sbt_tape' parms
  'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
  ENV=(OB2BARTYPE=Oracle8,OB2PROXYCOPY=1, OB2APPNAME=DB_NAME)';
allocate channel 'dev2' type 'sbt_tape' parms
  'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
  ENV=(OB2BARTYPE=Oracle8, OB2APPNAME=DB_NAME)';
sql 'alter tablespace TEMP offline immediate';
restore tablespace TEMP;
release channel 'dev1';
recover tablespace TEMP;
sql 'alter tablespace TEMP online';
release channel 'dev2';
}
```

- If the database is in the mount state, the script to restore and recover the tablespace should have the following format:

For a non-ZDB or ZDB backup set session:

```
run{
allocate channel dev1 type 'sbt_tape' parms
  'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
  ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME)';
restore tablespace 'TEMP';
recover tablespace 'TEMP';
release channel dev1;
}
```

For a ZDB proxy-copy session, allocate one channel for restoring proxy-copy sessions and one channel for database recovery. Release the proxy-copy channel before the recovery:

```
run{
allocate channel 'dev1' type 'sbt_tape' parms
  'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
  ENV=(OB2BARTYPE=Oracle8,OB2PROXYCOPY=1, OB2APPNAME=DB_NAME)';
allocate channel 'dev2' type 'sbt_tape' parms
  'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
  ENV=(OB2BARTYPE=Oracle8, OB2APPNAME=DB_NAME)';
restore tablespace 'TEMP';
release channel 'dev1';
recover tablespace 'TEMP';
release channel 'dev2';
}
```

You can also save the script into a file and perform a tablespace restore using the saved files:

1. Create a file `restore_TAB` in the `/var/opt/omni/tmp` (UNIX systems) or `Data_Protector_home\tmp` (Windows systems) directory.
2. Start the tablespace restore.

Windows systems: `ORACLE_HOME\bin\rman target sys/manager@PROD catalog rman/rman@CATAL cmdfile=Data_Protector_home\tmp\restore_TAB`

UNIX systems: `ORACLE_HOME/bin/rman target sys/manager@PROD catalog rman/rman@CATAL cmdfile=/var/opt/omni/tmp/restore_TAB`

Example of datafile restore and recovery

To restore and recover a datafile, you may take only a part of the database offline.

To restore and recover a datafile:

1. Log in to the Oracle RMAN.

Windows systems: `ORACLE_HOME\bin\rman target sys/manager@PROD catalog rman/rman@CATAL`

UNIX systems: `ORACLE_HOME/bin/rman target sys/manager@PROD catalog rman/rman@CATAL`

2. Start the datafile restore and recovery:

- If the database is in an open state, the script to restore the datafile should have the following format:

UNIX systems

For a non-ZDB or ZDB backup set session:

```
run{
allocate channel dev1 type 'sbt_tape' parms
  'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
  ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME)';
sql "alter database datafile
  '/opt/oracle/data/oradata/DATA/temp01.dbf' offline";
restore datafile
  '/opt/oracle/data/oradata/DATA/temp01.dbf';
recover datafile
  '/opt/oracle/data/oradata/DATA/temp01.dbf';
sql "alter database datafile
  '/opt/oracle/data/oradata/DATA/temp01.dbf' online";
release channel dev1;
}
```

For a ZDB proxy-copy session, allocate one channel for restoring proxy-copy sessions and one channel for database recovery. Release the proxy-copy channel before the recovery:

```
run{
allocate channel dev1 type 'sbt_tape' parms
  'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
  ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME)';
allocate channel dev2 type 'sbt_tape' parms
  'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
  ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME, OB2PROXYCOPY=1)';
sql "alter database datafile
  '/opt/oracle/data/oradata/DATA/temp01.dbf' offline";
restore datafile
  '/opt/oracle/data/oradata/DATA/temp01.dbf';
release channel dev2;
recover datafile
  '/opt/oracle/data/oradata/DATA/temp01.dbf';
sql "alter database datafile
  '/opt/oracle/data/oradata/DATA/temp01.dbf' online";
release channel dev1;
}
```

Windows systems

For a non-ZDB or ZDB backup set session:

```
run{
allocate channel dev1 type 'sbt_tape' parms
  'SBT_LIBRARY=Data_Protector_home\bin\orasbt.dll,
  ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME)';
sql "alter database datafile
  'C:\oracle\data\oradata\DATA\temp01.dbf' offline";
restore datafile
  'C:\oracle\data\oradata\DATA\temp01.dbf';
recover datafile
  'C:\oracle\data\oradata\DATA\temp01.dbf';
sql "alter database datafile
  'C:\oracle\data\oradata\DATA\temp01.dbf' online";
release channel dev1;
}
```

For a ZDB proxy-copy session, allocate one channel for restoring proxy-copy sessions and one channel for the recovery process. Release the proxy-copy channel before the recovery:

```
run{
allocate channel dev1 type 'sbt_tape' parms
  'SBT_LIBRARY=Data_Protector_home\bin\orasbt.dll,
  ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME)';
allocate channel dev2 type 'sbt_tape' parms
  'SBT_LIBRARY=Data_Protector_home\bin\orasbt.dll,
  ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME, OB2PROXYCOPY=1)';
sql "alter database datafile
  'Oracle_home\data\oradata\DATA\temp01.dbf' offline";
restore datafile
  'Oracle_home\data\oradata\DATA\temp01.dbf';
release channel dev2;
recover datafile
  'Oracle_home\data\oradata\DATA\temp01.dbf';
sql "alter database datafile
  'Oracle_home\data\oradata\DATA\temp01.dbf' online";
release channel dev1;
}
```

- If the database is in a mount state, the script to restore and recover the datafile should have the following format:

UNIX system

For a non-ZDB or ZDB backup set session:

```
run{
allocate channel dev1 type 'sbt_tape' parms
  'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
  ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME)';
restore datafile
  '/opt/oracle/data/oradata/DATA/temp01.dbf';
recover datafile
  '/opt/oracle/data/oradata/DATA/temp01.dbf';
release channel dev1;
}
```

For a ZDB proxy-copy session, allocate one channel for restoring proxy-copy sessions and one channel for the recovery process. Release the proxy-copy channel before the recovery:

```
run{
allocate channel dev1 type 'sbt_tape' parms
  'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
  ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME)';
allocate channel dev2 type 'sbt_tape' parms
  'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
  ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME, OB2PROXYCOPY=1)';
restore datafile
  '/opt/oracle/data/oradata/DATA/temp01.dbf';
release channel dev2;
recover datafile
  '/opt/oracle/data/oradata/DATA/temp01.dbf';
release channel dev1;
}
```

Windows system

For a non-ZDB or ZDB backup set session:

```
run{
allocate channel dev1 type 'sbt_tape' parms
  'SBT_LIBRARY=Data_Protector_home\bin\orasbt.dll,
  ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME) ';
restore datafile
  'Oracle_home\data\oradata\DATA\temp01.dbf';
recover datafile
  'Oracle_home\data\oradata\DATA\temp01.dbf';
release channel dev1;
}
```

For a ZDB proxy-copy session, allocate one channel for restoring proxy-copy sessions and one channel for the recovery process. Release the proxy-copy channel before the recovery:

```
run{
allocate channel dev1 type 'sbt_tape' parms
  'SBT_LIBRARY=Data_Protector_home\bin\orasbt.dll,
  ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME) ';
allocate channel dev2 type 'sbt_tape' parms
  'SBT_LIBRARY=Data_Protector_home\bin\orasbt.dll,
  ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME, OB2PROXYCOPY=1) ';
restore datafile
  'Oracle_home\data\oradata\DATA\temp01.dbf';
release channel dev2;
recover datafile
  'Oracle_home\data\oradata\DATA\temp01.dbf';
release channel dev1;
}
```

You can also save the script into a file and perform a datafile restore using the saved files:

1. Create a file `restore_dbf` in the `/var/opt/omni/tmp` (UNIX systems) or `Data_Protector_home\tmp` (Windows systems) directory.
2. Start the datafile restore:

Windows systems: `ORACLE_HOME\bin\rman target sys/manager@PROD catalog rman/rman@CATAL cmdfile=/var/opt/omni/tmp/restore_dbf`

UNIX systems: `ORACLE_HOME\bin\rman target sys/manager@PROD catalog rman/rman@CATAL cmdfile=Data_Protector_home\tmp\restore_dbf`

Example of archive log restore

To restore an archive log:

1. Log in to the Oracle RMAN:

Windows systems: `ORACLE_HOME\bin\rman target sys/manager@PROD catalog rman/rman@CATAL`

UNIX systems: `ORACLE_HOME\bin\rman target sys/manager@PROD catalog rman/rman@CATAL`

2. Start the archive log restore:

```
run{
allocate channel dev1 type 'sbt_tape' parms
  'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
  ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME) ';
restore archivelog all;
release channel dev1;}
```

You can also save the script into a file and perform an archive log restore using the saved files:

1. Create a file `restore_arch` in the `/var/opt/omni/tmp` (UNIX systems) or `Data_Protector_home\tmp` (Windows systems) directory.
2. Start the archive log restore:
Windows systems: `ORACLE_HOME\bin\rman target sys/manager@PROD catalog rman/rman@CATAL cmdfile=Data_Protector_home\tmp\restore_arch`
UNIX systems: `ORACLE_HOME/bin/rman target sys/manager@PROD catalog rman/rman@CATAL cmdfile=/var/opt/omni/tmp/restore_arch`

Example of database restore using a different device (with the automatic device selection functionality disabled)

Suppose a database was backed up with the device `dev1`. To restore the database with the device `dev2`, add the line `send device type 'sbt_tape' 'CHDEV=dev1>dev2';` to the RMAN script:

1. Log in to the Oracle RMAN:
Windows systems: `ORACLE_HOME\bin\rman target sys/manager@TIN`
UNIX systems: `ORACLE_HOME/bin/rman target sys/manager@TIN`
2. Run:

```
run {
  allocate channel 'dev_0' type 'sbt_tape'
    parms 'SBT_LIBRARY=C:/PROGRA~1/OmniBack/bin/orasbt.dll,
  ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=CAN,OB2BARLIST=test)';
  allocate channel 'dev_1' type 'sbt_tape'
    parms 'SBT_LIBRARY=C:/PROGRA~1/OmniBack/bin/orasbt.dll,
  ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=CAN,OB2BARLIST=test)';
  allocate channel 'dev_2' type 'sbt_tape'
    parms 'SBT_LIBRARY=C:/PROGRA~1/OmniBack/bin/orasbt.dll,
  ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=CAN,OB2BARLIST=test)';
  send device type 'sbt_tape' 'NO_AUTO_DEVICE_SELECTION=1';
  send device type 'sbt_tape' 'CHDEV=dev1>dev2';
  restore database;
}
```

NOTE: The line `device type 'sbt_tape' 'NO_AUTO_DEVICE_SELECTION=1';` disables the automatic device selection.

Restoring using another device

Data Protector supports the restore of Oracle database objects from devices other than those on which the database objects were backed up.

Specify these devices in the `/etc/opt/omni/server/cell/restoredev` (UNIX systems) or `Data_Protector_home\Config\server\Cell\restoredev` (Windows systems) file in the following format:

```
"DEV 1" "DEV 2"
```

where

`DEV 1` is the original device and `DEV 2` the new device.

On Windows systems, this file must be in the Unicode format.

Note that this file should be deleted after it is used.

Example

Suppose you have Oracle objects backed up on a device called `DAT1`. To restore them from a device named `DAT2`, specify the following in the `restoredev` file:

```
"DAT1" "DAT2"
```

Instant recovery and database recovery

For general information on instant recovery, see the *HP Data Protector Zero Downtime Backup Concepts Guide* and the *HP Data Protector Zero Downtime Backup Administrator's Guide*. For information on instant recovery in cluster environment (Cluster File System (CFS), MC/ServiceGuard, and Microsoft Cluster Server), see *HP Data Protector Zero Downtime Backup Administrator's Guide*.

The Data Protector instant recovery functionality is used only to restore the target volumes on which the database files are located. The database recovery part is performed after instant recovery by the RMAN utility. During database recovery, incremental backups and archive log backups performed after ZDB to disk or ZDB to disk+tape are restored from tape. Only those archive logs that do not reside on the target volumes are restored.

-
- ❗ **IMPORTANT:** If the Oracle control file, online redo logs, and SPFILE are on the same source volumes as datafiles and you enable instant recovery by setting the `omnirc` variables, note that the control file, SPFILE, and online redo logs are overwritten during the instant recovery.
-

Prerequisites

- The control file that reflects the internal database structure at the time of backup must be available on the application system. If necessary, restore the appropriate control file from a tape backup.
- The recovery catalog must be open.

Limitations

- For ZDB-to-disk sessions, only archived redo logs can be used for a database recovery after an instant recovery.
- The recovery process will fail if the log entry with the specified logseq number or SCN number was created before the target volume.

RAC preparation steps

In case of RAC, set the following variable in the `omnirc` file:

```
ZDB_IR_VGCHANGE=vgchange -a s
```

The instant recovery procedure is the same as without RAC.

However, if instant recovery is to be performed to some other node than the one that was backed up, the following procedure must be performed before the standard instant recovery procedure:

1. Make sure that the MC/SG virtual package is running on the target node.
2. Set the `OB2BARHOSTNAME` environment variable as the virtual hostname before running the configuration from the command line:

```
export OB2BARHOSTNAME=virtual_hostname
```

Instant recovery using the Data Protector GUI

To perform an instant recovery:

1. Shut down the Oracle database instance using `sqlplus`. In case of RAC, shut down all instances.

For example:

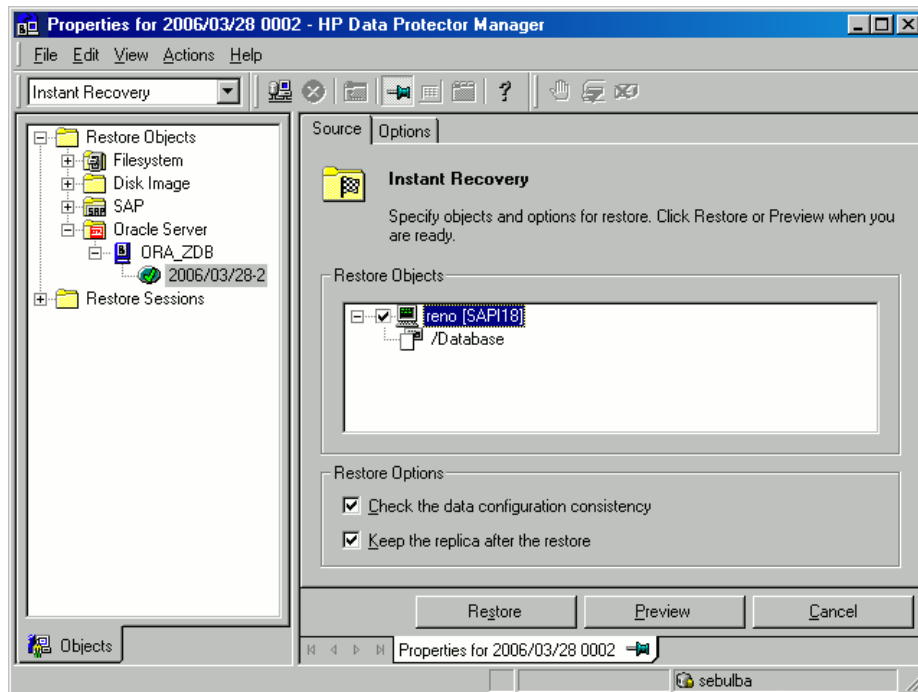
```
sqlplus
sql> shutdown immediate
sql> exit
```

2. In the Context List, click **Instant Recovery**.

3. Expand **Oracle Server** and select the ZDB-to-disk or ZDB-to-disk+tape session from which you want to perform the restore.
4. In the **Source** tab, select the objects to recover. Only whole databases can be selected. With HP P9000 XP Disk Array Family, it is recommended to leave the **Keep the replica after the restore** option set to enable a restart of an instant recovery session. With HP P6000 EVA Disk Array Family, replica is kept on the disk array only if the **Copy replica data to the source location** is selected.

Set other HP P6000 EVA Disk Array Family or HP P9000 XP Disk Array Family options. For details, press **F1**.

Figure 34 Selecting backup sessions (P9000 XP Array example)



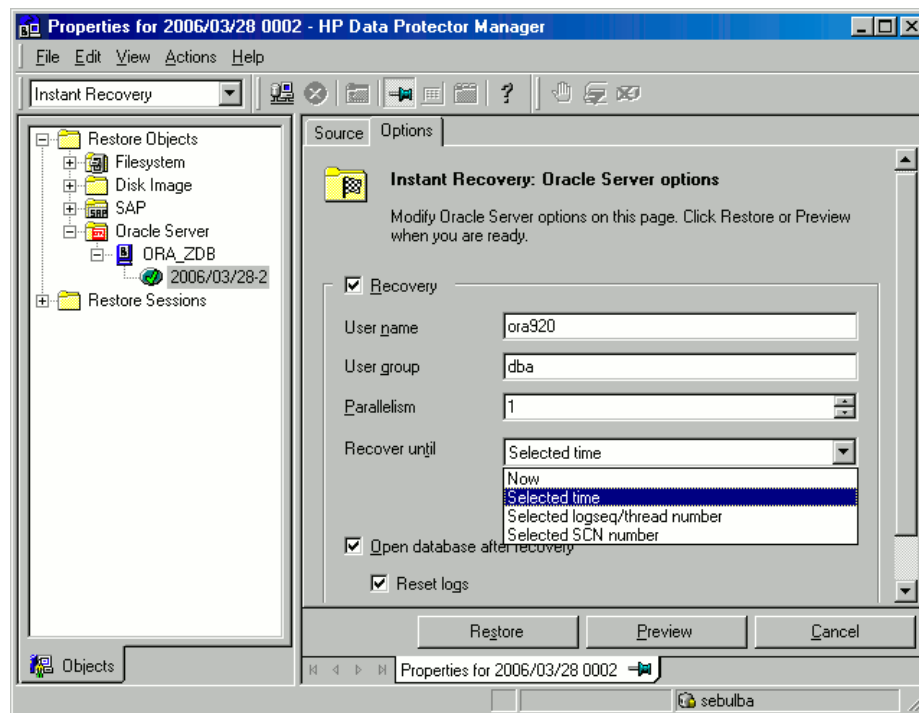
5. At this point, you can decide whether to perform a database recovery immediately after an instant recovery or not:

- To perform only an instant recovery, click **Restore**.

NOTE: You can perform a database recovery at a later time either from the Data Protector Manager Restore Context or manually using the RMAN CLI. See [“Oracle database recovery after the instant recovery”](#) (page 87).

- To perform a database recovery immediately after an instant recovery, click on the **Options** tab, select **Recovery** and then select the database recovery options. For a recovery until a selected time, logseq/thread number, or SCN number, it is recommended to reset the log files. See [“Oracle recovery options”](#) (page 87) and [“Restore, recovery, and duplicate options”](#) (page 71) for details on available options.

Figure 35 Oracle recovery options



6. Click **Restore** or **Preview**. Note that preview only checks if the replica can be restored. It does not check if the database recovery will be successful.

Data Protector recovers the database after performing instant recovery by switching the database to a mount state, restoring the necessary incremental backups and archived redo logs from tape, and applying the redo logs.

If you reset the logs, reset the database; otherwise, Oracle will during the next backup try to use the logs that were already reset and the backup will fail. Login to the target and recovery catalog database and run:

```
rman target Target_Database_Login catalog Recovery_Catalog_Login
RMAN> RESET DATABASE;
RMAN> exit
```

Oracle database recovery after the instant recovery

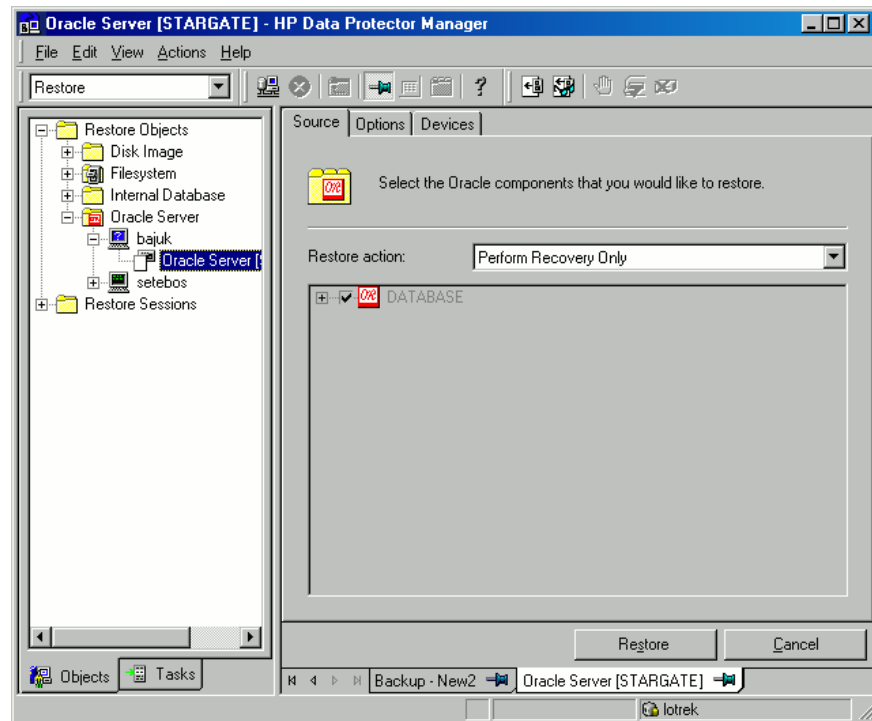
To recover the Oracle database after the instant recovery has been performed, perform the following steps:

1. Put the Oracle database in a mount state by connecting to the target database from the sqlplus and then running the following command:

```
startup mount
```

2. To recover the database, the following two options are available:
 - Perform a recovery from the Data Protector Manager Restore Context:
 - a. Expand **Oracle Server** and select the database to recover. In the **Source** tab, under **Restore action**, select **Perform recovery only**.

Figure 36 Selecting the database for recovery



- b. In the **Options** tab, select the recovery options. For details, see [“Restore, recovery, and duplicate options”](#) (page 71).
 - c. Click **Restore**.
- Perform a manual database recovery using RMAN.

Run the following RMAN script to recover the database:

```
run {
  allocate channel dev1 type 'sbt_tape' parms
    'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
    ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME)';
  recover database;
  sql 'alter database open';
  release channel dev1;
}
```

For additional examples on how to recover the database after an instant recovery, see [“Restoring Oracle using RMAN”](#) (page 74).

Oracle in Veritas Cluster instant recovery

If Oracle on the application system runs in a Veritas Cluster, the following two Veritas Cluster resources must be disabled before instant recovery is performed, and enabled after instant recovery has finished to prevent the failover of the Oracle Veritas Cluster Service Group:

- Veritas Cluster application resource for the Oracle application and
- Veritas Cluster mountpoint resource for the Oracle database files.

Follow the steps below to perform an instant recovery to the application system with Oracle in a Veritas Cluster:

1. On the application system, enter the following commands to disable the two Veritas Cluster resources:
 - a. `hares -offline application_resource_name -sys system`
 where *application_resource_name* is the name of the Veritas Cluster application resource for the Oracle application and *system* is the name of the active node.
`hares -offline mountpoint_resource_name -sys system`
 where *mountpoint_resource_name* is the name of the Veritas Cluster mountpoint resource for the Oracle database files and *system* is the name of the active node.
 - b. `hares -modify application_resource_name Enabled 0`
 where *application_resource_name* is the name of the Veritas Cluster application resource for the Oracle application.
`hares -modify mountpoint_resource_name Enabled 0`
 where *mountpoint_resource_name* is the name of the Veritas Cluster mountpoint resource for the Oracle database files.
2. Perform an instant recovery.
3. If you performed only an instant recovery without the database recovery, use RMAN as described in [“Oracle database recovery after the instant recovery” \(page 87\)](#) to bring the Oracle database to a consistent state.
4. On the application system, enter the following commands to enable the two Veritas Cluster resources:
 - a. `hares -modify mountpoint_resource_name Enabled 1`
 where *mountpoint_resource_name* is the name of the Veritas Cluster mountpoint resource for the Oracle database files.
`hares -modify application_resource_name Enabled 1`
 where *application_resource_name* is the name of the Veritas Cluster application resource for the Oracle application.
 - b. `hares -online application_resource_name -sys system`
 where *application_resource_name* is the name of the Veritas Cluster application resource for the Oracle application and *system* is the name of the active node.
`hares -online mountpoint_resource_name -sys system`
 where *mountpoint_resource_name* is the name of the Veritas Cluster mountpoint resource for the Oracle database files and *system* is the name of the active node.

Aborting sessions

You can abort currently running sessions by clicking the abort button.

If, during a session, RMAN or SQL*Plus do not respond when requested, Data Protector automatically aborts the session. By default, Data Protector waits for the response for 5 minutes. Using `omnirc` or environment variables `OB2_RMAN_COMMAND_TIMEOUT` and `OB2_SQLP_SCRIPT_TIMEOUT`, you can modify this time interval.

For details of how to set environment variables, see [“Setting environment variables” \(page 45\)](#). For details of how to set the corresponding `omnirc` options, see the online Help index: “`omnirc` option”. Note that environment variables override `omnirc` options.

Troubleshooting

This section contains a list of general checks and verifications and a list of problems you might encounter when using the Data Protector Oracle integration. You can start at [“Problems” \(page 95\)](#) and if you cannot find a solution there, perform general checks and verifications.

For general Data Protector troubleshooting information, see the *HP Data Protector Troubleshooting Guide*.

For general ZDB, restore, and instant recovery troubleshooting information, see the troubleshooting sections in the *HP Data Protector Zero Downtime Backup Administrator's Guide*.

Before you begin

- Ensure that the latest official Data Protector patches are installed. See the online Help index: “patches” on how to verify this.
- See the *HP Data Protector Product Announcements, Software Notes, and References* for general Data Protector limitations, as well as recognized issues and workarounds.
- See <http://www.hp.com/support/manuals> for an up-to-date list of supported versions, platforms, and other information.

Checks and verifications

For more detailed information about how to perform any of the following procedures, see the Oracle documentation.

If your configuration, backup, or restore failed:

- On the application system, verify that you can access the Oracle target database and that it is opened as follows:
 1. **UNIX systems:** Export the `ORACLE_HOME` and `DB_NAME` variables as follows:
 - if you are using an sh - like shell, enter the following commands:

```
ORACLE_HOME="ORACLE_HOME"
export ORACLE_HOME
DB_NAME="DB_NAME"
export DB_NAME
```
 - if you are using a csh - like shell, enter the following commands:

```
setenv ORACLE_HOME "ORACLE_HOME"
setenv DB_NAME "DB_NAME"
```
 2. **Windows systems:** Set the `ORACLE_HOME` and `DB_NAME` variables.
 2. Start SQL*Plus from the bin directory in the `ORACLE_HOME` directory:

```
sqlplus /nolog
```
 3. Start SQL*Plus and type:

```
connect user_name/password@service as sysdba;
select * from dba_tablespace;
exit
```

If this fails, open the Oracle target database.
- On the application system, verify that you can access the recovery catalog (if used) and that it is opened as follows:
 1. Export or set the `ORACLE_HOME` and `DB_NAME` variables as described in [Step 1](#).
 2. Start SQL*Plus from the bin directory in the `ORACLE_HOME` directory:

```
sqlplus /nolog
```

3. Start SQL*Plus and type:


```
connect Recovery_Catalog_Login
select * from rcver;
exit
```

If this fails, open the recovery catalog.

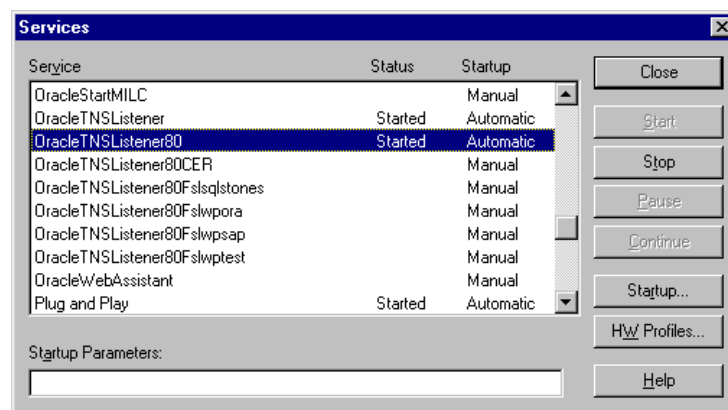
- Verify that the listener is correctly configured for the Oracle target database and the recovery catalog database. This is required to properly establish network connections:
 1. Export or set the `ORACLE_HOME` variable as described in [Step 1](#).
 2. Start the listener from the `bin` directory in the `ORACLE_HOME` directory:

```
lsnrctl status service
```

If this fails, startup the listener process and see the Oracle documentation for instructions on how to create a configuration file (`LISTENER.ORA`).

On Windows, the listener process can be started in the Control Panel > Administrative Tools > Services.

Figure 37 Checking the status of the Oracle listener



The status of the respective listener service in the **Services** window should be **started**, otherwise you must start it manually.

3. Start SQL*Plus from the `bin` directory in the `ORACLE_HOME` directory:


```
sqlplus /nolog
```
4. Start SQL*Plus and type:

```
connect Target_Database_Login
exit
and then
connect Recovery_Catalog_Login
exit
```

If this fails, see the Oracle documentation for instructions on how to create a configuration file (`NAMES.ORA`).

- From the application system, verify that the target database and recovery catalog database are configured to allow remote connections with the system privileges and to allow backup:
 - If you use the recovery catalog database:

Export or set the `ORACLE_HOME` and `DB_NAME` variables as described in [Step 1](#).

```
SQL> connect login_to_recovery_catalog_or_target_database as sysdba;
```

```
> exit
```

```
ORACLE_HOME/bin/rman target login_to_target database catalog  
login_to_Recovery_Catalog
```

- If you do not use the recovery catalog database:

Export or set the `ORACLE_HOME` and `DB_NAME` variables as described in [Step 1](#).

```
ORACLE_HOME/bin/rman target login_to_target_database nocatalog
```

See the Oracle documentation for how to set up a password file and parameters in the `initDB_NAME.ora` file and how to add system privileges for a user.

For information, see the section “Recovery Manager Connection Options” in the *Oracle Backup and Recovery Guide*.

- If you use the recovery catalog database, verify that the target database is registered in the recovery catalog:

1. Export or set the `ORACLE_HOME` variable as described in [Step 1](#).

2. Start SQL*Plus from the `bin` directory in the `ORACLE_HOME`; directory:

```
sqlplus /nolog
```

3. Start SQL*Plus and type:

```
connect Recovery_Catalog_Login;
```

```
select * from rc_database;
```

```
exit
```

If this fails, start the configuration using Data Protector on the application system, or see the Oracle documentation for information on how to register an Oracle target database in the recovery catalog database.

- On the application system, verify backup and restore directly to disk using an RMAN channel type disk:

If you use the recovery catalog:

1. Export or set the `ORACLE_HOME` variable as described in [Step 1](#).

2. Start RMAN from the `bin` directory in the `ORACLE_HOME` directory:

```
rman target Target_Database_Login catalog Recovery_Catalog_Login
```

If you do not use the recovery catalog:

1. Export or set the `ORACLE_HOME` variable as described in [Step 1](#).

2. Start RMAN from the `bin` directory in the `ORACLE_HOME` directory:

```
rman target Target_Database_Login nocatalog
```

An example of the RMAN backup script is presented below:

```
run {  
  allocate channel 'dev0' type disk;  
  backup tablespace tablespace_name format  
    'ORACLE_HOME/tmp/datafile_name';  
}
```

After a successful backup, try to restore the backed up tablespace by running the following restore script:

```
run {  
  allocate channel 'dev0' type disk;  
  sql 'alter tablespace tablespace_name offline immediate';  
  restore tablespace tablespace_name;  
  recover tablespace tablespace_name;  
  sql 'alter tablespace tablespace_name online';  
}
```

If this fails, see the Oracle documentation for details of how to execute a backup and restore directly to disk using RMAN.

Additionally, if your configuration or backup failed:

- Verify that the Data Protector software has been installed properly.
For details, see the *HP Data Protector Installation and Licensing Guide*.
- Check if the `SYSDBA` privilege is granted to the Oracle administrator.
- If you have special Oracle environment settings, ensure that they are entered in the Data Protector Oracle configuration files on the Cell Manager. For information on setting the variables in the Data Protector Oracle configuration files, see the `util_cmd` man page or the *HP Data Protector Command Line Interface Reference*.
- Perform a filesystem backup (non-ZDB) of the Oracle Server system so that you can eliminate any potential communication problems between the Oracle Server and the Data Protector Cell Manager system.
For details about how to do a filesystem backup, see the online Help index: "standard backup procedure".
Ensure that the hostname defined in the backup specification as a system to be backed up is the name of the application system.
- On Windows systems, check the Data Protector Inet service startup parameters on the Oracle Server system:
Go to **Control Panel > Administrative Tools > Services > Data Protector Inet**.
The service must run under a specified user account. Make sure that the same user is also added to the Data Protector `admin` or `user` group.
- Examine the system errors reported in the following file on the application system (Oracle proxy-copy ZDB method) or backup system (Oracle backup set ZDB method):
UNIX systems: `/var/opt/omni/log/debug.log`
Windows systems: `Data_Protector_home\log\debug.log`

Additionally, if your backup or restore failed:

- Test the Data Protector internal data transfer using the `testbar2` utility:
 1. Verify that the Cell Manager name is correctly defined on the Oracle Server system. Check the following file, which contains the name of the Cell Manager system:
UNIX systems: `/etc/opt/omni/client/cell_server`
Windows systems: `Data_Protector_home\Config\client\cell_server`
 2. From the `bin` directory in the `ORACLE_HOME` directory, run:
If backup failed:

```
testbar2 -type:Oracle8 -appname:DB_NAME-perform:backup  
-bar:backup_specification_name
```


If restore failed:

```
testbar2 -type:Oracle8 -appname:DB_NAME-perform:restore
```


The hostname should not be specified in the `object` option. It is automatically provided by `testbar2`.
 3. You should see only `NORMAL` messages displayed on your screen, otherwise examine the errors reported by the `testbar2` utility by clicking the **Details** button in the Data Protector **Monitor** context.

If the messages indicate problems on the Data Protector side of the integration, proceed as follows:

- Check if the user under which the backup or restore session was started has appropriate Oracle permissions (for example, belongs to the `DBA` group). This user must also be in the Data Protector operator or admin user group.
- Check that the respective Data Protector user group has the `See private objects` user right enabled.
- **If backup failed:**
Create an Oracle backup specification to back up to a null device or file. If the backup succeeds, the problem may be related to the backup devices. See the *HP Data Protector Troubleshooting Guide* for instructions on troubleshooting devices.
- **If restore failed:**
Run the `omnidb` command to see objects in the database.

If the test fails again, call a support representative for assistance.

Additionally, if your restore failed:

- Verify that an object exists on the backup media.
This can be done by running the following command on the Oracle server system from the `bin` directory in the `ORACLE_HOME` directory:

```
omnidb -oracle8 "object_name" -session "Session_ID" -media
```


The output of the command lists detailed information about the specified Oracle object, as well as the session IDs of the backup sessions containing this object and a list of the media used. For detailed syntax of the `omnidb` command, see its man page.
- Ensure that the database is in the correct state.
If you are trying to restore a database item using the Data Protector GUI and the GUI stops responding, try one of the following:
 - If you are restoring the control file, the database should be in the `NoMount` state.
Open a command window and enter the following:

```
sqlplus/nolog
SQL>connect user/password@service as sysdba
SQL>shutdown immediate
SQL>startup nomount
```

- If you are restoring datafiles, the database should be in the Mount state. Open a command window and enter the following:

```
sqlplus/nolog
SQL>connect user/password@service as sysdba
SQL>shutdown immediate
SQL>startup mount
```

- If there is a problem you cannot resolve while you are trying to restore a database item using the Data Protector GUI, try using the RMAN CLI to restore the database items.

For information, see [“Restoring Oracle using RMAN” \(page 74\)](#).

- Try putting the database into the Open state manually after using the Data Protector GUI to recover and restore a backup session.

If you have used the Data Protector GUI to recover and restore a backup session and you see the following error message:

Oracle Error: ORA-1589: must use RESETLOGS or NORESETLOGS option for database open.

Open a SQLplus window and use the following command:

```
sqlplus/nolog
SQL>connect user/password@service as sysdba
SQL>alter database open noresetlogs;
```

If this does not work, try using the following command:

```
SQL>alter database open resetlogs;
```

Problems

Problem

SQL*Plus is unable to connect to destination

Action

Check if the Oracle listener process is up and running. Check if there are any environment variables you need to enter (for example, TNS_ADMIN). Enter these variables in the Data Protector Oracle configuration files on the Cell Manager. For information, see the `util_cmd` man page.

Problem

The following error is displayed: ORA-12532: : invalid argument

If this is reported by SQL*Plus in the Data Protector monitor, the application system may be low on resources (CPU, memory, and so on).

Action

Try to configure the application system in such a way that it consumes as little resources as possible. This error can be reproduced without using Data Protector by starting SQL*Plus on the application system, and connecting to the target database on the application system.

Problem

Backup set ZDB is aborted after 10 minutes

While performing a backup set ZDB, the following warning is displayed for each database datafile:

```
RMAN-06554: WARNING: file n is in backup mode
```

The ZDB session then aborts with the following message:

```
Bar backup session was started but no client connected in 600 seconds.
```

Action

Increase the value of the following variables in the global options file (by default, these variables are set to 10):

- If you upgraded Data Protector from a previous version of Data Protector:

```
SmWaitForFirstClient=minutes
```

- If you performed a clean installation:

```
SmWaitForFirstBackupClient=minutes
```

See the *HP Data Protector Troubleshooting Guide* for more information on the global options file.

Problem

Backup set ZDB fails after changing the physical schema of the Database

Backup fails after you have modified the physical schema of a database, for example, if you added or dropped a tablespace, added a new datafile, or added or dropped a rollback segment.

Depending on the performed modification, different error messages are displayed, for example:

```
RMAN-06056: could not access datafile datafile
```

The problem occurs because the physical schema of target database is not updated in the recovery catalog.

Action

Manually re-synchronize the recovery catalog database with the current control file.

Problem

On UNIX systems, a backup set ZDB-to-disk+tape session fails

While performing a backup set ZDB to disk+tape, the session fails with the following error when Oracle Server attempts to start an instance on the backup system:

```
[Major] From: ob2rman@computer.company.com DB_NAME Time: Date Time
```

The database reported error while performing requested operation.

This problem occurs when either the user ID or the group ID number of the Oracle operating system user account on the application and backup systems do not match. Under such circumstances, Oracle Server is unable to start the instance on the backup system due to missing privileges.

Action

Configure the user accounts as described in [“Configuring Oracle operating system user accounts” \(page 37\)](#), and restart the session.

Problem

Proxy copy restore fails

Proxy copy restore fails with the following error:

```
RMAN-10035: exception raised in RPC: ORA-27197: skgfprs: sbtpcrestore  
returned error
```


RMAN-10031: ORA-27197 occurred during call to
DBMS_BACKUP_RESTORE.PROXYRESTOREDATAFILE

Action

Check the IDB for the session and the objects of the latest backup. You might check if a more recent session exists in the recovery catalog. Connect to the RMAN prompt:

```
rman target user/password@TGT_DB catalog user/password@CDB
```

At the RMAN> prompt, enter

```
list backup;
```

to display a list of the objects in the recovery catalog. Check the list of Proxy Copy sessions, listed at the end.

To synchronize the recovery catalog and the IDB, run the RMAN command:

```
resync catalog;
```


After the synchronization is performed, restore should be possible.

Problem

Restore after a switch between ZDB methods fails

If you perform a restore to a specified time ($T_RESTORE$) that lies in the time interval between the start of the first backup of the entire database using the new method (T_START), and before this backup is finished (T_FINISH), RMAN may try to restore the backup files made with the new method using a channel allocated for backup files made using the previous method. As a result, the restore procedure fails.

Figure 38 Restore after a switch between ZDB methods fails



Action

Restore the backup session manually using RMAN scripts. Add the required parameter to the allocated channels, that is `OB2PROXYCOPY=1` for the channel which will be used for restoring the backup made using the proxy-copy ZDB method. Then restore the backup files using the correct channels.

For example, if you switched from the backup set to the proxy-copy ZDB method, the script may look similar to the following one:

```
run {  
  ALLOCATE CHANNEL 'dev_0' TYPE 'sbt_tape'  
  PARMS 'SBT_LIBRARY=Path_to_Data_Protector_MML,  
        ENV(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME)';  
  ALLOCATE CHANNEL 'dev_1' TYPE 'sbt_tape'  
  PARMS 'SBT_LIBRARY=Path_to_Data_Protector_MML,  
        ENV(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME,  
        OB2PROXYCOPY=1)';  
  RESTORE DATAFILE list_of_backup_set_backups UNTIL  
    T RESTORE CHANNEL 'dev_0';  
  RESTORE DATAFILE list_of_proxy-copy_backups UNTIL  
    T RESTORE CHANNEL 'dev_1';  
  RELEASE 'dev_0';  
}
```

```
RECOVER DATABASE UNTIL T_RECOVER ...
RELEASE 'dev_1';
}
```

Where:

T_RESTORE specifies the time to which to restore and *T_RECOVER* the time to which to apply the transactions.

list_of_backup_set_backups is a list of backups of the entire database using the backup set ZDB method.

list_of_proxy_copy_backups is a list of datafile backups completed after the start of the backup of the entire database (*T_START*) and before *T_RESTORE*.

Problem

Data Protector reports errors when calling SYS.LT_EXPORT_PKG.schema_inf_exp during Oracle backup

The following errors are listed in the Data Protector monitor:

```
EXP-00008: ORACLE error 6550 encountered
ORA-06550: line 1, column 13:
PLS-00201: identifier 'SYS.LT_EXPORT_PKG' must be declared
ORA-06550: line 1, column 7:
PL/SQL: Statement ignored
EXP-00083: The previous problem occurred when calling
SYS.LT_EXPORT_PKG.schema_info_exp
. exporting statistics
Export terminated successfully with warnings.
[Major] From: ob2rman.pl@machine "MAKI" Time: 10/01/01 16:07:53
Export of the Recovery Catalog Database failed.
```

Action

Start SQL*Plus and grant the execute permission to the LT_EXPORT_PKG as follows (make sure that the user sys has the SYSDBA privilege granted beforehand):

```
sqlplus 'sys/password@CDB as sysdba'
```

```
SQL> grant execute on sys.lt_export_pkg to public;
```

Restart the failed backup session.

Problem

On a UNIX system, Data Protector reports "Cannot allocate/attach shared memory"

Backup fails and the following error message is displayed:

```
Cannot allocate/attach shared
memory (IPC Cannot Allocate Shared Memory Segment)
System error: [13] Permission denied) => aborting
```

Action

Set the OB2SHMEM_IPCGLOBAL omnirc variable in the /opt/omni/.omnirc file to 1 to use the memory windowing properly, and restart the failed backup session. See the *HP Data Protector Troubleshooting Guide* for details on using the omnirc file.

Problem

Backup fails after a point in time restore and recovery

The following error is displayed:

```
RMAN-06004: ORACLE error from recovery catalog database: RMAN-20003:
target database incarnation not found in recovery catalog
```

Action

Connect to the target and recovery catalog database using RMAN and reset the database to register the new incarnation of database in the recovery catalog:

```
rman target Target_Database_Login catalog Recovery_Catalog_Login
RMAN> RESET DATABASE;
RMAN> exit
```

Problem

Oracle online backup fails with the following error:

RMAN-06004: ORACLE error from recovery catalog database: RMAN-20220: controlfile copy not found in the recovery catalog

When running an online backup, Data Protector adds the filename of the *controlfilecopy* to the RMAN backup script. This filename has to be cataloged to the RMAN catalog prior to the backup command.

Action

To catalog the *controlfilecopy* to the RMAN catalog:

1. Connect to RMAN on the application system.
2. Run the following command:

```
RMAN> catalog controlfilecopy 'CONTROL_FILE_LOCATION/ctrlDB_NAME.ctl'
```

Problem

Backup of archive logs on RAC cannot be performed

On RAC, the archive logs are not installed on a NFS mounted disk. Backup of archive logs cannot be performed.

Action

Edit the archive logs backup specification:

- Add an additional `allocate channel` command for each node.
- Add a command to connect to each instance. The connection parameters should be given as *username/passwd@INSTANCE*.

For example, if you are using two nodes, the backup specification might look as follows:

```
run {
allocate channel 'dev_0' type 'sbt_tape' parms
'SBT_LIBRARY=Path_to_Data_Protector_MML,
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME,OB2BARLIST=RAC_arch) '
connect username/passwd@INSTANCE_1;
allocate channel 'dev_2' type 'sbt_tape' parms
'SBT_LIBRARY=Path_to_Data_Protector_MML,
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME,OB2BARLIST=RAC_arch) '
connect username/passwd@INSTANCE_2;
backup
format 'RAC_arch<QU_%s:%t:%p>.dbf'
archivelog all;
}
```

Problem

The Recovery Catalog was lost and the control file cannot be restored from Data Protector managed backup

The Recovery Catalog was not used, the RMAN autobackup feature was not used, and the control file cannot be restored from Data Protector managed backup. A valid control file backup exists on tape.

Action

- Restore the control file from RMAN backup set, mount and restore the database, and perform database recovery:

```
run {
  allocate channel 'dev_0' type 'sbt_tape' parms
    'SBT_LIBRARY=Path_to_Data_Protector_MML';
  restore controlfile from 'backup piece handle';
  sql 'alter database mount';
  set until time 'MMM DD YY HH24:MM:SS';
  restore database;
  recover database;
  sql 'alter database open resetlogs';
  release channel 'dev_0';
}
```

At this point you must manually register any backups made after the control file backup that was restored. After that, continue with the restore procedure.

For the *backup piece handle* search the Data Protector internal database and session outputs of previous backup sessions.

Problem

How to modify the RMAN restore script

When you start a restore of an Oracle database using the Data Protector GUI or CLI, an RMAN restore script is created, which is instantly run, so you cannot edit it first.

Action

To edit the script before it is run, set the Data Protector `omnirc` variable `OB2RMANSAVE` to point to an existing directory. When the variable is set and you start a restore, the RMAN restore script, which is created at run time, is saved to the specified location under the name `RMAN_restore_backup_specification_name.rman`, and the actual restore is skipped. Then you can edit the script and run it manually afterwards. On how to set the `omnirc` variable, see the online Help index: "omnirc options".

To start a restore using Data Protector again, clear the `OB2RMANSAVE` variable by deleting its content or commenting or removing the whole variable. If you comment or remove the variable on a Windows system, restart the Data Protector `Inet` service for the settings to take effect.

Problem

Instant recovery session for an Oracle database fails

On a Windows Server 2008 x64 system, when you run an instant recovery from a backed up Oracle 11g application database that resides on a disk array of the HP P6000 EVA Disk Array Family, the session may fail with the following error messages:

```
[Major] From: SMISA@appsystem.company.com "SMISA"
Time: 5/17/2010 5:00:50 AM
A filesystem could not be mounted.
Filesystem name :
Mount point : I:\

[Critical] From: SMISA@appsystem.company.com "SMISA"
Time: 5/17/2010 5:00:50 AM
Failed to resume the application system.
```

```
[Critical] From: SMISA@appsystem.company.com "SMISA"  
Time: 5/17/2010 5:00:50 AM  
Instant Recovery failed.
```

Such an instant recovery failure occurs when the “copy-back” instant recovery method (the default) is chosen and the following instant recovery options are selected in the Source and Options panes of the Data Protector GUI:

- Wait for the replica to complete (with the default waiting period used)
- Retain source for forensics
- Check the data configuration consistency
- Force the removal of all replica presentations
- Recovery
- Open database after recovery

Action

On the application system and the backup system, set the values of the `omnirc` variables `ZDB_DELAY_BEFORE_RESCAN` and `ZDB_DELAY_AFTER_RESCAN` to 300 and restart the instant recovery session.

Problem

Instant recovery of an Oracle database fails

An instant recovery session for an Oracle database fails with a message similar to the following:

```
[Normal] From: ob2rman@x64-node1.x64ring.com "testdb" Time:  
2/7/2008 10:48:19 AM  
Starting target database instant recovery.
```

```
Net service name: testdb.  
Instance status: .  
Instance name: .  
Database DBID = .  
Database control file type: .  
Database log mode: .
```

```
[Major] From: ob2rman@x64-node1.x64ring.com "testdb" Time:  
2/7/2008 10:48:20 AM  
The database reported error while performing requested operation.
```

Note that the database parameters are empty, which happens when Data Protector is not able to connect to the Oracle database.

Action

Data Protector must be able to connect to the Oracle database even when the database is in the Mount or Nomount state. One way of achieving this is by configuring static service information for your Oracle listener. For details, see the Oracle documentation.

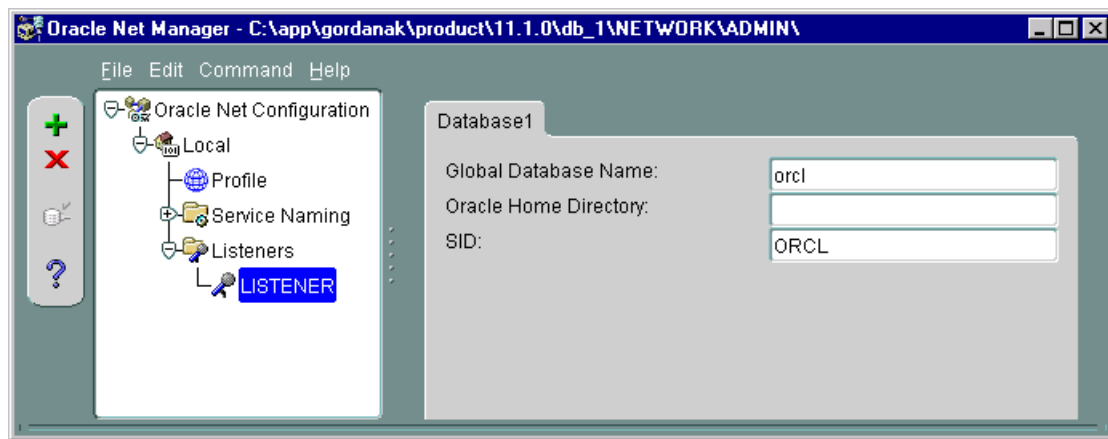
The following example shows how static service information is configured using the Oracle Net Manager.

Suppose you have the following environment:

```
Listener name: LISTENER  
Global database name: orcl  
Oracle SID: ORCL
```

To configure static service information for the listener, open the Oracle Net Manager, select the listener, go to the **Database Services** context, add a database and specify the Oracle database parameters.

Figure 39 Configuring an Oracle listener



As a result, the listener.ora file gets updated with the SID_LIST_LISTENER section.

```
SID_LIST_LISTENER =
  (SID_LIST =
    (SID_DESC =
      (GLOBAL_DBNAME = orcl)
      (SID_NAME = ORCL)
    )
  )
```

At the end, restart the Oracle listener service to apply the changes.

Problem

“Binary util_orarest is missing” error is displayed when browsing Oracle restore sessions

The following error message is displayed when browsing Oracle database for restore sessions in the Data Protector GUI Restore context:

Binary util_orarest is missing. Cannot get information from the remote host.

The problem can appear in the following cases:

- When restoring database items to a new host.
- When restoring 64-bit Oracle version 10.2.0.4 in the RAC environment, on HP-UX 11.23 PA-RISC systems. If util_orarest is present on the system, this error can mean that the util_orarest agent ends abnormally while trying to load the 32-bit OCI library from the ORACLE_HOME/lib32 directory.

Actions

- When restoring database items to a new host, resolve the problem as follows:
 1. Close Data Protector.
 2. Set the environment variable on the system where the Cell Manager resides:
`OB2_ORARESTHOSTNAME = target_Oracle_host`
 3. Restart Data Protector and try to restore the database items again.
 4. When the restore is complete, close Data Protector and re-set the following environment variable:
`OB2_ORARESTHOSTNAME = empty`

5. Restart Data Protector.

- When restoring 64-bit Oracle database in RAC environment, on HP-UX 11.23, resolve the problem as follows:
In the directory `ORACLE_HOME/lib` remove the soft link `libclntsh.sl`, which points to the 64-bit OCI library `ORACLE_HOME/lib/liblntsh.sl.10.1`.

2 Data Protector SAP R/3 ZDB integration

Introduction

This chapter explains how to configure and use the Data Protector SAP R/3 ZDB integration (**SAP R/3 ZDB integration**). It describes concepts and methods you need to understand to back up and restore the following files of the SAP R/3 database environment (**SAP R/3 objects**):

- data files
- control files
- online redo logs
- offline (archived) redo logs
- SAP R/3 logs and parameter files

Data Protector supports offline and online backups. During an online backup, the SAP R/3 application is actively used.

Data Protector offers interactive and scheduled backups of the following types:

- ZDB to disk
- ZDB to tape
- ZDB to disk+tape

Data Protector supports only a filesystem restore. You can restore SAP R/3 files:

- To the original location
- To another client
- To another directory

“[Backup and restore sessions](#)” (page 104) shows which restore methods are available, depending on the ZDB session you restore from.

Table 8 Backup and restore sessions

ZDB type	Restore methods
ZDB to tape	Standard restore
ZDB to disk	Instant recovery
ZDB to disk+tape	Standard restore, instant recovery

When the instant recovery completes, you can recover the database to a specific point in time using the SAP BRTOOLS interface.

Table 9 Supported disk arrays and array configurations

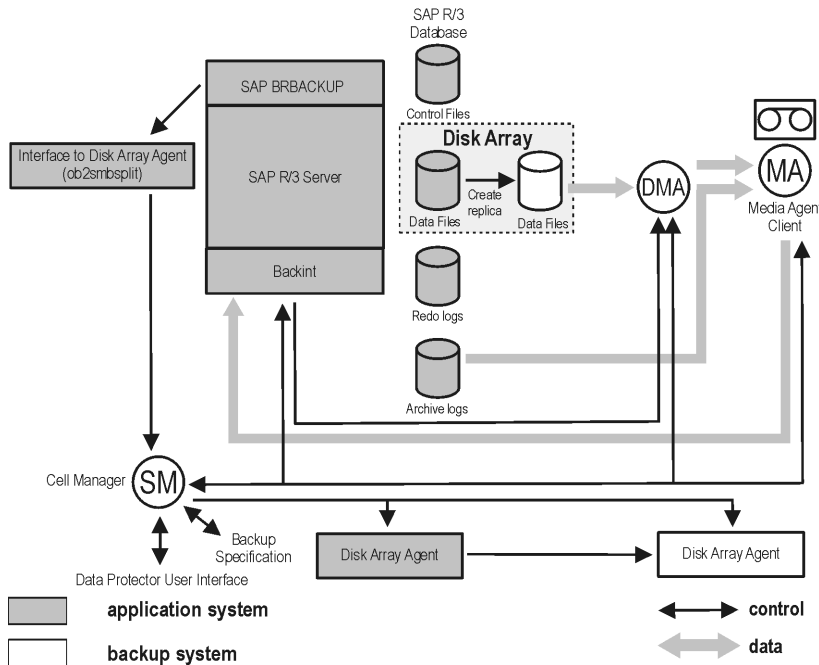
Supported array	Supported configurations
EMC Symmetrix (EMC)	TimeFinder, SRDF, combined SRDF+TimeFinder
HP P9000 XP Disk Array Family (P9000 XP Array)	HP BC P9000 XP, HP CA P9000 XP, combined HP CA+BC P9000 XP
HP P6000 EVA Disk Array Family (P6000 EVA Array)	HP BC P6000 EVA, combined HP CA+BC P6000 EVA

This chapter provides information specific to the Data Protector SAP R/3 ZDB integration. For general Data Protector procedures and options, see the *online Help*. For details on ZDB terminology, ZDB types, advantages of offline and online backups, and instant recovery concepts, see the *HP Data Protector Zero Downtime Backup Concepts Guide*.

Integration concepts

“SAP R/3 integration architecture” (page 105) shows the architecture of the Data Protector SAP R/3 ZDB integration. The figure illustrates the preferred configuration, in which the Oracle control file, online redo log files, and Oracle SPFILE reside on a different volume group than the Oracle data files.

Figure 40 SAP R/3 integration architecture

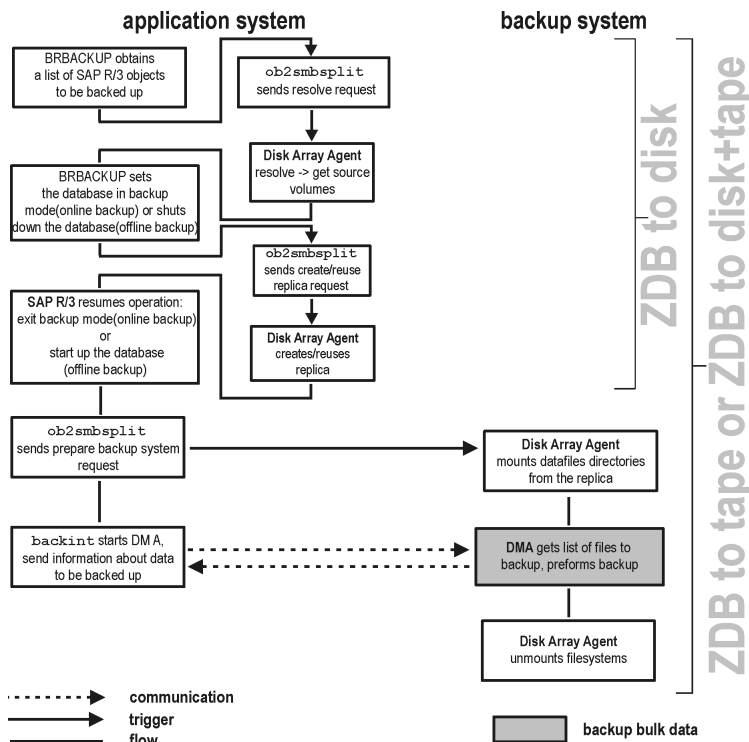


For other supported configurations, see “ZDB integrations omnirc variables” (page 271).

ZDB flow

For details of how ZDB options affect the ZDB flow, including mounting, activating volume or disk groups and so on, see the *HP Data Protector Zero Downtime Backup Administrator's Guide*.

Figure 41 SAP R/3 ZDB session flow (BRBACKUP started on the application system)



1. The SAP R/3 backup specification is read and the Data Protector `omnisap.exe` program is started on the application system.
2. `omnisap.exe` starts BRBACKUP, which switches the database to the backup mode (online backup) or shuts down the database (offline backup), and starts the split command on the application system, with the list of files to be included in the replica creation.
The database is put out of the backup mode (online backup) or is restarted (offline backup) after the replica is created.

NOTE: Data Protector can use the `splitint` interface (if BRTOOLS supports it) to reduce the time during which the database is in the backup mode.

3. The Data Protector `ob2smbsplit` command resolves the backup configuration, creates a replica, and preparing the replica for backup.
The mountpoints for the backed up object are created on the backup system.
The backup volume/disk groups are activated and the filesystems are mounted on the backup system.
4. **ZDB to disk only:** The remaining ZDB options are processed and the details on the session are written to the ZDB database. The session then finishes.
5. **ZDB to tape, ZDB to disk+tape:** BRBACKUP starts the Data Protector `backint` program, which starts establishing a connection between the Data Protector Data Movement Agents (DMA) on the backup system and the General Media Agents (MA). The process is coordinated by the Data Protector Backup Session Manager (BSM). When the connection is established, the data specified for backup is streamed to tape.

NOTE: You can configure SAP to use a third-party `backint` to perform ZDB to disk+tape backup. To enable this feature, set the `OB2_3RD_PARTY_BACKINT` environment variable to 1 and copy the `backint` you want to use to the SAP BRTOOLS directory. In case of the ZDB to disk+tape backup, a `disk_only` backup object is created in the Data Protector IDB. The third-party `backint` starts after the split and is responsible for backing up the needed files.

6. When the data transfer completes, the backup system is disabled (filesystems are unmounted on all platforms and volume/disk groups deactivated on UNIX).
7. **EMC and P9000 XP Array only:** Links are re-established, depending on how ZDB options are specified.

Data Protector SAP R/3 configuration file

Data Protector stores the integration parameters for every configured SAP R/3 database in the following file on the Cell Manager:

- On UNIX: `/etc/opt/omni/server/integ/config/SAP/client_name%ORACLE_SID`
- On Windows:
`Data_Protector_home\Config\Server\Integ\Config\Sap\client_name%ORACLE_SID`

The parameters stored are:

- Oracle home directory
- encoded connection string to the target database
- BRTOOLS home directory
- the variables which need to be exported prior to starting a backup
- SAPDATA home directory
- user name and user group
- temporary directory used for the copy of the control file or redo logs
- list of control files and redo logs that will be copied to a safe location
- character set (ORA_NLS_CHARACTERSET)
- speed parameters (time needed for a specific file to back up - in seconds)
- manual balancing parameters

The configuration parameters are written to the Data Protector SAP R/3 configuration file:

- during configuration of the integration
- during creation of a backup specification
- when the configuration parameters are changed

❗ **IMPORTANT:** To avoid problems with your backups, take extra care to ensure the syntax and punctuation of your configuration file match the examples.

NOTE: You can set up the parameters in the `Environment` section (sublist) of the file by referring to other environment variables in the following way:

```
SAPDATA_HOME=${ORACLE_HOME}/data
```

Syntax

The syntax of the Data Protector SAP R/3 configuration file is as follows:

```
ORACLE_HOME='ORACLE_HOME';
ConnStr='ENCODED_CONNECTION_STRING_TO_THE_TARGET_DATABASE';
BR_directory='BRTOOLS_HOME';
SAPDATA_HOME='SAPDATA_HOME';
ORA_NLS_CHARACTERSET='CHARACTER_SET';
OSUSER='USER_NAME';
OSGROUP='USER_GROUP';
Environment={
  [ENV_var1='value1';]
  [ENV_var2='value2';]
```

```

    ...]
}
SAP_Parameters={backup_spec_name=(' -concurrency #_of_concurrency
' | '-time_balance' | '-load_balance' | '-manual_balance');
}
speed={
AVERAGE=1;
'filename'=#_of_seconds_needed_to_back_up_this_file;
}
compression={'filename'=size_of_the_file_in_bytes_after_the
_compression;
}
manual_balance={backup_specification_name={
'filename'=device_number;
}
}
}

```

The `ORA_NLS_CHARACTERSET` parameter is set automatically by Data Protector during SAP R/3 database configuration. For details of how to configure SAP R/3 database for use with Data Protector, see [“Configuring SAP R/3 databases” \(page 117\)](#).

Example

This is an example of the file:

```

ORACLE_HOME='/app/oracle805/product';
ConnStr='EIBBKIBBEIBBFIBBGHBBBOHBB
QDBBOFBBCFBFBPFBBBCFBBIFFBBGFBBBDGBBBFBBCFBBDFFBBCFBFB';
BR_directory='/usr/sap/ABA/SYS/exe/run';
SAPDATA_HOME='/sap';
ORA_NLS_CHARACTERSET='USASCII7';
OSUSER='orasisd';
OSGROUP='dba';

Environment={
  SAP_Parameters={
    sap_weekly_offline=(' -concurrency 1', '-no_balance');
    sap_daily_online=(' -concurrency 3', '-load_balance');
    sap_daily_manual=(' -concurrency 3', '-manual_balance');
  }
  speed={
    AVERAGE=203971;
    '/file1'=138186;
    '/file2'=269756;
  }
  compression={
    '/file1'=1234;
    '/file2'=5678;
  }
  manual_balance={
    sap_daily_manual={
      '/file1'=1; /* file 1 is backed up by the first sapback */
      '/file2'=2; /* file 2 is backed up by the second sapback */
      '/file3'=1; /* file 3 is backed up by the first sapback */
      '/file4'=1;
    }
  }
}

```

Setting, retrieving, listing, and deleting Data Protector SAP R/3 configuration file parameters using the CLI

The Data Protector SAP R/3 configuration file parameters are normally written to the Data Protector SAP R/3 configuration file after:

- the Data Protector configuration of the Oracle instance that is run by SAP R/3 is completed.
- a new backup specification is created.
- a backup that uses balancing by time algorithm is completed.

The `util_cmd` command

You can set, retrieve, list, or delete the Data Protector SAP R/3 configuration file parameters using the `util_cmd -putopt` (setting a parameter), `util_cmd -getopt` (retrieving a parameter), or `util_cmd -getconf` (listing all parameters) command on the Data Protector SAP R/3 client. The command resides in the `Data_Protector_home\bin` (Windows systems) or `/opt/omni/lbin` (HP-UX, Solaris systems) directory.

Cluster-aware clients

In a cluster environment, the environment variable `OB2BARHOSTNAME` must be defined as the virtual hostname before running the `util_cmd` command from the command line (on the client). The `OB2BARHOSTNAME` variable is set as follows:

- On UNIX: `export OB2BARHOSTNAME=virtual_hostname`
- On Windows: `set OB2BARHOSTNAME=virtual_hostname`

The `util_cmd` synopsis

The syntax of the `util_cmd` command is as follows:

```
util_cmd -getconf[ig] SAP_oracle_instance [-local filename]
util_cmd -getopt[ion] [SAP_oracle_instance] option_name [-sub[list]
sublist_name] [-local filename]
util_cmd -putopt[ion] [SAP_oracle_instance] option_name [option_value]
[-sub[list] sublist_name] [-local filename]
```

where:

`option_name` is the name of the parameter

`option_value` is the value for the parameter

`[-sub[list] sublist_name]` specifies the sublist in the configuration file to which a parameter is written to or taken from.

`[-local filename]` specifies one of the following:

- When it is used with the `-getconf[ig]` option, it specifies the filename for the output of the command to be written to. If the `-local` option is not specified, the output is written to the standard output.
- When it is used with the `-getopt[ion]`, it specifies the filename of the file from which the parameter and its value are to be taken and then written to the standard output. If the `-local` option is not specified, the parameter and its value are taken from the Data Protector SAP R/3 configuration file and then written to the standard output.
- When it is used with the `-putopt[ion]` option, it specifies the filename for the output of the command to be written to. If the `-local` option is not specified, the output is written to the Data Protector SAP R/3 configuration file.

NOTE: If you are setting the *option_value* parameter as a number, the number must be put in single quotes, surrounded by double quotes.

Return values

The `util_cmd` command displays a short status message after each operation (writes it to the standard error):

- Configuration read/write operation successful.
This message is displayed when all the requested operations have been completed successfully.
- Configuration option/file not found.
This message is displayed when either an option with the specified name does not exist in the configuration, or the file specified as the `-local` parameter does not exist.
- Configuration read/write operation failed.
This message is displayed if any fatal errors occurred, for example: the Cell Manager is unavailable, the Data Protector SAP R/3 configuration file is missing on the Cell Manager, and so on.

Setting parameters

To set the Data Protector `OB2OPTS` and the Oracle `BR_TRACE` parameters for the Oracle instance `ICE` that is run by SAP R/3, use the following commands on the Data Protector SAP R/3 client:

Windows

```
Data_Protector_home\bin\util_cmd -putopt SAP ICE OB2OPTS '-debug \ 1-200  
debug.txt' -sublist Environment
```

```
Data_Protector_home\bin\util_cmd -putopt SAP ICE BR_TRACE "'10'" \  
-sublist Environment
```

HP-UX, Solaris

```
/opt/omni/lbin/util_cmd -putopt SAP ICE OB2OPTS '-debug 1-200 \  
debug.txt' -sublist Environment
```

```
/opt/omni/lbin/util_cmd -putopt SAP ICE BR_TRACE "'10'" \ -sublist  
Environment
```

Retrieving parameters

To retrieve the value of the `OB2OPTS` parameter for the Oracle instance `ICE`, use the following command on the Data Protector SAP R/3 client:

- On Windows: `Data_Protector_home\bin\util_cmd -getopt SAP ICE \ OB2OPTS -sublist Environment`
- On HP-UX, Solaris: `/opt/omni/lbin/util_cmd -getopt SAP ICE OB2OPTS \ -sublist Environment`

Listing parameters

To list all the Data Protector SAP R/3 configuration file parameters for the Oracle instance `ICE`, use the following command on the Data Protector SAP R/3 client:

- On Windows: `Data_Protector_home\bin\util_cmd -getconf SAP ICE`
- On HP-UX, Solaris: `/opt/omni/lbin/util_cmd -getconf SAP ICE`

Deleting parameters

To remove the value of the OB2OPTS parameter for the Oracle instance ICE, use the following command on the Data Protector SAP R/3 client:

- On Windows: `Data_Protector_home\bin\util_cmd -putopt SAP ICE \ OB2PTS "" -sublist Environment`
- On HP-UX, Solaris: `/opt/omni/lbin/util_cmd -putopt SAP ICE OB2OPTS "" \ -sublist Environment`

Configuring the integration

To configure the integration:

1. Configure the required user accounts. See “Configuring user accounts” (page 113).
2. Configure SQL*Net V2 or Net8 TNS listener. See “Configuring SQL*Net V2 or Net8 TNS listener” (page 113).
3. Check the connection to the Oracle database from the application system. See “Checking the connection” (page 114).
4. Enable the use of the authentication password file. See “Authentication password file” (page 115).
5. Optionally, set the archived logging mode to enable online backups. See “Enabling archived logging” (page 115).
6. Share directories on the application system. See “Sharing directories on the application system” (page 116).
7. Configure every SAP R/3 database you intend to back up from or restore to. See “Configuring SAP R/3 databases” (page 117).
8. Configure the SAP R/3 parameter file. See “Configuring the SAP R/3 parameter file” (page 123).

Prerequisites

- Ensure that you have correctly installed and configured the SAP R/3 application. The database used by the SAP R/3 application must be an Oracle database. If any other database is used, you can back it up using the corresponding Data Protector integration. It is assumed that you are familiar with the SAP R/3 application and Oracle database administration.
 - For supported versions, platforms, devices, and other information, see the latest support matrices at <http://www.hp.com/support/manuals>.
 - For information on installing, configuring, and using the SAP R/3 application and the SAP backup and restore tools (BRBACKUP, BRRESTORE, and BRARCHIVE), see the SAP R/3 application documentation.
- Ensure that you have a license to use the Data Protector SAP R/3 ZDB integration. For information, see the *HP Data Protector Installation and Licensing Guide*.
- Ensure that you have correctly installed Data Protector.
 - For information on how to install the Data Protector disk array integration (P6000 EVA Array, P9000 XP Array, or EMC) with SAP R/3 in various architectures, see the *HP Data Protector Installation and Licensing Guide*.
 - On how to configure the Data Protector ZDB integration (EMC, P9000 XP Array, or P6000 EVA Array), see the *HP Data Protector Zero Downtime Backup Administrator's Guide*.
 - For information on the Data Protector Cell Manager package configuration in the MC/SG cluster, see the online Help index: “MC/ServiceGuard integration”.

NOTE: You cannot run ZDB sessions in the RMAN mode.

- The SAP R/3 directories `SAPBACKUP`, `SAPARCH`, `SAPREORG`, `SAPCHECK`, and `SAPTRACE` must not reside on the same disk array source volumes as the data files. Otherwise, the `BRTOOLS` data needed for complete recovery of a database is overwritten during instant recovery. You can set the locations for these directories in the `initDBSID.sap` file.

Before you begin

- Configure devices and media for use with Data Protector.
- To test whether the SAP R/3 system and the Cell Manager communicate properly, configure and run a Data Protector filesystem backup and restore.
- **Windows systems except Windows Server 2008:** Restart the Data Protector `Inet` service under the Oracle operating system user account described in “[Configuring user accounts](#)” (page 113). For information on changing the Data Protector `Inet` account, see the online Help index: “changing Data Protector `Inet` account”.

If there are several SAP R/3 instances running on the same system with different SAP administrator accounts configured for each instance, create an additional, common SAP administrator account. Configure the Data Protector `Inet` service to use this account as the service startup account.

Cluster-aware clients

- Configure SAP R/3 databases only on one cluster node, since the configuration files reside on the Cell Manager.
UNIX: During the configuration, Data Protector creates a link to the Data Protector `backint` and `splitint` programs on the currently active node. On all the other nodes, do it manually. Run:

```
ln -s /opt/omni/sbin/backint \ /usr/sap/ORACLE_SID/sys/exe/run
```

If `splitint` is supported by `BRTOOLS`, also run:

```
ln -s /opt/omni/sbin/ob2smbsplit \
/usr/sap/ORACLE_SID/sys/exe/run/splitint
```

Windows: During the configuration, Data Protector copies the Data Protector `backint` and `ob2smbsplit.exe` programs (the latter only if `splitint` is supported by `BRTOOLS`) from `Data_Protector_home\bin` to the directory that stores the SAP backup tools and renames `ob2smbsplit.exe` to `splitint.exe`. This is done only on the currently active node. On the other node, do it manually.

- If you intend to use the Data Protector CLI, set the Data Protector environment variable `OB2BARHOSTNAME` to the virtual server name as follows:

Windows: set `OB2BARHOSTNAME=virtual_server_name`

UNIX: export `OB2BARHOSTNAME=virtual_server_name`

- **Tru64:** Create the following links:

```
ln -s /sapfiles/admin/dbs/initsap.dba initSAP.dba
```

```
ln -s /sapfiles/admin/dbs/initsap.ora initSAP.ora
```

```
ln -s /sapfiles/admin/dbs/initsap.sap initSAP.sap
```

NOTE: SAP recommends to install SAP backup utilities on all cluster nodes.

Configuring user accounts

To enable backup and restore of SAP R/3 database files, you need to configure or create several user accounts.

Oracle operating system user account	<p>Operating system user account that is added to the following user groups:</p> <ul style="list-style-type: none">• UNIX systems: dba and sapsys• Windows systems: ORA_DBA and ORA_SID_DBA local groups <p>For example, user oraSID.</p> <p>UNIX systems only: Ensure that this user is the owner of the filesystem or of the raw logical volume on which the database is mounted. The minimum permissions should be 740.</p>
User account root (UNIX systems only)	<p>Default operating system administrator's user account added to the dba user group.</p>
Oracle database user account	<p>Database user account granted at least the following Oracle roles:</p> <ul style="list-style-type: none">• sysdba• sysoper <p>For example, user system.</p> <p>Do not configure the SYS user for backing up SAP R/3 objects. When using SYS, the SAP backup fails with ORA-28009 error.</p>

Add the following user accounts to the Data Protector admin or operator user group:

- Oracle operating system user account
- **UNIX systems only:** User account root (for both the application system and backup system)

In cluster environments, add these user accounts to the Data Protector admin or operator user group for the following clients:

- virtual server
- every node in the cluster

For information on adding Data Protector users, see the online Help index: "adding users".

Configuring SQL*Net V2 or Net8 TNS listener

1. Ensure that the listener.ora and tnsnames.ora files on the application system are configured as shown in the following example. The files are located in:

UNIX: ORACLE_HOME/network/admin

Windows: ORACLE_HOME\network\admin

Example

Oracle instance: PRO

Application system: alpha.hp.com

listener.ora	<pre> LISTENER = (DESCRIPTION_LIST = (DESCRIPTION = (ADDRESS = (PROTOCOL = TCP) (HOST = alpha.hp.com) (PORT = 1522)))) SID_LIST_LISTENER = (SID_LIST = (SID_DESC = (GLOBAL_DBNAME = PRO) (SID_NAME = PRO) (ORACLE_HOME = /app/oracle815/product))) </pre>
tnsnames.ora	<pre> PRO = (DESCRIPTION = (ADDRESS_LIST = (ADDRESS = (PROTOCOL = TCP) (HOST = alpha.hp.com) (PORT = 1522))) (CONNECT_DATA = (SERVICE_NAME = PRO))) </pre>

2. Start the SQL*Net V2 or Net8 TNS listener by running the following on the Oracle Server system:

UNIX: `ORACLE_HOME/bin/lsnrctl start`

Windows: `ORACLE_HOME\bin\lsnrctl start`

Checking the connection

To check the connection to the Oracle instance from the application system:

1. Log in to the application system as the Oracle OS user.
2. Export/set the `ORACLE_HOME` and `ORACLE_SID` variables.
3. Start sqlplus.
4. Connect to the Oracle target database as the Oracle database user, first with the `sysdba` role and then with the `sysoper` role.

Example

For the following configuration:

Oracle instance: PRO ORACLE_HOME: /app/oracle816/product

run:

```
id
uid=102(oracle) gid=101(dba)
export ORACLE_SID=PRO
export ORACLE_HOME=/app/oracle816/product
export SHLIB_PATH=/app/oracle816/product/lib:/opt/omni/lib
sqlplus /nolog
SQLPLUS> connect system/manager@PRO as sysdba;
Connected.
SQLPLUS> connect system/manager@PRO as sysoper;
Connected.
```

Authentication password file

Enable the use of the authentication password file for the database administrator:

1. Shut down the Oracle target database on the application system.
2. In the `initORACLE_SID.ora` file, specify:

```
remote_login_passwordfile = exclusive
```

For instructions on how to set up the password file, see the Oracle documentation.

Enabling archived logging

When you set the database to the archived logging mode, you protect the unsaved online redo logs from being overwritten. Online backup of data files is useless without the related redo logs because you cannot recover the database to a consistent state.



TIP: Archive the redo log files generated during the online backup immediately after BRBACKUP completes.

To protect the archive directory from overflowing, clear the directory regularly.

To enable archived logging:

1. In the `initORACLE_SID.ora` file, set
`log_archive_start = true`
and specify the `log_archive_dest` option.

Example

This is an example of the `initORACLE_SID.ora` file for the Oracle instance PRO:

```
# @(#)initSID.ora      20.4.6.1      SAP      98/03/30
#####
# (c)Copyright SAP AG, Walldorf
#####
. . . .
. . . . .
. . . . .
. . . . .
### ORACLE Authentication Password File
remote_login_passwordfile = exclusive
### ORACLE archiving
log_archive_dest = /oracle/PRO/saparch/PROarch
log_archive_start = true
. . . .
```

2. Mount the Oracle database and start the archived logging mode using the Oracle Server Manager. Run:

```
startup mount
alter database archivelog;
archive log start;
alter database open;
```

Example

For the Oracle instance PRO, run:

UNIX: export ORACLE_SID=PRO

Windows: set ORACLE_SID=PRO

```
sqlplus /nolog
SQLPLUS> connect user/passwd@PRO;
Connected.
SQLPLUS> startup mount
ORACLE instance started.
Total System Global Area          6060224 bytes
Fixed Size                        47296 bytes
Variable Size                    4292608 bytes
Database Buffers                 1638400 bytes
Redo Buffers                      81920 bytes
Database mounted.
SQLPLUS> alter database archivelog;
Statement processed.
SQLPLUS> archive log start;
Statement processed.
SQLPLUS> alter database open;
```

Sharing directories on the application system

The following directories on the application system must be accessible:

- sapbackup
- sapreorg
- Oracle home directory
- BR*Tools home directory

NOTE: The sapreorg and BR*Tools home directories must be accessible only if you want to run SAP compliant ZDB sessions (BRBACKUP is started on the backup system and not on the application system).

UNIX application system

1. Share the directories on the application system through NFS with root permissions.

For example, suppose that the sapbackup directory on the application system points to /oracle/SID/sapbackup and the backup system is backup.company.com. To share the sapbackup directory:

HP-UX: In the file /etc/exports on the application system, add the line:

```
/oracle/SID/sapbackup -root=backup.company.com
```

Solaris: In the file /etc/dfs/dfstab on the application system, add the line:

```
share -F nfs -o root=backup.company.com /oracle/SID/sapbackup
```

2. Mount the directories on the backup system. Ensure that you have the same directory structure on both the application and backup system.

For example, suppose that you have an HP-UX application system `app.company.com`. To mount the `/oracle/SID/sapbackup` directory on the backup system, add the following line to the file `/etc/fstab` on the backup system:

```
app.company.com:/oracle/SID/sapbackup  
/oracle/SID/sapbackup nfs defaults 0 0
```

Windows application system

- On the application system, find the location of the `sapbackup` and `sapreorg` directories. If they reside inside the SAP data home directory, share this directory under any name you want. Then, on the backup system, create the `HKEY_LOCAL_MACHINE\SOFTWARE\SAP\ORACLE_SID\Environment\SAPDATA_HOME` Windows registry key, specifying the SAP data home directory path as seen from the backup system.

For example, suppose your application system is `mycomputer.company.com` and the SAP data home is `K:\oracle\my_instance`, which you share under the name `my_SAPinstance`. Then, the `SAPDATA_HOME` Windows registry key on the backup system must have the value `\\mycomputer.company.com\my_SAPinstance`.

If the `sapbackup` and `sapreorg` directories do not reside together, share each directory separately and also create a separate Windows registry key on the backup system (`SAPBACKUP`, `SAPREORG`).

- Share the Oracle home directory on the application system and specify its path in the `ORACLE_HOME` Windows registry key on the backup system, similarly as described above.
- Ensure that the `BR*Tools` home directory on the application system is accessible from the backup system as follows:

```
\\application_system\sapmnt\SAP_SID\SYS\exe\run
```



TIP:

Instead of creating registry keys, you can also set the Data Protector SAP R/3 integration environment variables (`ORACLE_HOME`, `SAPDATA_HOME`, `SAPBACKUP`, `SAPREORG`).

Choosing authentication mode

Data Protector SAP R/3 ZDB integration supports two authentication modes for accessing Oracle databases that are used by SAP R/3:

- database authentication mode
- operating system authentication mode

With database authentication mode, you need to re-configure the SAP R/3 integration for an SAP R/3 database with the new Oracle login information each time the corresponding Oracle database user account changes. Such a reconfiguration is not needed if operating system authentication mode is used.

You select the preferred authentication mode when you configure a particular SAP R/3 database.

Configuring SAP R/3 databases

You need to provide Data Protector with the following configuration parameters:

- Oracle Server home directory
- SAP R/3 data home directory

- Optionally, if you choose database authentication mode, Oracle database user account. The user account is used by BRBACKUP and BRARCHIVE during backup.
- Directory in which the SAP backup utilities are stored

Data Protector then creates the configuration file for the SAP R/3 database on the Cell Manager and verifies the connection to the database. On UNIX, Data Protector also creates a soft link for the `backint` program from the directory that stores the SAP backup utilities to `/opt/omni/libin`.

On Windows, Data Protector copies the `backint` and `programs` (the latter only if `splitint` is supported by BR*Tools) from `Data_Protector_home\bin` to the directory that stores the SAP backup tools and renames `ob2smbsplit.exe` to `splitint.exe`.

To configure an SAP R/3 database, use the Data Protector GUI or CLI.

Before you begin

- Ensure that the SAP R/3 database is open.

Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **SAP R/3**, and click **Add Backup**.
3. In the **Create New Backup** dialog box, select the template.

From the **Backup type** drop-down list, select the **Snapshot or split mirror backup** option, and from the **Sub type** drop-down list, select the appropriate disk array agent. The agent must be installed on the application system and the backup system.

Click **OK**.

4. In **Application system**, select the SAP R/3 client to be backed up. In cluster environments, select the virtual server.

In **Backup system**, select the backup system.

Specify other disk array specific backup options (see [“EMC backup options” \(page 126\)](#) for EMC, [“P9000 XP Array backup options” \(page 126\)](#) for P9000 XP Array, or [“P6000 EVA Array backup options” \(page 127\)](#) for P6000 EVA Array). For details on the options, press **F1**.

NOTE: *P9000 XP Array and P6000 EVA Array only:* To enable instant recovery, select **Track the replica for instant recovery**.

5. In **Application database**, type the Oracle instance name (`ORACLE_SID`).

Specify the **User and group/domain** options, which are available on UNIX and Windows Server 2008 clients, as follows:

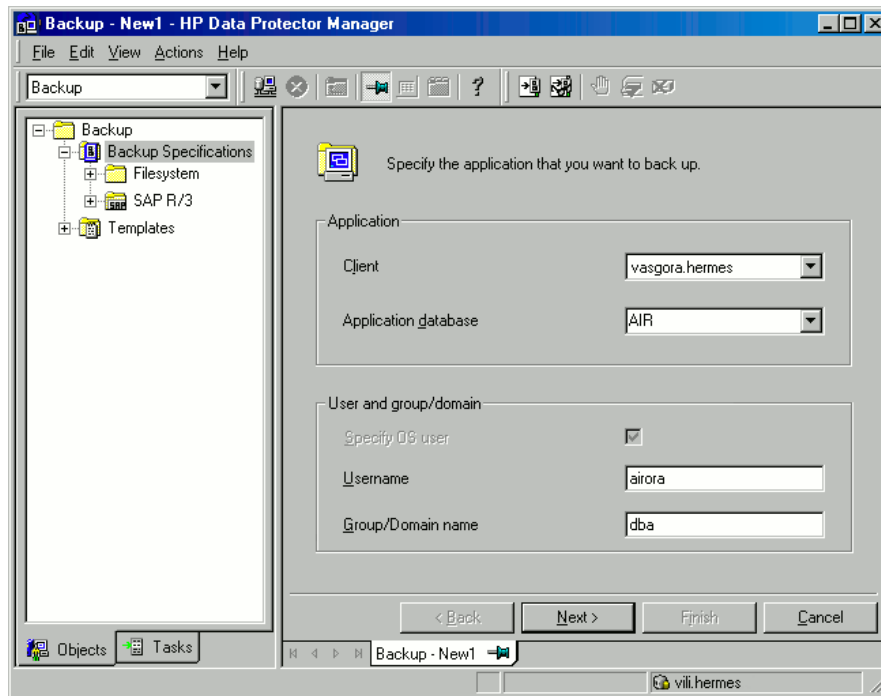
- **UNIX:** In **Username**, type the Oracle OS user described in [“Configuring user accounts” \(page 113\)](#). In **Group/Domain name**, type `dba`.

- **Windows Server 2008:**

In **Username** and **Group/Domain name**, specify the operating system user account under which you want the backup session to run (for example, the user name `Administrator`, domain `DP`).

Ensure that this user has been added to the Data Protector `admin` or `operator` user group and has the SAP R/3 backup rights. This user becomes the backup owner.

Figure 42 Specifying an SAP R/3 system and Oracle instance



Click **Next**.

6. In the **Configure SAP** dialog box, specify the pathname of the Oracle Server home directory and SAP R/3 data home directory. If you leave the fields empty, the default `ORACLE_HOME` directory is used.

Under **Oracle login information to target database**, specify the following:

- For the database authentication mode, specify **Username**, **Password**, and **Service**.
- For the local operating system authentication mode, leave **Username**, **Password**, and **Service** empty.
- For the remote operating system authentication mode, specify only **Service** (leave **Username** and **Password** empty).

The following are the option descriptions:

- **Username** and **Password**: Specify the user name and password of the Oracle database user account described in [“Configuring user accounts”](#) (page 113).
- **Service**: Specify the Oracle service name.

In **Backup and restore executables directory**, specify the pathname of the directory in which the SAP backup utilities reside. By default, the utilities reside in:

UNIX: `/usr/sap/ORACLE_SID/SYS/exe/run`

Windows: `\\SAP_system\sapmnt\ORACLE_SID\sys\exe\run`

Figure 43 Configuring an SAP R/3 database on a UNIX system (operating system authentication mode)

The 'Configure SAP' dialog box for UNIX system configuration contains the following fields and values:

- Client:** vasgora.company.com
- Oracle SID:** AIR
- Oracle Server home directory:** /app/oracle/product/81
- SAP Data home directory:** /app/oracle/product/AIR
- Oracle login information to target database:**
 - Username:** (empty)
 - Password:** (empty)
 - Service:** AIR
- Backup and restore executables directory:** /usr/sap/CER/SYS/exe/run

Buttons at the bottom: OK, Cancel, Help.

Figure 44 Configuring an SAP R/3 database on a Windows system (database authentication mode)

The 'Configure SAP' dialog box for Windows system configuration contains the following fields and values:

- Client:** prem.company.com
- Oracle SID:** CER
- Oracle Server home directory:** F:\Oracle\Ora81
- SAP Data home directory:** F:\Oracle\CER
- Oracle login information to target database:**
 - Username:** system
 - Password:** (masked with asterisks)
 - Service:** CER
- Backup and restore executables directory:** F:\Oracle\CER\sys\exe\run

Buttons at the bottom: OK, Cancel, Help.

Click **OK**.

7. The SAP R/3 database is configured. Exit the GUI or proceed with creating the backup specification at [Step 7](#).

Using the Data Protector CLI

1. Log in to the SAP R/3 system using the Oracle operating system user account.
2. At the command prompt, change current directory to the following directory:

Windows systems: `Data_Protector_home\bin`

HP-UX, Solaris systems: /opt/omni/lbin

3. Run:

```
util_sap.exe -CONFIG ORACLE_SID ORACLE_HOME  
targetdb_connection_string SAPTOOLS_DIR [SAPDATA_HOME] [SQL_PATH]
```

Parameter description

ORACLE_SID

Oracle instance name.

ORACLE_HOME

Pathname of the Oracle Server home directory.

targetdb_connection_string

This argument value determines the authentication mode used for accessing the Oracle database:

- To select the database authentication mode, specify the login information to the target database in the format *user_name/password@Oracle_service*.
- To select the local operating system authentication mode, specify only the character */*.
- To select the remote operating system authentication mode, specify the login information to the target database in the format */@Oracle_service*.

SAPTOOLS_DIR

Pathname of the directory that stores the SAP backup utilities.

SAPDATA_HOME

Pathname of the directory where the SAP R/3 data files are installed. By default, this parameter is set to *ORACLE_HOME*.

The message *RETVAL*0 indicates successful configuration.

Handling errors

If you receive the message *RETVAL**error_number* where *error_number* is different than zero, an error occurred.

To get the error description:

Windows:

```
Data_Protector_home\bin\omnigetmsg 12 error_number
```

which is located on the Cell Manager.

HP-UX, Solaris: Run:

```
/opt/omni/lbin/omnigetmsg 12 error_number
```



TIP: To get a list of Oracle instances that are used by the SAP R/3 application, run:

```
util_sap.exe -APP
```

To get a list of tablespaces of an Oracle instance, run:

```
util_sap.exe -OBS0 ORACLE_SID
```

To get a list of database files of a tablespace, run:

```
util_sap.exe -OBS1 ORACLE_SID TABLESPACE
```

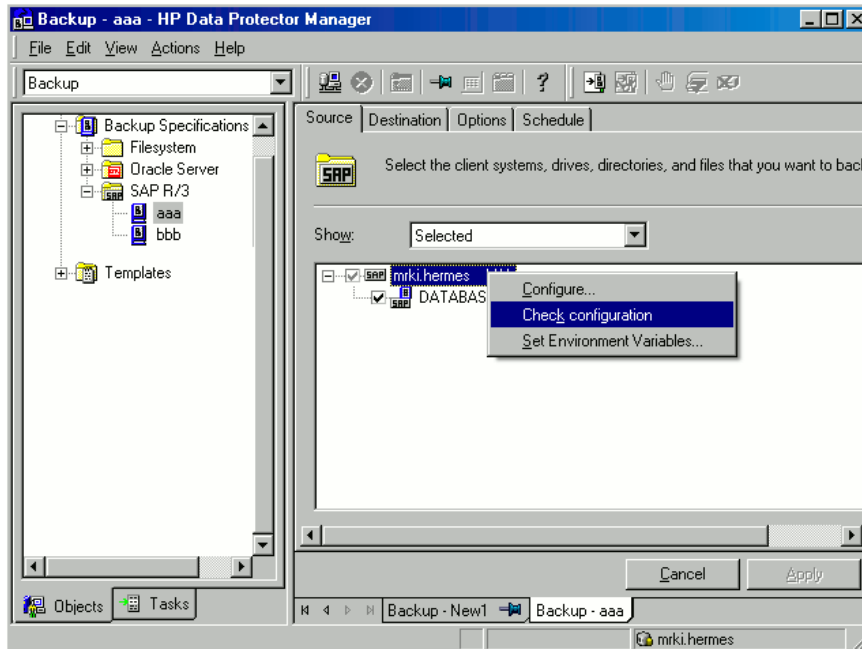
Checking the configuration

You can check the configuration of an SAP R/3 database after you have created at least one backup specification for this database. Use the Data Protector GUI or CLI.

Using the Data Protector GUI

1. In the Context List, select **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **SAP R/3**. Click the backup specification to display the Oracle instance to be checked.
3. Right-click the Oracle instance and click **Check configuration**.

Figure 45 Checking the SAP R/3 configuration



Using the Data Protector CLI

Log in to the SAP R/3 system as the Oracle OS user. From the directory:

Windows: `Data_Protector_home\bin`

HP-UX, Solaris: `/opt/omni/sbin`

run:

```
util_sap.exe -CHKCONF ORACLE_SID
```

where `ORACLE_SID` is the name of the Oracle instance.

A successful configuration check displays the message `*RETVAL*0`.

If you receive the message `*RETVAL*error_number` where `error_number` is different than zero, an error occurred. On how to get the error description, see [“Handling errors”](#) (page 121).

To check if the SAP R/3 configuration is suitable for instant recovery, from the directory,

Windows: `Data_Protector_home\bin`

HP-UX, Solaris: `/opt/omni/sbin`

run:

```
util_sap.exe -CHKCONF_IR ORACLE_SID [-verbose]
```

The `-verbose` option creates a file with a list of control files and redo log files that are on the same source volumes as the database files. If this list is not empty, a warning is displayed, stating that instant recovery is impossible. On how to solve this problem, see [“Reconfiguring an Oracle instance for instant recovery”](#) (page 269).

Configuring the SAP R/3 parameter file

To configure the integration, you need to set some parameters in the SAP R/3 parameter file on both the application system and backup system. The file template is located in:

UNIX: `ORACLE_HOME/dbs/initORACLE_SID.sap`

Windows: `ORACLE_HOME\database\initORACLE_SID.sap`

Table 10 SAP parameter file settings

Parameter	Value/Description
split_cmd	<p>On the application system:</p> <p>UNIX: <code>/opt/omni/lbin/ob2smbssplit \$"</code></p> <p>Windows: <code>"Data_Protector_home\bin\ob2smbssplit \$"</code></p> <p>On the backup system, you do not need to set the parameter.</p> <p>BRBACKUP uses this parameter to trigger the replica creation. At run time, the optional sign "\$" is replaced with the name of the text file containing the names of files to be backed up.</p> <p>Windows only: If the pathname contains spaces, use Windows short names instead.</p>
primary_db	<p>On the application system: LOCAL</p> <p>On the backup system: name of the service used for connecting to the Oracle database.</p> <p>This parameter defines the service name of the Oracle database to link the backup system to the application system.</p>

Backup

The integration provides online and offline database backups of the following types:

- ZDB to disk
- ZDB to tape
- ZDB to disk+tape

To configure a backup, create a ZDB backup specification.

Archived logs can only be backed up in a ZDB to disk+tape, ZDB to tape, or non-ZDB (standard backup) session. If you try to back up archived logs in a ZDB to disk session, the session fails.



TIP: Create a separate (non-ZDB) backup specification for backing up archived logs.

What is backed up depends on your selection in the backup specification. For details, see [“What is backed up”](#) (page 123).

Table 11 What is backed up

Selected items	Backed up files
ARCHIVELOGS	<ul style="list-style-type: none">• offline (archived) redo logs• control files
DATABASE or individual tablespaces	<ul style="list-style-type: none">• data files• control files• SAP R/3 logs and parameter files• online redo logs (only during offline backups)

You can specify SAP R/3 backup options in two different ways:

- Using the BRBACKUP options.
- Using the SAP parameter file.

NOTE: The BRBACKUP options override the settings in the SAP parameter file.

You can specify BRBACKUP options when you create a backup specification. If no options are specified, the SAP R/3 application refers to the current settings in the SAP parameter file. In such a case, before running a backup, ensure that the SAP parameter file is correctly configured.



TIP: When you create a backup specification, select a backup template that already contains the desired BRBACKUP options.

Considerations

- Before you start a backup, ensure that the SAP R/3 database is in the open or shutdown mode.
- ZDB, restore, and instant recovery sessions that use the same source volume on the application system cannot run simultaneously.
- You cannot start a ZDB to disk session if another session is backing up the archived logs, even if the Oracle data files and the archived logs reside on different source volumes.
- **P9000 XP Array only:** If the LVM Mirroring configuration is used, Data Protector displays a warning during a backup because the volume group source volumes on the application system do not have their BC pairs assigned. The message should be ignored.
- **ZDB to disk only:** Archived logs cannot be backed up. To back up archived logs, create a non-ZDB (standard) SAP R/3 backup specification. For information, see the *HP Data Protector Integration Guide for Oracle and SAP*.
- Configurable backup modes are supported only by using templates.
- By default, Data Protector supports all BRTOOL options except -a and -b. To enable support for -a and -b, set the OB2BRTNOSECU omnirc variable to 1. On how to set the omnirc variable, see the online Help index: "omnirc options".

Creating backup specifications

Create a backup specification using the Data Protector Manager.

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **SAP R/3**, and click **Add Backup**.
3. In the **Create New Backup** dialog box, select a template ("[Selecting a template](#)" (page 125)) and click **OK**.

Figure 46 Selecting a template

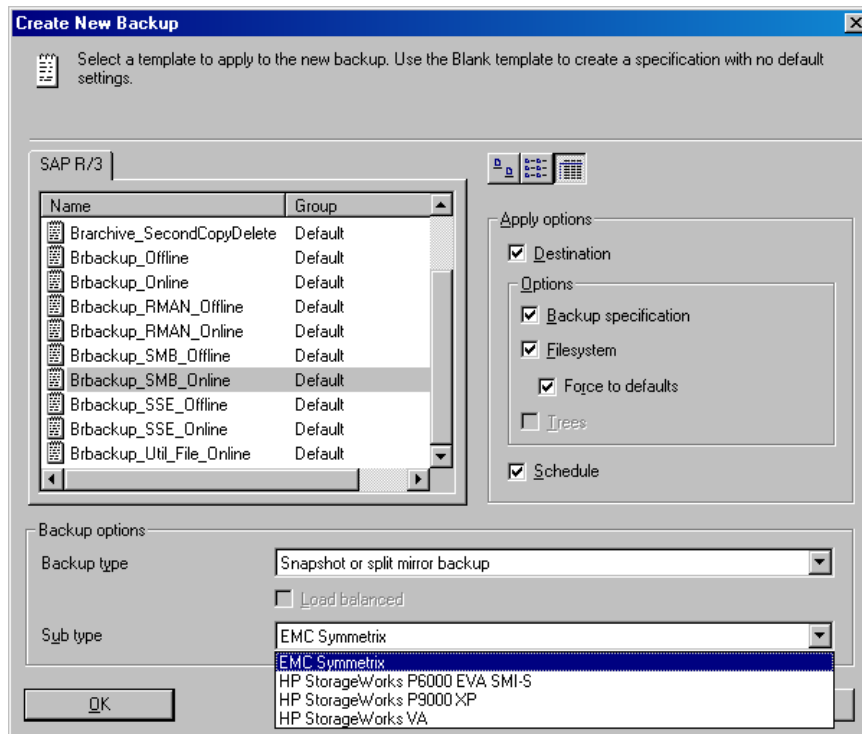


Table 12 Backup templates

Blank SAP Backup	No predefined options.
Brbackup_SMB_Offline	Used for an offline ZDB (split mirror or snapshot backup). The database is stopped during the creation of a replica.
Brbackup_SMB_Online	Used for an online ZDB (split mirror or snapshot backup). The database is active during the creation of a replica.
Brbackup_SPLITINT_Offline	Used for an offline ZDB (split mirror or snapshot backup) using <code>splitint</code> . The database is stopped during the creation of a replica. The database is offline for a shorter period of time than if <code>SMB_offline</code> was used, but <code>splitint</code> must be supported by BRTOOLS.
Brbackup_SPLITINT_Online	Used for an online ZDB (split mirror or snapshot backup) using <code>splitint</code> . The database is active during the creation of a replica. The database is in the backup mode for a shorter period of time than if <code>SMB_online</code> was used, but <code>splitint</code> must be supported by BRTOOLS.
Brbackup_offline_mirror	Used for an offline ZDB (split mirror or snapshot backup). The database is stopped during the split phase of creation of a replica.
Brbackup_online_mirror	Used for an online ZDB (split mirror or snapshot backup). The database is active during the split phase of creation of a replica.

From the **Backup type** drop-down list, select **Snapshot or split mirror backup**.

From the **Sub type** drop-down list, select the appropriate disk array agent. The agent must be installed on the application system and the backup system.

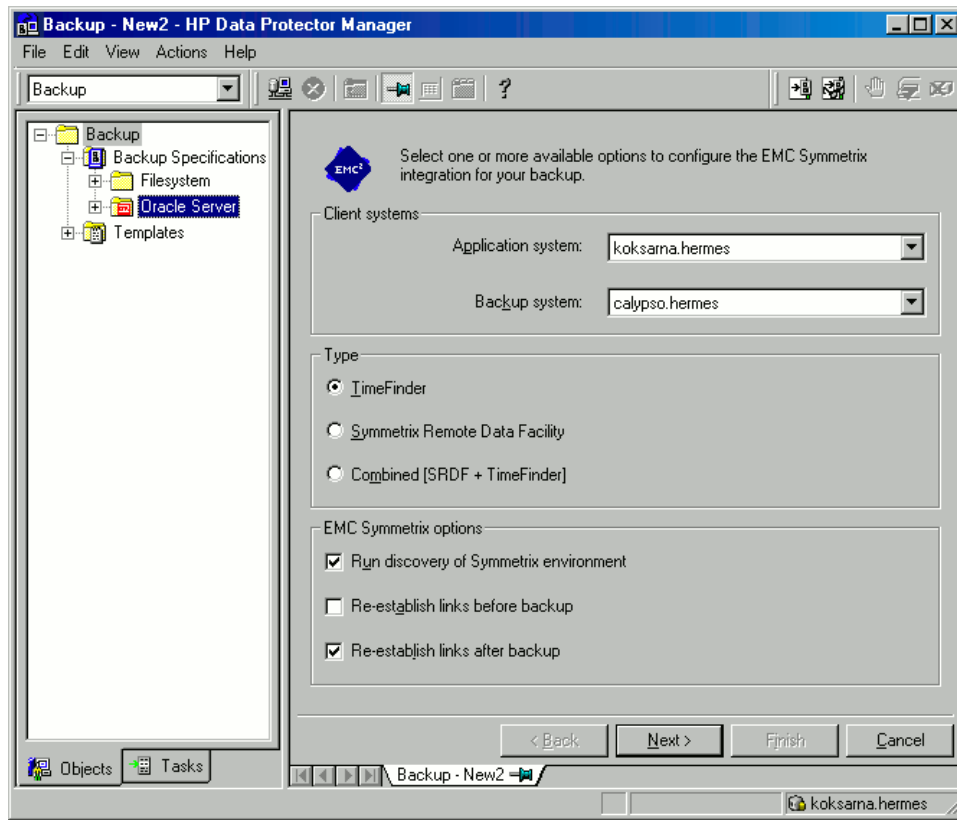
4. In **Application system**, select the SAP R/3 client to be backed up. In cluster environments, select the virtual server.

In **Backup system**, select the backup system.

Select other disk array-specific backup options (see [“EMC backup options”](#) (page 126) for EMC, [“P9000 XP Array backup options”](#) (page 126) for P9000 XP Array, or [“P6000 EVA](#)

Array backup options” (page 127) for P6000 EVA Array). For detailed information on the backup options, press **F1**.

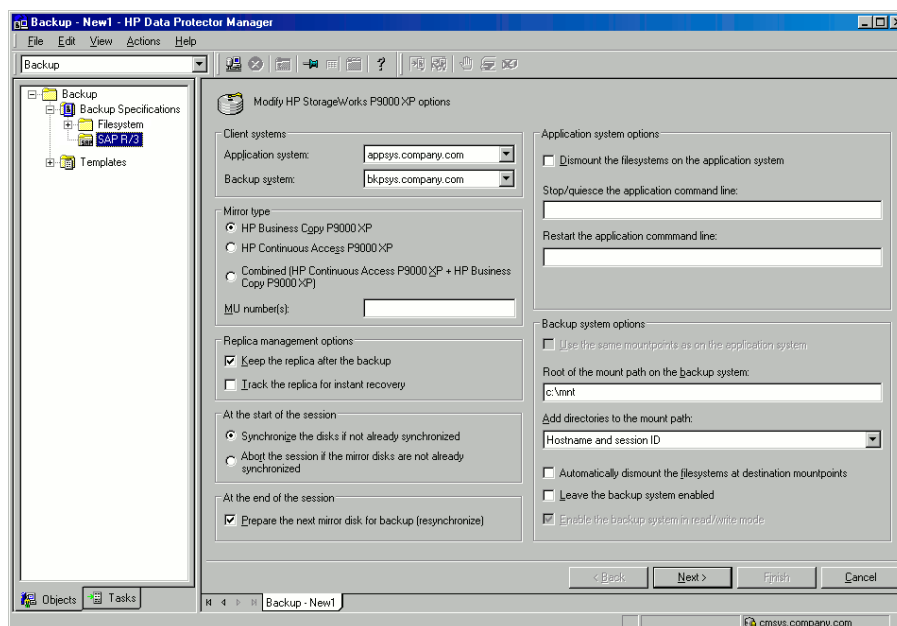
Figure 47 EMC backup options



P9000 XP Array specifics

To enable instant recovery, leave the **Track the replica for instant recovery** option selected. It is not possible to run instant recovery with Data Protector if this option is cleared.

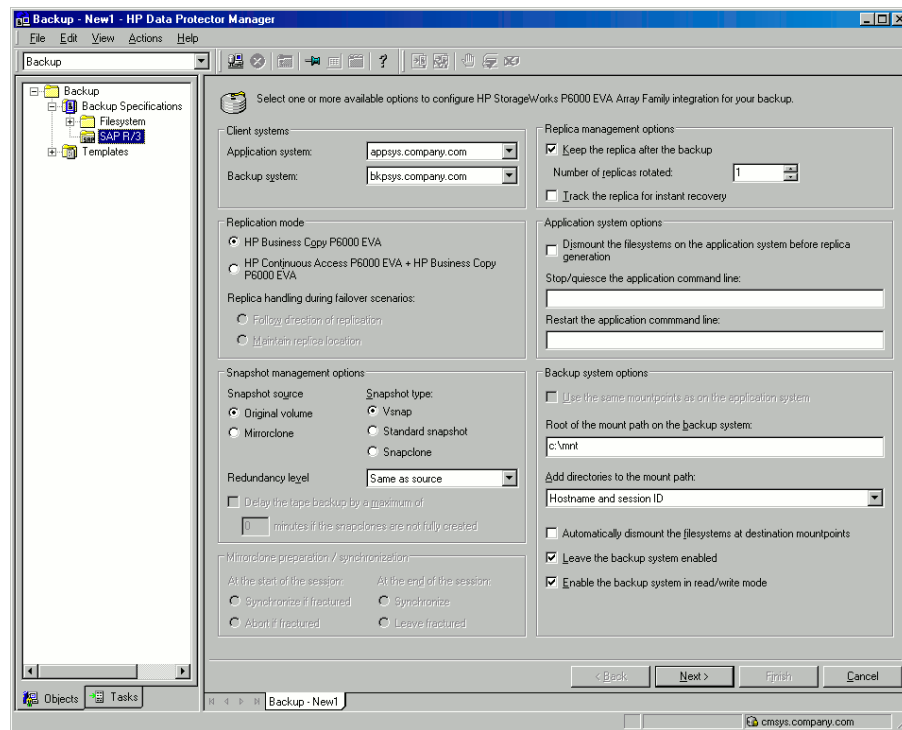
Figure 48 P9000 XP Array backup options



P6000 EVA Array specifics

To enable instant recovery, select the **Track the replica for instant recovery** option.

Figure 49 P6000 EVA Array backup options



Click **Next**.

5. In **Application database**, select the Oracle instance (ORACLE_SID) to be backed up. Specify the **User and group/domain** options, which are available on UNIX and Windows Server 2008 clients, as follows:
 - **UNIX:** In **Username**, type the Oracle OS user described in [“Configuring user accounts”](#) (page 113). In **Group/Domain name**, type dba.
 - **Windows Server 2008:**
In **Username** and **Group/Domain name**, specify the operating system user account under which you want the backup session to run (for example, the user name Administrator, domain DP).

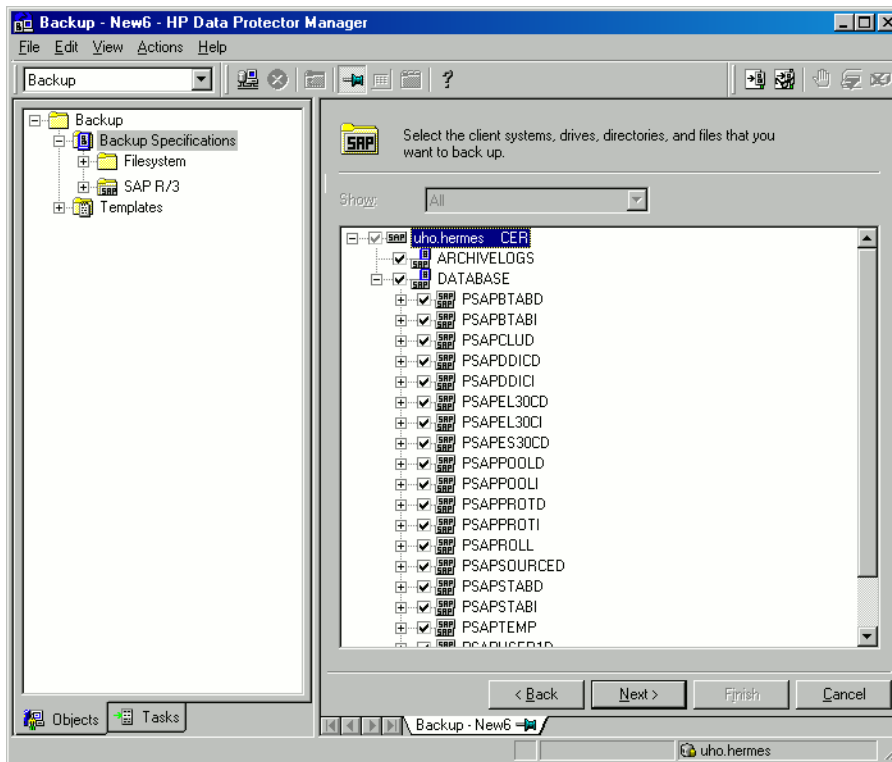
Ensure that this user has been added to the Data Protector admin or operator user group and has the SAP R/3 backup rights. This user becomes the backup owner.

Click **Next**.

6. If the SAP R/3 database is not configured yet for use with Data Protector, the **Configure SAP** dialog box is displayed. Configure it as described in [“Configuring SAP R/3 databases”](#) (page 117).
7. Select SAP R/3 objects to be backed up. You can select individual tablespaces, data files, or archived logs.

NOTE: If you plan to do instant recovery, select the whole **DATABASE** item.

Figure 50 Selecting backup objects



Click **Next**.

8. Select devices to use for the backup.

To specify device options, right-click the device and click **Properties**. Specify the number of parallel backup streams in the **Concurrency** tab and the media pool.

NOTE: Parallelism (the number of streams your SAP R/3 database is backed up with) is set automatically. If load balancing is used, the number of streams equals the sum of concurrencies of the selected devices.

Click **Next**.

9. Set backup options. For information on the application-specific options, see [“SAP R/3 backup options” \(page 129\)](#).

Figure 51 Application-specific options

Application Specific Options

SAP integration

SAP 4.5 integration specific options

Options

Log file:

BR Backup:

Backup Objects:

BR Archive:

Balancing:

Pre-exec:

Pgst-exec:

Backup mode: ☒ All ☐ Full

☒ Use default RMAN channels

Objects outside database:

OK Cancel Help

Click **Next**.

10. Optionally, schedule the backup. See [“Scheduling backup specifications”](#) (page 130).

Click **Next**.

11. Save the backup specification, specifying a name and a backup specification group.



TIP: Preview your backup specification before using it for real. See [“Previewing backup sessions”](#) (page 131).

Table 13 SAP R/3 backup options

Option	Description
Log file	If you want to create a backint log file during backup, specify a pathname for the file. By default, this file is not created because Data Protector stores all relevant information about backup sessions in the database.
BR Backup	Specifies BRBACKUP options. For example, for online backup using the <code>splitint</code> interface, type <code>-t online_mirror</code> . If <code>splitint</code> is not supported by BRTOOLS, type <code>-t online_split</code> . To run BRBACKUP under a different Oracle database user than the one specified during the configuration, type <code>-u user_name</code> .
Backup Objects	Lists BRBACKUP options passed by <code>omnisap.exe</code> . The list is displayed after you save the backup specification.
BR Archive	Specifies BRARCHIVE options. Not applicable for ZDB to disk.

Table 13 SAP R/3 backup options (*continued*)

Option	Description
Balancing: By Load	Groups files into subsets of approximately equal sizes. The subsets are then backed up concurrently by Data Protector <code>sapback</code> programs. If your backup devices use hardware compression, the sizes of the original and backed up files differ. To inform Data Protector of this, specify the original sizes of the backed up files in the <code>compression</code> section of the Data Protector SAP R/3 configuration file. See “Data Protector SAP R/3 configuration file” (page 107).
Balancing: By Time	Groups files into subsets that are backed up in approximately equal periods of time. The duration depends on the file types, speed of the backup devices, and external influences (such as mount prompts). This option is best for environments with large libraries of the same quality. The subsets are backed up concurrently by Data Protector <code>sapback</code> programs. Data Protector automatically stores backup speed information in the <code>speed</code> section of the Data Protector SAP R/3 configuration file. It uses this information to optimize the backup time. This type of balancing may lead to non-optimal grouping of files in the case of an online backup or if the speed of backup devices varies significantly.
Balancing: Manual	Groups files into subsets as specified in the manual balancing section of the Data Protector SAP R/3 configuration file. For more information, see “Manual balancing” (page 135). Not applicable for ZDB to disk.
Balancing: None	No balancing is used. The files are backed up in the same order as they are listed in the internal Oracle database structure. To check the order, use the Oracle Server Manager SQL command: <code>select * from dba_data_files</code>
Pre-exec, Post-exec	The command specified here is started by <code>omnisap.exe</code> on the SAP R/3 system before the backup (<code>pre-exec</code>) or after it (<code>post-exec</code>). Do not use double quotes. Provide only the name. The command must reside in the directory: Windows: <code>Data_Protector_home\bin</code> HP-UX, Solaris: <code>/opt/omni/bin</code>
Backup mode	Not applicable for ZDB.
Use default RMAN channels	Not applicable for ZDB.
Objects outside database	Specifies non-database files of the Oracle SAP R/3 environment to be saved. Save these files in a separate backup session.

NOTE: The total number of `sapback` processes started in one session using Data Protector is limited to 256.

Modifying backup specifications

To modify your backup specification, click its name in the Scoping Pane of the Backup context, then click the appropriate tab, and apply the changes.

Scheduling backup specifications

You can run unattended backups at specific times or periodically. For details on scheduling, see the online Help index: “scheduled backups”.

Scheduling example

To schedule ZDB to disk+tape at 8:00, 13:00, and 18:00 during week days:

1. In the **Schedule** property page, select the starting date in the calendar and click **Add** to open the **Schedule Backup** dialog box.

- Under **Recurring**, select **Weekly**. Under **Time options**, select **8:00**. Under **Recurring Options**, select **Mon, Tue, Wed, Thu, and Fri**.
From the **Split mirror/snapshot backup** drop-down list, select **To disk+tape**. See “[Scheduling backups](#)” (page 131).
Click **OK**.
- Repeat [Step 1](#) and [Step 2](#) to schedule backups at 13:00 and 18:00.
- Click **Apply** to save the changes.

Figure 52 Scheduling backups

Previewing backup sessions

Preview the backup session to test it. You can use the Data Protector GUI or CLI.

Using the Data Protector GUI

- In the Context List, click **Backup**.
- In the Scoping Pane, expand **Backup Specifications** and then **SAP R/3**. Right-click the backup specification you want to preview and click **Preview Backup**.
- Specify **Backup type** and **Network load**. Click **OK**.

The message `Session completed successfully` is displayed at the end of a successful preview.

Using the Data Protector CLI

From the directory:

Windows: `Data_Protector_home\bin`

HP-UX, Solaris: `/opt/omni/bin/`

run:

```
omnib -sap_list backup_specification_name -test_bar
```

What happens during the preview?

The `omnisap.exe` command is started, which starts the Data Protector `testbar` command to test the following:

- The syntax of the backup specification
- If devices are correctly specified
- If the necessary media are in the devices

Starting backup sessions

Interactive backups are run on demand. They are useful for urgent backups or restarting failed backups.

Backup methods

Start a backup of SAP R/3 objects in any of the following ways:

- Using the Data Protector GUI.
- Using the Data Protector CLI.
- Using the SAP BR*Tools.

Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, and then **SAP R/3**. Right-click the backup specification you want to use and click **Start Backup**.
3. Specify **Network load**. Click **OK**.

NOTE: Only the `Full` backup type is supported.

ZDB to disk, ZDB to disk+tape only: Specify the **Split mirror/snapshot backup** option.

The message `Session completed successfully` is displayed at the end of a successful backup session.

Using the Data Protector CLI

From the directory:

Windows: `Data_Protector_home\bin`

HP-UX, Solaris: `/opt/omni/bin/`

run:

ZDB to tape, ZDB to disk+tape:

`omnib -sap_list backup_specification_name`

ZDB to disk:

`omnib -sap_list backup_specification_name -disk_only`

For details, see the `omnib` man page or the *HP Data Protector Command Line Interface Reference*.

Using the SAP BRTOOLS

1. Log in to the SAP R/3 backup system or SAP R/3 application system as the Oracle OS user.
2. Export/set the following environment variables:

ORACLE_SID=*SAP_instance_name*

ORACLE_HOME=*Oracle_software_home_directory*

[SAPBACKUP_TYPE=OFFLINE]

Default is ONLINE.

SAPDATA_HOME=*database_files_directory*

SAPBACKUP=*BRTOOLS_logs_and_control_file_copy_directory*

SAPREORG=*BRSPACE_logs_directory*

OB2BARLIST=*backup_specification_name*

The backup specification is needed only to specify which Data Protector devices should be used for backup. Other information from the backup specification, like SAP R/3 objects to be backed up or the BRBACKUP options, is ignored and has to be specified manually at run time.

[OB2_3RD_PARTY_BACKINT=1]

Specifies usage of a third-party backint to perform ZDB to disk+tape backup. After setting the variable, copy the backint you want to use to the SAP BRTOOLS directory.

[OB2BARHOSTNAME=*application_system_name*]

Required if you are logged in to the backup system. Optional if you want to specify a virtual server name in cluster environments.

[OB2BACKUPHOSTNAME=*backup_system_name*]

Required if you are logged in to the application system.

OB2SMB=1

Specifies a ZDB session.

[OB2SMBIR=1]

Specifies tracking replica for instant recovery.

[ZDB_ORA_INCLUDE_CF_OLF=1]

Required if you are logged in to the backup system. For details, see [“ZDB integrations omnirc variables”](#) (page 271).

[OB2DISKONLY=1]

Specifies a ZDB-to-disk session.

If you are logged on to the backup system, ensure that the NLS_LANG environment variable is set to the same value as the NLS_LANG environment variable on the application system.

Alternatively, these variables can be specified in the backint parameter file. If this is required, the location of the file must be specified in the SAP configuration file using the *util_par_file* parameter:

util_par_file = *path\filename*

If you do not supply the path, the system searches for the parameter file in the directory:

UNIX: *ORACLE_HOME*/dbs

Windows: *SAPDATA_HOME*\database

3. Run the BRBACKUP command. The command syntax depends on whether you are logged in to the application system or backup system:

Application system:

```
brbackup -t {online_split | offline_split | online_mirror | \
offline_mirror} [-q split] -d \ util_file -m all -c -u user/password
```

Backup system:

```
brbackup -t {online_mirror | offline_mirror} [-q split] -d util_file
-m all -c -u user/password
```

The `-q split` option is required if `OB2DISKONLY` is set to 1.

Configuring SAP compliant ZDB sessions

SAP R/3 standards recommend that, in ZDB sessions that use the `splitint` backup interface, BRBACKUP is started on the backup system and not on the application system. You can configure Data Protector to comply with these standards by setting the Data Protector `OB2_MIRROR_COMP` environment variable to 1. The variable is saved in the Data Protector SAP R/3 instance configuration file. Consequently, in all `splitint` ZDB sessions for this SAP R/3 instance, BRBACKUP will be started on the backup system. By default, BRBACKUP is started on the application system.

Set the `OB2_MIRROR_COMP` environment variable using the Data Protector GUI or CLI.

NOTE: If no backup specification for the related SAP R/3 instance exists, you cannot use the Data Protector CLI to set the `OB2_MIRROR_COMP` variable.

Using the Data Protector GUI

You can set the `OB2_MIRROR_COMP` variable when you create a backup specification or modify an existing one:

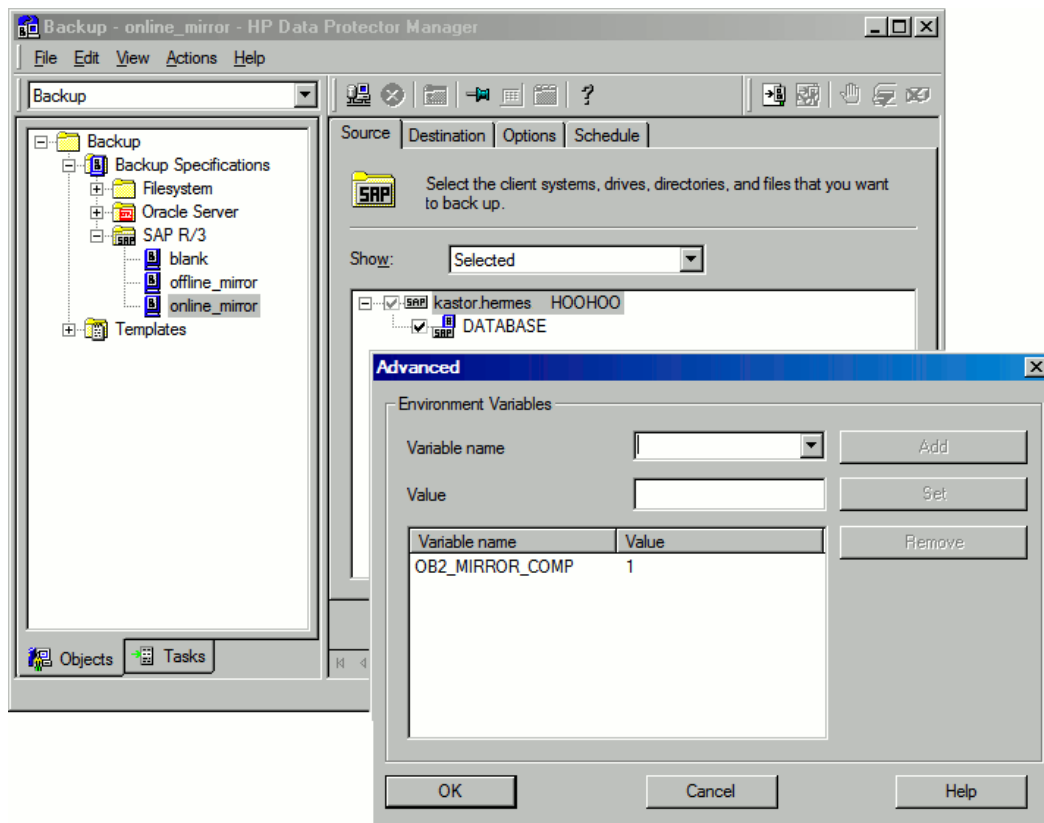
1. Proceed to the Source page of the backup specification.

NOTE: In environments in which the control file and datafiles reside on the same source disk, Data Protector does not let you proceed to the Source page if the **Track the replica for instant recovery** option is selected. Specifically, the Data Protector instant recovery check fails. In such a case, clear the option first. You can select the option later if needed, when the `OB2_MIRROR_COMP` variable is already set and, consequently, the instant recovery check is no longer performed.

Right click the SAP R/3 instance at the top and click **Set Environment Variables**.

2. In the Advanced dialog box, set `OB2_MIRROR_COMP` to 1. See “[Setting environment variables](#)” (page 135).
Click **OK**.

Figure 53 Setting environment variables



Using the Data Protector CLI

From the directory:

Windows: `Data_Protector_home\bin`

HP-UX, Solaris, and Linux: `/opt/omni/bin/`

Other UNIX systems: `/usr/omni/bin/`

run:

```
util_cmd -putopt SAP instance_name OB2_MIRROR_COMP 1 -sublist Environment
```

Manual balancing

Manual balancing means that you manually group files into subsets, which are then backed up in parallel. To group files into subsets, add the `manual_balance` section to the Data Protector SAP R/3 configuration file as described in the following example.

Example

Suppose that we have a backup specification named `SAP-R3` with the following files to be backed up: `fileA`, `fileB`, `fileC`, `fileD`. To group the files into three subsets (`0={fileA, fileC}`, `1={fileB}`, `2={fileD}`), add the following lines to the Data Protector SAP R/3 configuration file:

```
manual_balance={
  SAP-R3={
    fileA=0;
    fileB=1;
    fileC=0;fileD=2;}}
```

When you group files into subsets, consider the following:

- Use only one file from the same hard disk at a time.
- The number of files in a subset must be equal to or smaller than the sum of the concurrencies of all devices specified for backup.
- If the backup specification contains files that are not allocated to any subset, Data Protector automatically adds these files to the list of files to be backed up using the load balancing principle. Before the backup, this list is logged in:

UNIX: `ORACLE_HOME/sapbackup/*.lst`

Windows: `SAPDATA_HOME\sapbackup*.lst`

Restore

You can restore SAP R/3 objects using any of the following methods:

- **Standard restore:** Data is restored from backup media created in ZDB to tape, ZDB to disk+tape, and non-ZDB (standard backup) sessions. See [“Standard restore” \(page 137\)](#).
- **Instant recovery:** Data is restored from a replica created in *online* ZDB-to-disk or ZDB-to-disk+tape sessions. See [“Instant recovery” \(page 140\)](#).

After the restore, you can recover the database to a specific point in time using the SAP BRTOOLS interface. Instant recovery method enables you to restore and recover the database within the same session. However, you can only restore (and recover) the whole database. To restore only a part of the database or the archived logs, use the standard restore method.

[“SAP recovery methods” \(page 136\)](#) shows which restore methods are available, depending on the backup session you restore from.

Table 14 SAP recovery methods

Disk array	Backup types	Recovery of the whole database		Recovery of a part of the database
		Until now	To a point in time, logseq/thread or SCN number	
P9000 XP, P6000 EVA, EMC	ZDB to tape - online	Restore	Restore	Restore
	ZDB to tape - offline	Restore	Restore	Restore
P9000 XP, P6000 EVA	ZDB to disk - online	Instant recovery + database recovery	Instant recovery + database recovery	N/A
	ZDB to disk - offline	N/A	N/A	N/A
	ZDB to disk+tape - online	<ul style="list-style-type: none"> • Instant recovery + database recovery or • Restore 	<ul style="list-style-type: none"> • Instant recovery + database recovery or • Restore 	Restore
	ZDB to disk+tape - offline	Restore	Restore	Restore

Table 15 Legend

Restore	You can do a standard restore from the Data Protector media using the Data Protector GUI or the SAP BRTOOLS. After the restore, you can recover the database using the SAP BRTOOLS.
Instant recovery + database recovery	You can do an instant recovery. You can include database recovery in the instant recovery session or do it afterwards, using the SAP BRTOOLS.
N/A	Not available.

Considerations

- SAP R/3 tablespaces located on raw partitions cannot be restored using the Data Protector GUI. Workaround: Use SAP restore commands (for example, `brrestore`).
- If your Oracle database is localized, you may need to set the appropriate Data Protector encoding before you start a restore. For details, see [“Localized SAP R/3 objects” \(page 144\)](#).
- Restore preview is not supported.

Standard restore

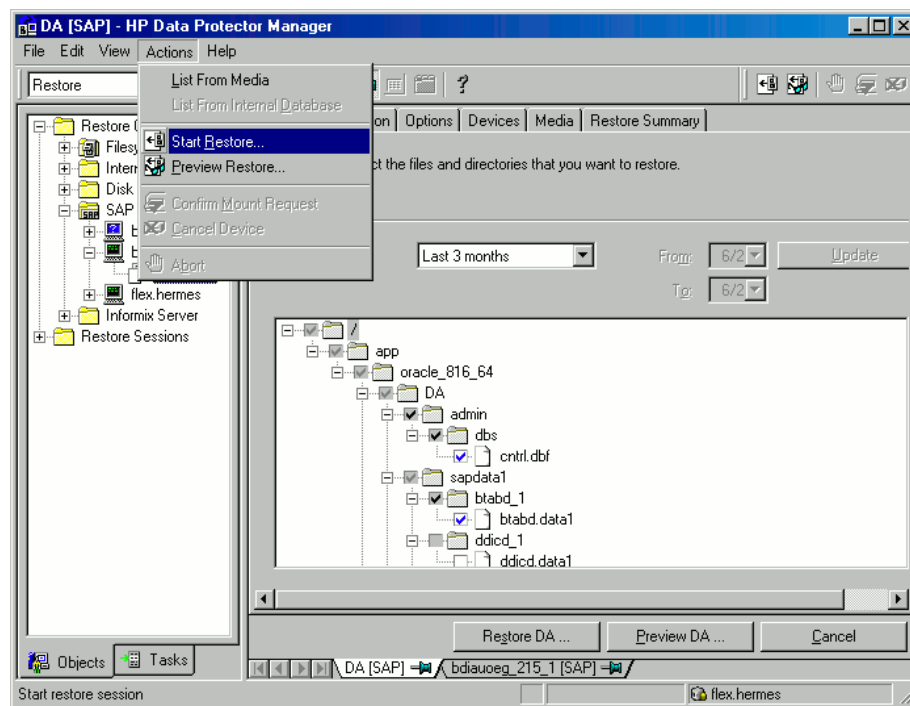
Restore SAP R/3 objects using the Data Protector Manager.

1. In the Context List, click **Restore**.
2. In the Scoping Pane, expand **SAP R/3**, expand the client (backup system) from which the data was backed up, and then click the Oracle instance you want to restore.
3. In the **Source** page, select SAP R/3 files to be restored.

To restore a file under a different name or to a different directory, right-click the file and click **Restore As/Into**.

To restore a file from a specific backup session, right-click the file and click **Restore Version**.

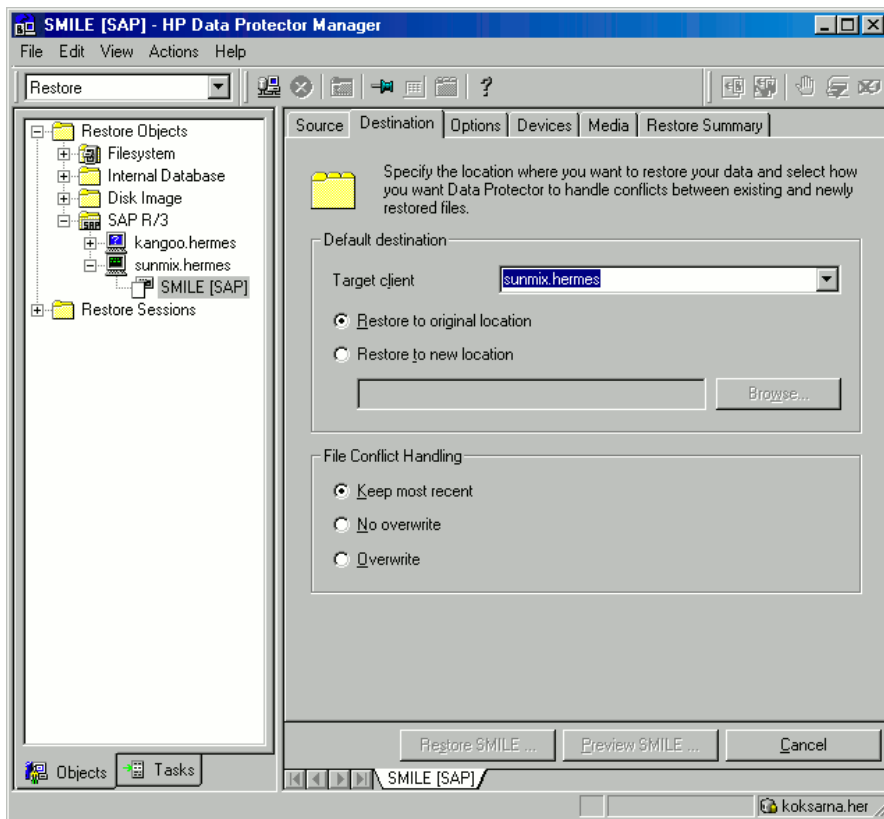
Figure 54 Selecting objects for restore



4. In the **Destination** tab, select the client to restore to (**Target client**). By default, this is the application system. See [“Selecting the target client” \(page 138\)](#).

For details on options, press **F1**.

Figure 55 Selecting the target client

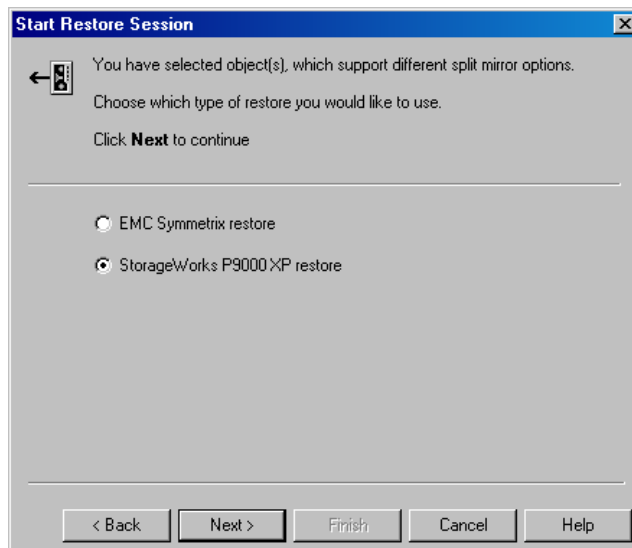


5. In the **Options** page, set the restore options. For information, press **F1**.
6. In the **Devices** page, select the devices to be used for the restore.
For more information on how to select devices for a restore, see the online Help index: “restore, selecting devices for”.
7. Click **Restore**.
8. In the **Start Restore Session** dialog box, click **Next**.
9. Specify **Report level** and **Network load**.
10. **EMC and P9000 XP Array:** This step is relevant only if you have both the EMC Symmetrix Agent and HP StorageWorks P9000 XP Agent components installed on the application system.

EMC: Select EMC Symmetrix restore.

P9000 XP Array: Select StorageWorks P9000 XP restore. See [“Selecting the StorageWorks P9000 XP restore” \(page 139\)](#).

Figure 56 Selecting the StorageWorks P9000 XP restore



Click **Next**.

11. **EMC and P9000 XP Array:** From the EMC Symmetrix mode or Mirror mode drop-down list, select **Disabled**. This sets the restore from backup media to the application system directly. See “[EMC – Selecting restore to the application system directly](#)” (page 139).

Figure 57 EMC – Selecting restore to the application system directly

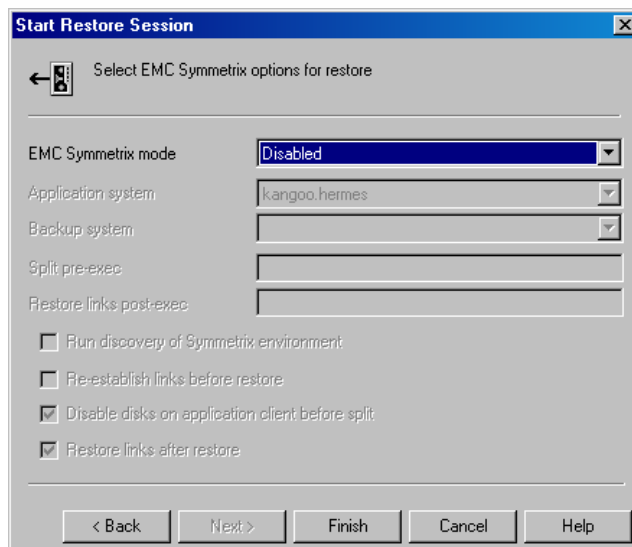
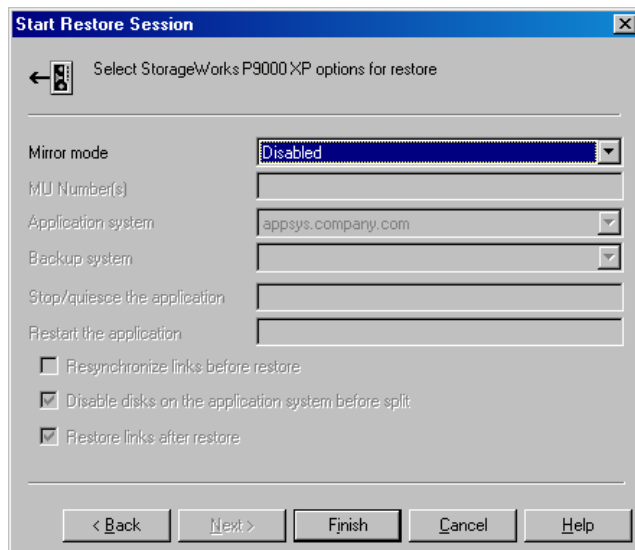


Figure 58 P9000 XP Array – Selecting restore to the application system directly



12. Click **Finish** to start the restore.

The message `Session completed successfully` is displayed at the end of a successful session.

Instant recovery

For general information on instant recovery, see *HP Data Protector Zero Downtime Backup Concepts Guide* and *HP Data Protector Zero Downtime Backup Administrator's Guide*. For information on instant recovery in cluster environment (Cluster File System (CFS), MC/ServiceGuard, and Microsoft Cluster Server), see *HP Data Protector Zero Downtime Backup Administrator's Guide*.

You can start an instant recovery using the Data Protector GUI or CLI.

Considerations

- The database recovery part is performed after the instant recovery procedure. During database recovery, archive log backups performed after the ZDB are restored from tape by the SAP BR*Tools utilities. If selected, the logs are reset and the database is opened.
- If the replica to be used for instant recovery contains the control file, first see [“Instant recovery from replicas containing the control file”](#) (page 143).

Instant recovery using the Data Protector GUI

To perform an instant recovery:

1. Shut down the Oracle database using `sqlplus`:

For example:

```
sqlplus
sqlplus> shutdown immediate
sqlplus> exit
```

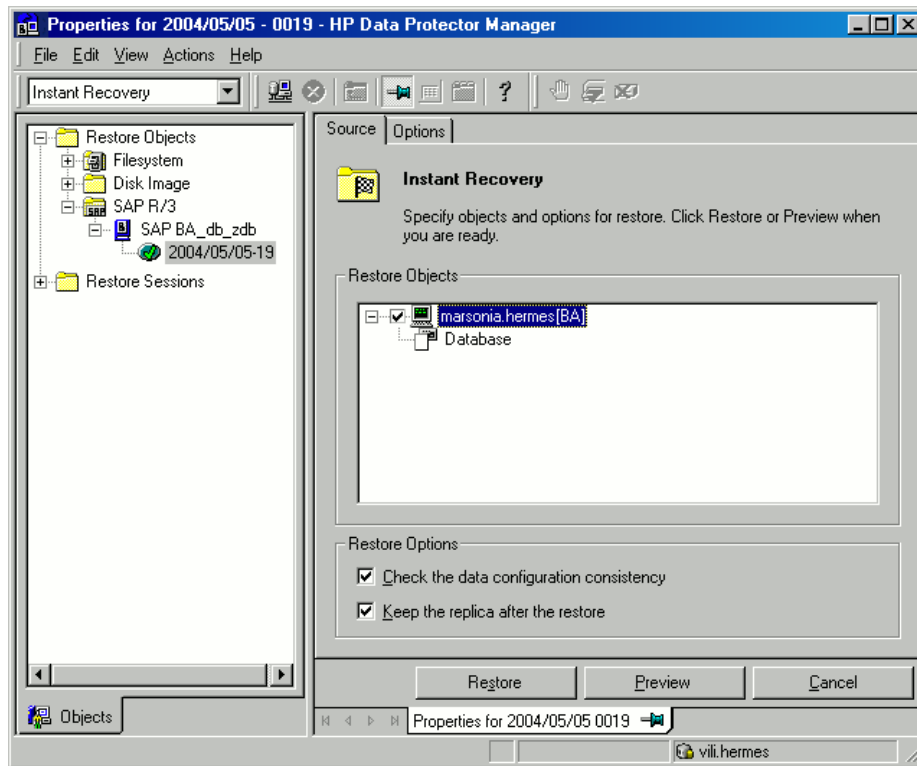
2. In the Data Protector Manager Context List, select **Instant Recovery**.
3. Expand **SAP R/3** and select the ZDB-to-disk or ZDB-to-disk+tape session from which you want to perform the restore.
4. In the **Source** tab, select the objects to recover. Only whole databases can be selected.

For HP P9000 XP Disk Array Family, it is recommended to leave the **Keep the replica after the restore** option selected to enable a restart of an instant recovery session. The option is selected by default, except for an offline backup where the database was in NOARCHIVELOG

mode during the backup. With HP P6000 EVA Disk Array Family, replica is kept on the disk array only if the **Copy replica data to the source location** is selected.

Set other HP P6000 EVA Disk Array Family or HP P9000 XP Disk Array Family options. For details, press **F1**.

Figure 59 SAP R/3 source options



5. At this point, you can decide whether to perform a database recovery immediately after an instant recovery or not:
 - To perform only an instant recovery, click **Restore**.
 - To automatically perform a database recovery after an instant recovery, select the recovery options. For details, see [“Database recovery options” \(page 141\)](#).

Click **Restore**.

Data Protector recovers the database after performing instant recovery by switching the database to mount state, restoring the necessary archive redo logs from tape, and applying the redo logs.

Database recovery options

User name (UNIX systems only)

Specifies the user name under which the instant recovery is performed. The user needs to be a member of the DBA group.

User group (UNIX systems only)

Specifies the user group the user in the **User name** field belongs to.

NOTE: The **User name** and the **User group** must be the same as defined in the backup ownership.

Recovery

Enables database recovery after instant recovery.

Recover until

The options in this drop-down list enables you to specify to which point in time you would like the recovery to be performed.

The following options are available:

Now: All existing archive logs are applied.

Selected time: Only archive logs until the specified time are applied.

Selected logseq/ thread number: Specifies an incomplete recovery. Only archive logs with a lower or equal number than the specified log sequence or thread number are applied.

Selected SCN number: Only archive logs until the specified SCN number are applied.

Open database after recovery

Opens the database after the recovery was performed.

Reset logs

Resets the archive logs after the database is opened. This option is by default not selected if the **Recover until** option is set to **Now**.

The following are Oracle recommendations on when to reset the logs:

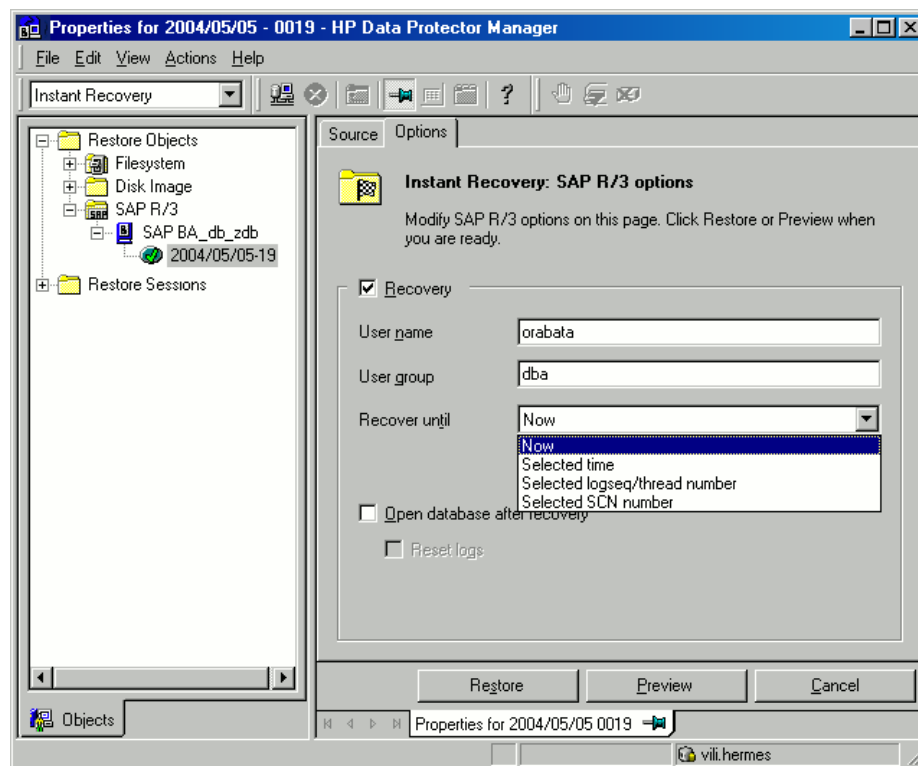
Always reset the logs:

- After an incomplete recovery, that is, if not all archive redo logs are applied.
- If a backup of the control file is used for recovery.

Do not reset the logs:

- After a complete recovery, when the control file is not used.
- If the archive logs are used for a standby database. However, if you must reset the archive logs, recreate the standby database.

Figure 60 SAP R/3 recovery options



Instant recovery using the Data Protector CLI

From the directory:

Windows: `Data_Protector_home\bin`

HP-UX: `/opt/omni/bin/`

```

run:
omnir
-host ClientName
-session SessionID
-instant_restore
[P9000_DISK_ARRAY_XP_OPTIONS | P6000_ENTERPRISE_VIRTUAL_ARRAY_OPTIONS]
-sap
-user UserName -group GroupName
-recover {now | time MM/DD/YY hh:mm:ss | logseq LogSeqNumber thread
ThreadNumber | SCN Number} [-open [-resetlogs]]
-appname ApplicationDatabaseName

```

The order of options is important. On Windows clients, the user name and group name options are not required. For a detailed description of the options, see the *HP Data Protector Command Line Interface Reference*.

Instant recovery from replicas containing the control file

During an instant recovery from a replica that contains the control file, the current control file and, possibly, online redo logs get overwritten. Therefore, before you start the session, copy the current control file and online redo logs to a safe location to be able to do a database recovery afterwards.

A replica contains the control file if it is created in any of the following sessions:

- Online ZDB session with the `omnirc` variable `ZDB_ORA_INCLUDE_CF_OLF` set to 1
- Online SAP compliant ZDB session
- Offline ZDB session (any configuration)

NOTE: An *offline* ZDB session also contains online redo logs. You can use such a session to restore the SAP R/3 database to a point in time at which the backup was performed. In this case, you do not need to follow the steps below.

To restore and recover the database:

1. Copy the current control files and online redo logs to a safe location.
2. Perform instant recovery (without database recovery). Use the Data Protector GUI or CLI.
3. Copy the current control files and online redo logs back to their original location.
4. Mount the target database.
5. Restore missing archived redo logs required for database recovery.

Example:

```

# sqlplus user/password@net_service_name
SQL> select SEQUENCE#, NAME from V$ARCHIVED_LOG where
(NEXT_TIME>to_date('2010/10/03','YYY/MM/DD')) and (FIRST_CHANGE#<='1000');
# brrestore -a log_no,... -d util_file -c force -u user/password

```

6. Recover the target database.

Example:

```

# rman target user/password@net_service_name
RMAN> run{
2> allocate channel dbrec type disk;
3> recover database until scn 1000;
4> release channel dbrec;
5> }

```

Restoring using another device

You can perform a restore using a device other than that used for the backup.

Using the Data Protector GUI

For information on how to select another device for a restore using the Data Protector GUI, see the online Help index: “restore, selecting devices for”.

Using the Data Protector CLI or SAP commands

If you are restoring using the Data Protector CLI or SAP R/3 commands, specify the new device in the file:

Windows: `Data_Protector_home\Config\Server\cell\restoredev`

UNIX: `/etc/opt/omni/server/cell/restoredev`

Use the format:

`"DEV 1" "DEV 2"`

where DEV 1 is the original device and DEV 2 the new device.

❗ **IMPORTANT:** Delete this file after use.

On Windows, use the Unicode format for the file.

Localized SAP R/3 objects

Oracle Server uses its own encoding, which may differ from the encoding used by the filesystem. In the Backup context, Data Protector displays the logical structure of the Oracle database (with Oracle names) and in the Restore context, the filesystem structure of the Oracle database. Therefore, to display non-ASCII characters correctly, ensure that the Data Protector encoding matches with the Oracle Server encoding during backup and with the filesystem encoding during restore. However, the incorrect display does not impact the restore.

UNIX: To be able to switch between the Data Protector encodings, start the GUI in UTF-8 locale.

Windows: If the current values of DBCS and the default Windows character set for non-Unicode programs do not match, problems arise. See [ZDB, restore, or instant recovery sessions fail due to invalid characters in filenames](#).

Monitoring sessions

You can monitor currently running sessions in the Data Protector GUI. When you run an interactive backup or a restore session, a monitor window shows you the progress of the session. Closing the GUI does not affect the session.

You can also monitor sessions from any Data Protector client with the `User Interface` component installed, using the Monitor context.

On how to monitor a session, see the online Help index: “viewing currently running sessions”.

System messages generated during backups are sent to both the SAP R/3 and the Data Protector monitor. However, mount requests are sent only to the Data Protector monitor.

Troubleshooting

This section lists general checks and verifications plus problems you might encounter when using the Data Protector SAP R/3 integration.

For general Data Protector troubleshooting information, see the *HP Data Protector Troubleshooting Guide*.

For general ZDB, restore, and instant recovery related troubleshooting, see the troubleshooting sections in the *HP Data Protector Zero Downtime Backup Administrator's Guide*.

See also the troubleshooting section in the SAP R/3 chapter of the *HP Data Protector Integration Guide*.

Before you begin

- Ensure that the latest official Data Protector patches are installed. See the online Help index: “patches” on how to verify this.
- See the *HP Data Protector Product Announcements, Software Notes, and References* for general Data Protector limitations, as well as recognized issues and workarounds.
- For an up-to-date list of supported versions, platforms, and other information, see the latest support matrices at <http://www.hp.com/support/manuals>.

General troubleshooting

Problem

Configuration fails due to a database operation failure

During configuration of an SAP R/3 database, Data Protector reports the following error:

Integration cannot be configured.

The database reported error while performing requested operation.

Action

Review user group membership for the user account which is used in Oracle database access authentication. For details, see “[Configuring user accounts](#)” (page 113).

Verifying the prerequisites (Oracle side)

Perform the following verification steps, in numerical order, to verify that Oracle is installed properly:

1. On the application system, verify that the target database is online, as follows:

UNIX:

```
export ORACLE_SID=Oracle_SID
export ORACLE_HOME=Oracle_home_path
$ORACLE_HOME/bin/sqlplus
```

Windows:

```
set ORACLE_SID=Oracle_SID
set ORACLE_HOME=Oracle_home_path
%ORACLE_HOME%\bin\sqlplus
```

At the SQLPlus prompt, type:

```
connect user/passwd@service
select * from dba_tablespaces
exit;
```

Try starting the target database.

2. In order to establish the TNS network connection, verify that Net8 software is configured correctly for the target database, as follows:
 - On the application system, perform the following:
UNIX:

```
$ORACLE_HOME/bin/lsnrctl status service
```


Windows:

```
%ORACLE_HOME%\bin\lsnrctl status service
```

If it fails, either start the TNS listener process or see the Oracle documentation on how to create the TNS configuration file (LISTENER.ORA).
 - On the application system, perform the following. Use sqlplus:
UNIX:

```
export ORACLE_SID=Oracle_SID  
export ORACLE_HOME=Oracle_home_path  
$ORACLE_HOME/bin/sqlplus
```


Windows:

```
set ORACLE_SID=Oracle_SID  
set ORACLE_HOME=Oracle_home_path  
%ORACLE_HOME%\bin\sqlplus
```

At the SQLPlus prompt, type:

```
connect user/passwd@service;  
exit;
```

If it fails, see the Oracle documentation on how to create the TNS configuration file (TSNAMES.ORA).

Verifying the prerequisites (SAP side)

Before you begin the steps in this section, be sure you have completed all the steps in [“Verifying the prerequisites \(Oracle side\)”](#) (page 145).

Perform the following verification steps, in numerical order, to verify that SAP is installed properly:

1. On the application system, verify a backup directly to disk, as follows:

```
brbackup -d disk -u user/password
```

If it fails, see the SAP Online Help for instructions on how to execute a backup to disk using the SAP backup utility.
2. On the application system, verify a restore from the disk, as follows:

```
brrestore -d disk -u user/password
```

If it fails, see the SAP Online Help for instructions on how to execute a restore to disk using the SAP restore utility.

3. On the application system, verify that SAP is configured properly, as follows:
Move the original `backint`. Create a test script with the name `backint` in the directory with the SAP backup utility, with the following entries:

```
#!/usr/bin/sh
echo "Test backint called as follows:"
echo "$0 $*"
echo "exiting 3 for a failure"
exit 3
```


Export all environment variables required by the SAP (`SAPDATA_HOME`, `SAPBACKUP...`) and then start the command with the backup owner user:
`brbackup -t offline_split -d util_file -u user/password -c`
or, if Data Protector uses `splitint`:
`brbackup -t offline_mirror -d util_file -u user/password -c`
If you receive arguments from `backint`, that means SAP is properly configured for backup using `backint`. Otherwise, you should reconfigure SAP.

Verifying the configuration

Before you begin this section, be sure that you completed all the steps provided in the sections [“Verifying the prerequisites \(Oracle side\)” \(page 145\)](#) and [“Verifying the prerequisites \(SAP side\)” \(page 146\)](#).

Perform the following verification steps, in numerical order, to verify that Data Protector is configured properly:

1. On the application system, verify a Data Protector filesystem backup of the SAP Database Server:

Perform a filesystem backup of the Oracle Server system so that you can eliminate any potential communication problems between the Oracle Server and the Data Protector Cell Manager system.

See the online Help index “standard backup procedure” for details about how to do a filesystem backup.

If it fails, see the *HP Data Protector Troubleshooting Guide* for help with troubleshooting a filesystem backup.
2. Verify the environment variable on the application system:

If you have to export some variables before starting the SAP backup utilities, Oracle Server Manager, or the TNS listener, set these environment variables using the Data Protector GUI.
3. Verify the permissions of the SAP user on application system:

SAP user permissions must be set to enable you to perform an SAP backup or restore with Data Protector. See [“Configuring user accounts” \(page 113\)](#). Use the `testbar2` to check the permissions:
 - Login in as an SAP user
 - Execute `/opt/omni/bin/testbar2 -perform:checkuser`
If the user account has all the required permissions, you will see only the usual messages displayed on the screen.
4. Examine the system errors:

System errors are reported in the following file on the Oracle Server:
`/var/opt/omni/log/debug.log`

Verifying the backup configuration

Before you begin this section, be sure that you completed all the steps provided in the sections “Verifying the prerequisites (Oracle side)” (page 145) and “Verifying the prerequisites (SAP side)” (page 146).

Perform the following verification steps, in numerical order, to verify that Data Protector is configured properly:

1. Verify the Data Protector SAP ZDB configuration on the application system:

Execute the following command:

HP-UX: `/opt/omni/sbin/util_sap -CHKCONF ORACLE_SID`

Windows: `Data_Protector_home\bin\util_sap.exe -CHKCONF ORACLE_SID`

If an error occurs, the error number is displayed in the form `*RETVAL*error_number`.

To get the error description, on the Cell Manager, run:

Windows: `Data_Protector_home\bin\omnigetmsg 12 error_number`

which is located on the Cell Manager.

HP-UX, Solaris: `/opt/omni/sbin/omnigetmsg 12 error_number`

2. Verify the SAP user.

Check that the respective user group has the `See Private Objects` user right selected. See also “Configuring user accounts” (page 113).

3. On the application system, verify the backup using `testbar2`:

Execute the following to ensure that communication within Data Protector is established:

- Create a non-ZDB backup specification on the application system.
- Run:

```
/opt/omni/bin/testbar2 -type:SAP -appname:ORACLE_SID \  
-perform:backup -file:file_name -bar barlist_name
```

If it fails, check the errors and try to fix them or call a support representative for assistance.

4. On the application system, verify the backup using `backint`:

Execute the following command to ensure that communication within Data Protector is established and that a backup of files can be performed:

- Create a non-ZDB backup specification on the backup system.
- ```
export OB2BARLIST=barlist_name
export OB2APPNAME=ORACLE_SID
/opt/omni/sbin/backint -f backup -t file -u ORACLE_SID -i \
input_file
```

  
where `input_file` is the file containing the full pathnames for backup.

If it fails, check the errors and try to fix them or call a support representative for assistance.

## Verifying restore

Before you begin this section, be sure that you completed all the steps provided in the sections “Verifying the prerequisites (Oracle side)” (page 145) and “Verifying the prerequisites (SAP side)” (page 146).

Perform the following verification steps, in numerical order, to verify that Data Protector is configured properly:

1. Verify the user for the restore

Verify that the user specified for the restore session is the user of the backup session and that they belong to the Data Protector operator or admin group. Check that the respective user group has the `See private objects` user right selected.

2. Verify that files are backed up and in the Data Protector database:

- Using the `omnidb` command;

See the appropriate man page on using the `omnidb` command.

- Using `backint`;

SAP tools also use this command to make a query.

```
/opt/omni/lbin/backint -f inquiry -u ORACLE_SID -i input_file
```

where *input\_file* is what will be queried. Backint expects a list of files in the following format:

```
backint_ID_1 pathName_1
```

```
backint_ID_2 pathName_2
```

```
backint_ID_3 pathName_3
```

To retrieve the *backint\_ID* numbers, enter the following command:

```
echo "#NULL #NULL" | backint -f inquiry -u ORACLE_SID
```

or, alternatively, you can just specify `#NULL` as *backint\_ID\_1* in the *input\_file*. In this case, the latest backup session for the file is used for the restore.

If it fails, proceed as follows:

- Check the backup session - was it successful?
- Check the user rights. Was the query started under the correct SAP user account?
- Call a support representative for assistance.

3. Verify the restore using Data Protector or CLI:

If it fails, proceed as follows:

- Check the backup session - was it successful?
- Check that the files are in the Data Protector database.
- Check the user rights. Was the restore started under the correct SAP user account?
- Call a support representative for assistance.

4. Verify the restore using `testbar2`:

Execute the following to ensure that restore is possible:

```
/opt/omni/bin/testbar2 -type:SAP -appname:ORACLE_SID \
-perform:restore -file:file_name -bar barlist_name -object objectName
```

If it fails, proceed as follows:

- Check the backup session - was it successful?
- Check that the files are in the Data Protector database.
- Check the user rights. Was the a restore started under the correct SAP user account
- Call a support representative for assistance.

5. Verify the restore using `backint`:

`backint` is the same command used by the SAP backup utility.

```
/opt/omni/lbin/backint -f restore -u ORACLE_SID -i input_file
```

where *input\_file* specifies what will be restored; `backint` expects the list of files in the following format:

```
backint_ID_1 pathName_1 [targetDirectory_1]
backint_ID_2 pathName_2 [targetDirectory_2]
backint_ID_3 pathName_3 [targetDirectory_3]
```

To retrieve the *backint\_ID* numbers, enter the following command:

```
echo "#NULL #NULL" | backint -f inquiry -u ORACLE_SID
```

or, alternatively, you can just specify `#NULL` as *backint\_ID\_1* in the *input\_file*. In this case, the latest backup session for the file is used for the restore.

If it fails, proceed as follows:

- Check the backup session - was it successful?
- Check that the files are in the Data Protector database.
- Check the user rights. Was the restore started under the correct SAP user account?
- Call a support representative for assistance.

## Configuration and backup problems

The following list gives a description of problems and actions to be taken to resolve them:

- **The Server Manager is unable to connect to the destination**

Check whether the Oracle TNS listener process is up and running. Check whether there are any environment variables required for a successful remote connection to the target database; for example, *TNS\_ADMIN* and *SHLIB\_PATH*. Set these environment variables using the Data Protector GUI.

For more information on the Data Protector SAP configuration file, see [“Data Protector SAP R/3 configuration file” \(page 107\)](#).

- **Configuration procedure fails**

Check whether the Oracle Server is up and running.

Check the login information for the target from the application system using Oracle Server Manager. If you cannot log in, then perform the following actions:

Check whether *sysoper* and *sysdba* rights are set for the Oracle administrator user.

Examine system errors reported in:

- On UNIX:  

```
/var/opt/omni/log/debug.log
/var/opt/omni/log/sap.log
/var/opt/omni/log/oracle8.log
```
- On Windows:  

```
Data_Protector_home\log\debug.log,
Data_Protector_home\log\sap.log
Data_Protector_home\log\oracle8.log
```

If you have special Oracle environment settings, ensure that they are registered in the Environment sublist of the Data Protector SAP configuration file:

/etc/opt/omni/server/integ/config/SAP/client\_name%ORACLE\_SID (UNIX Cell Manager), or Data\_Protector\_home\Config\server\integ\config\sap\client\_name%ORACLE\_SID (Windows Cell Manager).

For more information on the Data Protector SAP configuration file, [“Data Protector SAP R/3 configuration file” \(page 107\)](#).

- **Starting the backup fails**

On UNIX systems, check the output of the following command on the application system:

```
/opt/omni/sbin/util_sap.exe -CHKCONF ORACLE_SID
```

In case of an error, the error number is displayed in the form:

```
*RETVAL*Error_number
```

To get the error description, start the following command on the application system:

```
/opt/omni/sbin/omnigetmsg 12 Error_number
```

On Windows systems, perform the following procedure using the Data Protector GUI:

1. In the Context List, select **Backup**.
2. In the Scoping Pane, expand **Backup, Backup Specifications**, and then **SAP R/3**. A list of SAP backup specifications is displayed.
3. In the Scoping Pane, select the failed backup specification and right-click on the SAP R/3 server item in the Results Pane to display a pop-up menu.
4. From the pop-up menu, select **Check Configuration**.

A short description of the problems and how to resolve them is displayed.

- **Backup does not work**

- Check whether the Cell Manager is correctly set on the application system. The file /etc/opt/omni/client/cell\_server (UNIX systems) or HKEY\_LOCAL\_MACHINE\SOFTWARE\Hewlett-Packard\OpenView\OmniBackII\Site\CellServer (Windows systems) must contain the name of the Cell Manager.
- Check that the primary\_db parameter in the initORACLE\_SID.sap file on the application system is set to LOCAL.
- On UNIX systems, check whether the users are properly configured in user groups. Both the UNIX Oracle administrator (oraORACLE\_SID) and UNIX SAP administrator (ORACLE\_SIDadm) have to be in the Data Protector operator class.
- On UNIX systems, check whether the permissions of the SAPDATA\_HOME/sapbackup/ directory are set to 755.
- On Windows systems, check that the user account that started the Data Protector Inet service is added in the Data Protector operator class.

- **Backup fails with “Connect to database instance failed”**

If you start a backup while the database instance is in the unmount or mount mode, the session fails with a message similar to the following:

```
BR0301E SQL error -1033 at location BrDbConnect-2
```

```
ORA-01033: ORACLE initialization or shutdown in progress
```

```
BR0310E Connect to database instance HOOHOO failed
```

Before you start a backup, ensure that the database instance is the open or shutdown mode.

## Problem

### Configuration fails due to a script failure

During configuration of an SAP R/3 database, Data Protector reports the following error:

Integration cannot be configured.

Script failed. Cannot get information from remote host.

#### Action

If the SAP R/3 database is located on a Windows system, check the environment settings and ensure Data Protector Inet is running under a user account which has the required privileges. For details, see [“Before you begin” \(page 112\)](#).

If the SAP R/3 database is located on a UNIX system, resolve the problem by reviewing the user account configuration. For details, see [“Configuring user accounts” \(page 113\)](#).

#### Problem

##### **Backup using backint fails on Solaris**

A backup using backint fails on Solaris systems with the following error:

```
[Major] From: OB2BAR_DMA@computer.company.com "SAP" Time: 4/29/09 3:55:52 PM
```

```
Cannot open file '/saphome/SAP/sapbackup/cntrlSAP.dbf'. Error: 13
```

#### Action

Share the directories on the application system through NFS with root permissions. On the application system, add the following line in the file `/etc/dfs/dfstab`:

```
share -F nfs -o anon=0 /usr/src
```

## Restore problems

#### Problem

##### **ZDB, restore, or instant recovery sessions fail due to invalid characters in filenames**

On Windows systems, where the Oracle Database Character Set (DBCS) is not set to the same value as the default Windows character set for non-Unicode programs, and where SAP tools are used to create Oracle datafiles, ZDB, restore, and instant recovery fail if the datafiles contain non-ASCII or non-Latin 1 characters.

#### Actions

Use any of the following solutions:

- For new Oracle installations, set the DBCS to UTF-8.
- If you do not use other non-Unicode programs, set the language for non-Unicode programs to the same value as DBCS.
- Do not use non-ASCII or non-Latin 1 characters for filenames.

#### Problem

##### **Restore of SAP R/3 tablespaces located on raw partitions fails**

When restoring SAP tablespaces that are located on raw partitions using the Data Protector GUI, the restore fails with a message similar to the following:

```
[Major] From: VRDA@joca.company.com "SAP" Time: 5/9/06 3:33:51 PM
/dev/sapdata/rsapdata Cannot restore -> rawdisk section ! [Warning] From:
VRDA@joca.company.com "SAP" Time: 5/9/06 3:42:45 PM Nothing restored.
```

#### Action

Use SAP commands (for example, `brrestore`) to restore these tablespaces.



# 3 Data Protector Microsoft SQL Server ZDB integration

## Introduction

This chapter explains how to configure and use the Data Protector Microsoft SQL Server ZDB integration. It describes the concepts and methods you need to understand to back up and restore the Microsoft SQL Server (**SQL Server**) database objects.

During the backup, an SQL Server snapshot is made (the database files are frozen and any transactions to them are cached), so the database is highly available (*online backup*). The I/O to it is suspended during the time it takes to create a **replica** (split the mirror disks or create snapshots).

**NOTE:** SQL Server snapshot is an SQL Server related term and does not mean the same as a disk array snapshot.

The following disk arrays and array configurations are supported:

| Supported array                                  | Supported configurations                                   |
|--------------------------------------------------|------------------------------------------------------------|
| EMC Symmetrix (EMC)                              | Dual-host TimeFinder                                       |
| HP P9000 XP Disk Array Family (P9000 XP Array)   | HP BC P9000 XP, HP CA P9000 XP, combined HP CA+BC P9000 XP |
| HP P6000 EVA Disk Array Family (P6000 EVA Array) | HP BC P6000 EVA, combined HP CA+BC P6000 EVA               |

All ZDB types (ZDB to tape, ZDB to disk, and ZDB to disk+tape) are supported by this integration. For ZDB types description, see the *HP Data Protector Zero Downtime Backup Concepts Guide* and the online Help.

Using Data Protector, you can restore your SQL Server data:

- From backup media to the application system on LAN (standard restore).
- Using the instant recovery functionality.

The following table gives an overview of SQL Server recovery methods:

| ZDB type         | Recovery method                    |
|------------------|------------------------------------|
| ZDB to tape      | Standard restore                   |
| ZDB to disk      | Instant recovery                   |
| ZDB to disk+tape | Standard restore, instant recovery |

For a description of ZDB and instant recovery concepts, see the *HP Data Protector Zero Downtime Backup Concepts Guide*.

## Integration concepts

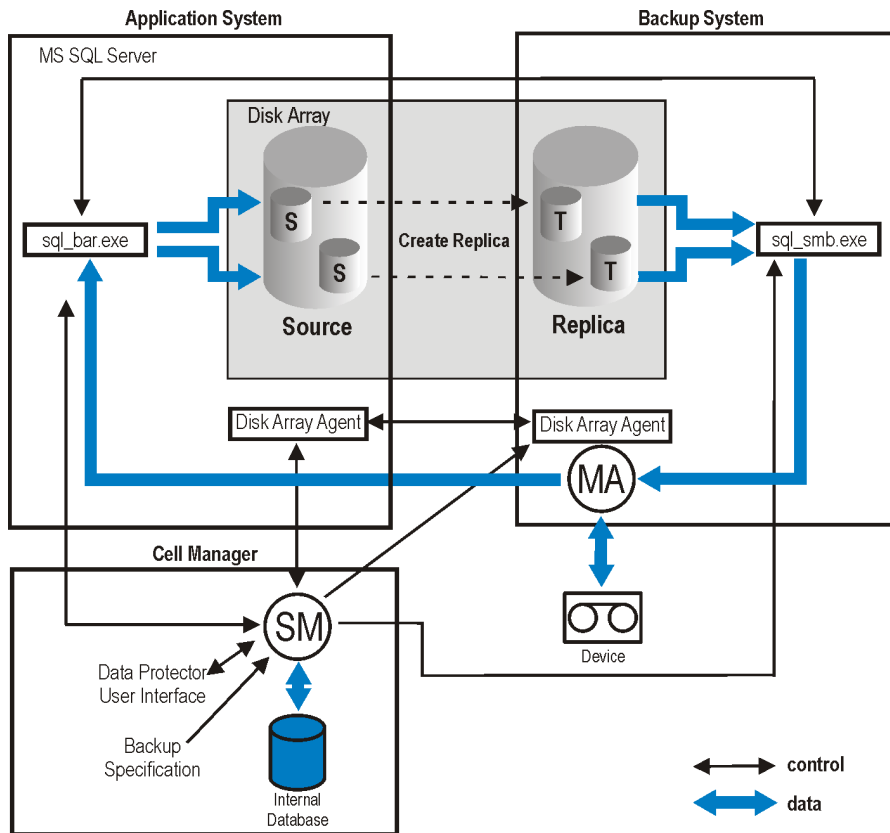
Data Protector integrates with SQL Server through the Data Protector `sql_bar.exe` executable, installed on SQL Server. During backup, `sql_bar.exe`, started on the application system, connects to SQL Server to find the locations of the database files. The integration then backs up SQL Server database(s), which are replicated within a disk array.

During restore, `sql_bar.exe` connects to SQL Server to receive the restore data, which is then written to disks.

ZDB process depends on whether you replicate your data in a split mirror or snapshot replication and on the selected ZDB type. Restore process depends on the restore type - standard restore or

instant recovery. See the *HP Data Protector Zero Downtime Backup Concepts Guide* for a detailed description of replication techniques and ZDB and restore processes.

**Figure 61 Backup and restore concepts**



## Configuring the integration

### Prerequisites

- You need a license to use the SQL Server ZDB integration. For information, see the *HP Data Protector Installation and Licensing Guide*.
- Ensure that you correctly installed and configured SQL Server.
  - For supported versions, platforms, devices, and other information, see the latest support matrices at <http://www.hp.com/support/manuals>.
  - For information on installing, configuring, and using SQL Server, see the SQL Server documentation.
- Ensure that you correctly installed Data Protector. For information on installing Data Protector in various architectures and installing a Data Protector disk array integration (P6000 EVA

Array, P9000 XP Array, or EMC) with SQL Server, see the *HP Data Protector Installation and Licensing Guide*.

Every SQL Server to be used with Data Protector must have the MS SQL Integration component installed.

- Install SQL Server on the application system. Install user databases on the disk array source volumes (system databases can be installed anywhere). If the system databases are also installed on a disk array, they *must* reside on *different* source volumes than user databases. If SQL Server is installed on the backup system as well, its databases *must* reside on source volumes *different* from the source volumes used for this integration. Drive letters/mount points assigned to those volumes must also be different from the drive letters/mount points assigned to the volumes used for this integration.

## Before you begin

- Configure devices and media for use with Data Protector. For instructions, see the online Help index: “configuring devices” and “creating media pools”.
- On Windows XP and Windows Server 2003 systems, if you plan to use **Integrated authentication** to connect to an SQL Server instance, you need to restart the Data Protector Inet service under a Windows domain user account that has the appropriate SQL Server permissions for running backups and restores. For information on changing the user account under which the Data Protector Inet service is running, see the online Help index: “Inet, changing account”.

However, for other supported Windows operating systems, you can use user impersonation instead. For details on setting accounts for the Inet service user impersonation, see the online Help index: “Inet user impersonation”.

- Using the SQL Server Management Studio, add the user account which you will use for backing up and restoring SQL Server data to the fixed server role sysadmin. For instructions, see the SQL Server documentation.
- To test whether SQL Server and Cell Manager communicate properly, configure and run a Data Protector filesystem ZDB and restore. For instructions, see the online Help.

## Data Protector SQL Server configuration file

Data Protector stores integration parameters for every configured SQL Server on the Cell Manager in:

### **HP-UX, Solaris, and Linux systems:**

`/etc/opt/omni/server/integ/config/MSSQL/client_name%instance_name`

### **Windows systems:**

`Data_Protector_home\Config\Server\Integ\Config\MSSQL\client_name% \instance_name`

Configuration parameters are the username and password of the SQL Server user, who must have permissions to run backups and restores within SQL Server (assuming the standard security is used). They are written to the Data Protector SQL Server configuration file during configuration of the integration.

The content of the configuration file is:

```
Login='user';
Password='encoded_password';
Domain='domain';
```

- 
- ❗ **IMPORTANT:** To avoid backup problems, ensure that the syntax of your configuration file matches the examples.
-

## Examples

- **SQL Server authentication:**

```
Login='sa';
Domain='';
Password='jsk74yh80fh43kdf';
```

- **Windows authentication:**

```
Login='Administrator';
Domain='IPR';
Password='dsjff08m80fh43kdf';
```

- **Integrated authentication:**

```
Login='';
Domain='';
Password='kf8u3hdgtfh43kdf';
```

## Configuring users

If you have restarted the Data Protector `Inet` service on the SQL Server system under a different user account, add this user to the Data Protector `admin` or `operator` Data Protector user group. For information on adding users to the Data Protector groups, see the online Help index: “adding users”.

## Configuring an SQL Server cluster

In a cluster, all the nodes must be installed as Data Protector cluster-aware clients and the Data Protector `Inet` service on all nodes must run under a Windows domain user account that has also cluster administrator rights.

You must configure the Data Protector `Inet` service user impersonation for all cluster nodes. The Windows domain user account that is used must be given the following Windows operating system Security Policy privileges:

- Impersonate a client after authentication
- Replace a process level token

For more information, see the online Help index: “cluster-aware client”, “Inet user impersonation”, and the SQL Server cluster documentation.

## Configuring SQL Server instances

An SQL Server instance is configured during the creation of the first backup specification. The configuration consists of setting the user account that Data Protector should use to connect to the SQL Server instance. The specified login information is saved to the Data Protector SQL Server instance configuration file on the Cell Manager.

---

**NOTE:** Ensure that the user account to be used has appropriate SQL Server permissions for running backups and restores. Check the permissions using SQL Server Enterprise Manager.

---

You can change configuration by following instructions described in [“Changing and checking configuration”](#) (page 158).

### Prerequisites

- SQL Server must be online during configuration.
- Configuration must be performed for every SQL Server instance separately.

## Using the Data Protector GUI

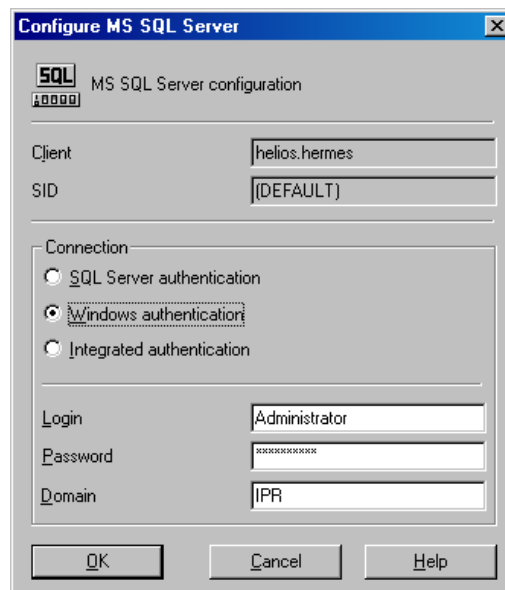
1. In the Context List, click **Backup**.

2. In the Scoping Pane, expand **Backup Specifications**, right-click **MS SQL Server**, and click **Add Backup**.
3. In the **Create New Backup dialog box**, select the **Blank Microsoft SQL Server Backup** template and specify backup type. For details, see [Step 3](#).  
Click **OK**.
4. Specify the ZDB-specific options. For details, see [Step 4](#).
5. In **Application database**, select or specify the name of the SQL Server instance.  
**Windows Server 2008 only:** If you intend to use the **Integrated authentication** option, specify the **User and group/domain** options. For information on the options, press **F1**.  
Click **Next**.
6. In the Configure MS SQL Server dialog box, specify the user account that Data Protector should use to connect to the SQL Server instance.
  - **SQL Server authentication:** SQL Server user account. Specify a username and password.
  - **Windows authentication:** Windows domain user account (preferred option). Specify a username, password, and the domain.
  - **Integrated authentication:** Select this option to enable Data Protector to connect to the SQL Server instance with the following Windows domain user account:
    - **Windows Server 2008:** The account specified in the **User and group/domain** options in the previous step or in the Client selection page.
    - **Other Windows systems:** The account under which the Data Protector Inet service on the SQL Server system is running.

Ensure that the user account you specify has the appropriate permissions for backing up and restoring the SQL Server databases.

See [“Configuring SQL Server” \(page 157\)](#).

**Figure 62** Configuring SQL Server



**NOTE:** It is recommended that the SQL Server system administrator configures the integration.

For details about security, see the SQL Server documentation.

Click **OK** to confirm the configuration.

7. The SQL Server instance is configured. Exit the GUI or proceed with creating the backup specification at [Step 7](#).

## Using the Data Protector CLI

From the *Data\_Protector\_home\bin* directory, run:

```
sql_bar config [-appsrv:SQL_Server_client] [-instance:instance_name]
[-dbuser:SQL_Server_user -password:password | -dbuser:Windows_user
-password:password -domain:domain]
```

### Parameter description

*-appsrv:SQL\_Server\_client*

Client system on which the SQL Server instance is running. This option is not required if you run the command locally.

*-instance:instance\_name*

SQL Server instance name. If you omit this option, the default SQL Server instance is configured.

*-dbuser:SQL\_Server\_user -password:password*

SQL Server user account (**SQL Server authentication**)

*-dbuser:Windows\_user -password:password -domain:domain*

Windows domain user account (**Windows authentication**)

---

**NOTE:** If no user account is specified, Data Protector uses **Integrated authentication**.

---

The message \*RETVAL\*0 indicates successful configuration.

## Changing and checking configuration

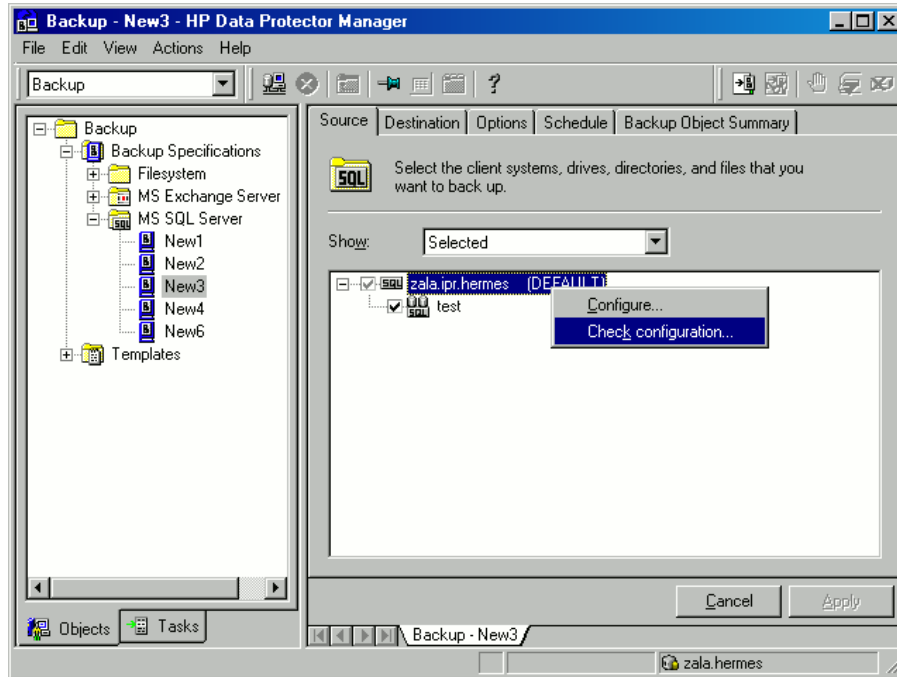
You can check and change configuration using the Data Protector GUI or CLI.

### Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **MS SQL Server**. Click a backup specification for which you want to change the configuration.
3. In the **Source** property page, right-click the SQL Server name and select **Configure**.
4. Configure SQL Server as described in [“Configuring SQL Server instances”](#) (page 156).

5. Right-click SQL Server and select **Check Configuration**. See “Checking configuration” (page 159).

**Figure 63 Checking configuration**



## Using the Data Protector CLI

To change the configuration, run the command for configuring SQL Server instances again, entering different data.

To check configuration, run:

```
sql_bar chkconf [-instance:instance_name]
```

If the optional parameter `-instance:instance_name` is not specified, the default instance is checked.

If the integration is not properly configured, the command returns:

```
*RETVAL*8523
```

To get the information about the existing configuration, run:

```
sql_bar getconf [-instance:instance_name]
```

If `-instance:instance_name` is not specified, Data Protector returns configuration for the default instance.

## Backup

To run ZDB of an existing SQL Server ZDB specification:

- Schedule a backup using the Data Protector Scheduler.
- Start an interactive backup using the Data Protector GUI or CLI.

## Prerequisites

- In case of nested mount points, the same drive letters, on which the source volumes to be replicated reside on the application system, must exist on the backup system to enable successful mounting of the target volumes. If the same drive letters do not exist on the backup system, the backup fails.

## Considerations

Your session will fail if:

- You start ZDB, restore, or instant recovery using the same source volume on the application system at the same time. A session must be started only after the preceding session using the same source volume on the application system finishes.
- SQL Server services are not running when the backup starts.

To configure a ZDB, create a Data Protector SQL Server ZDB specification.

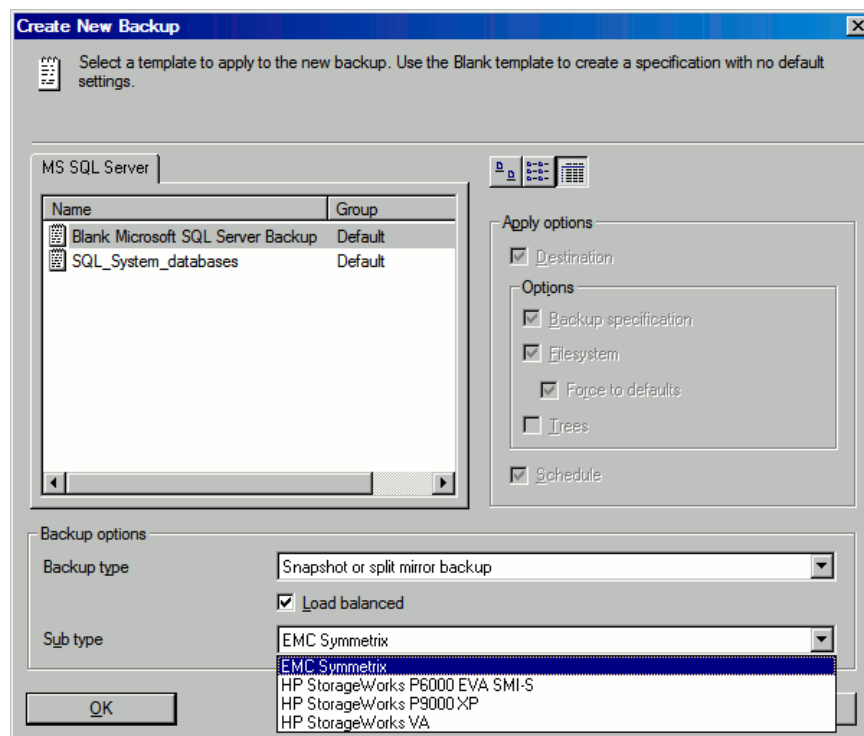
## Creating ZDB specifications

Create a ZDB specification, using the Data Protector Manager.

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **MS SQL Server**, and click **Add Backup**.
3. In the **Create New Backup** dialog box, select the **Blank Microsoft SQL Server Backup** template.

From the **Backup type** drop-down list, select **Snapshot or split mirror backup**, and from the **Sub type** drop-down list, select the appropriate disk array agent. The agent must be installed on the application system and the backup system. See [“Selecting a Microsoft SQL Server backup template and the snapshot or split mirror backup”](#) (page 160).

**Figure 64** Selecting a Microsoft SQL Server backup template and the snapshot or split mirror backup



Click **OK**.



4. Under **Client systems**, select the SQL Server system. In cluster environments, select the virtual server of the SQL Server resource group.

In the **Backup system** drop-down list, select the backup system.

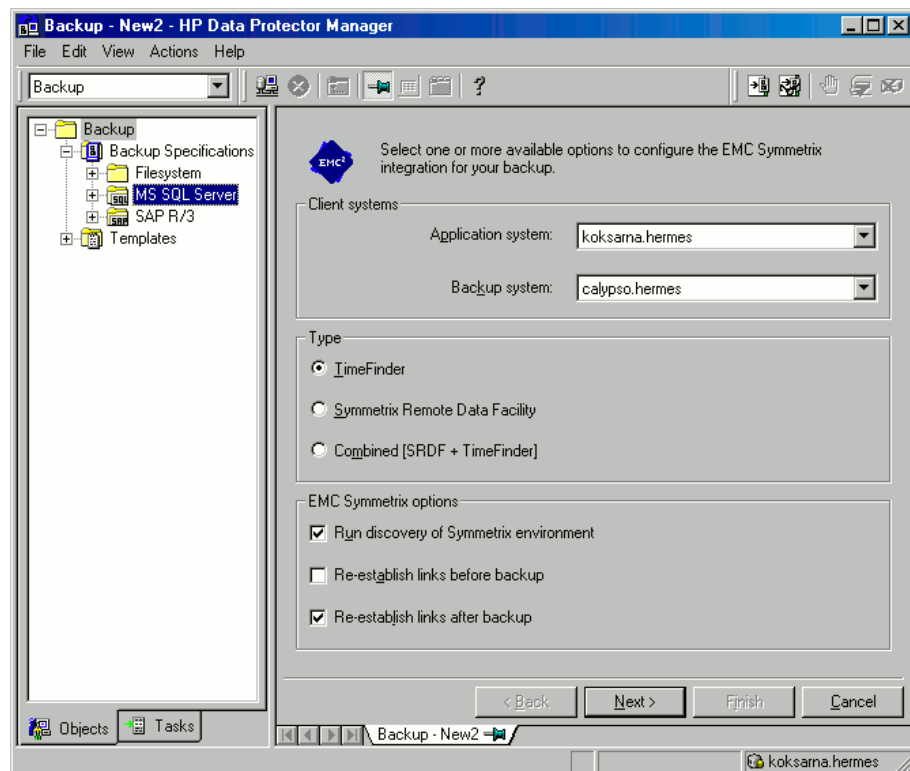
Select other disk array-specific backup options (see for EMC, “P9000 XP Array backup options” (page 162) for P9000 XP Array, or “P6000 EVA Array backup options” (page 162) for P6000 EVA Array). For detailed information on the backup options, press **F1**.

#### **EMC:**

In the EMC GeoSpan for Microsoft Cluster Service environment, select the backup system for the active node and specify the TimeFinder configuration.

After a failover in EMC GeoSpan for MSCS, select the backup system for the currently active node and save the backup specification.

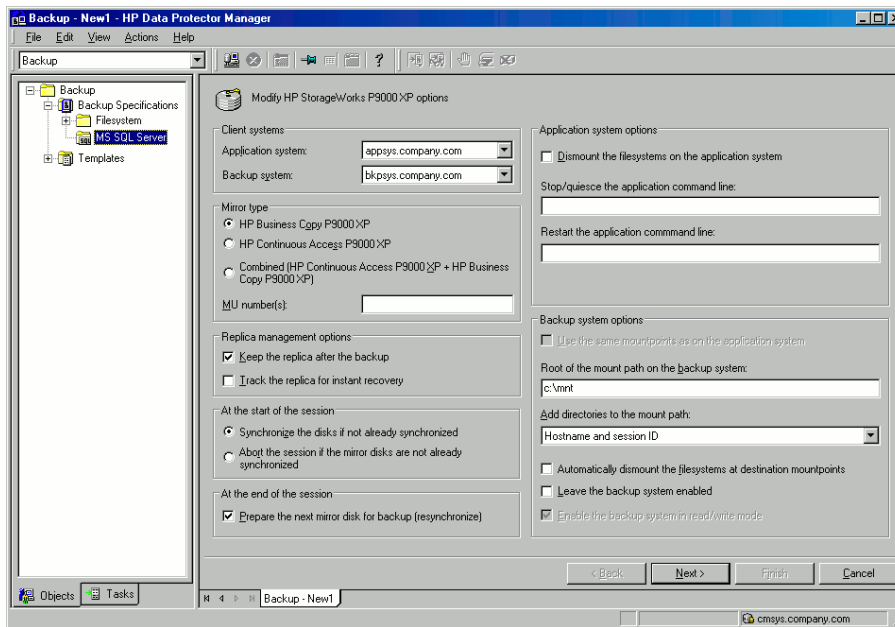
**Figure 65 EMC backup options**



#### **P9000 XP Array:**

To enable instant recovery, leave **Track the replica for instant recovery** selected.

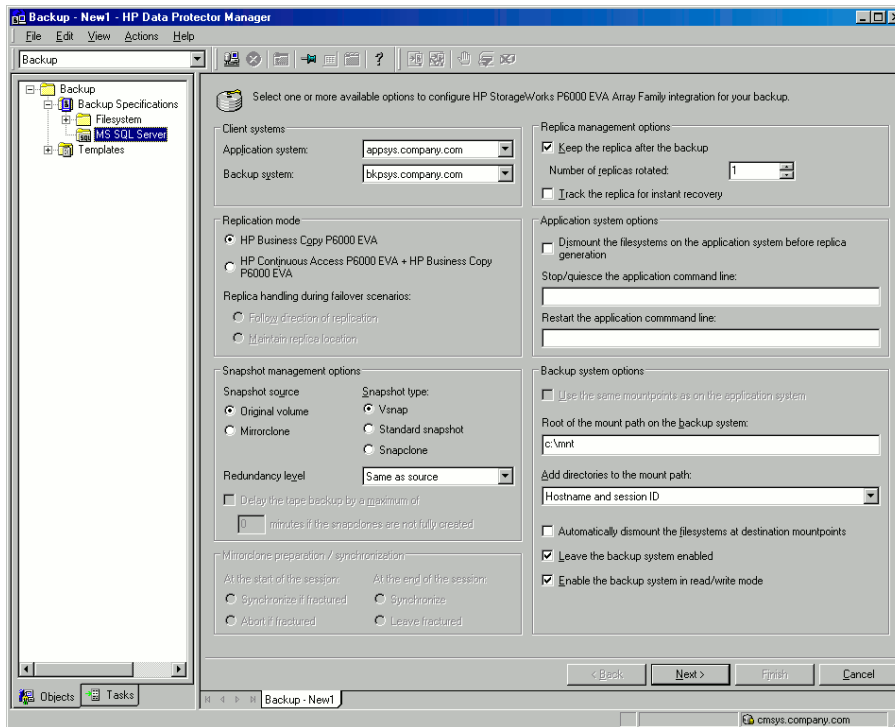
**Figure 66 P9000 XP Array backup options**



### **P6000 EVA Array:**

To enable instant recovery, select **Track the replica for instant recovery**.

**Figure 67 P6000 EVA Array backup options**



Click **Next**.

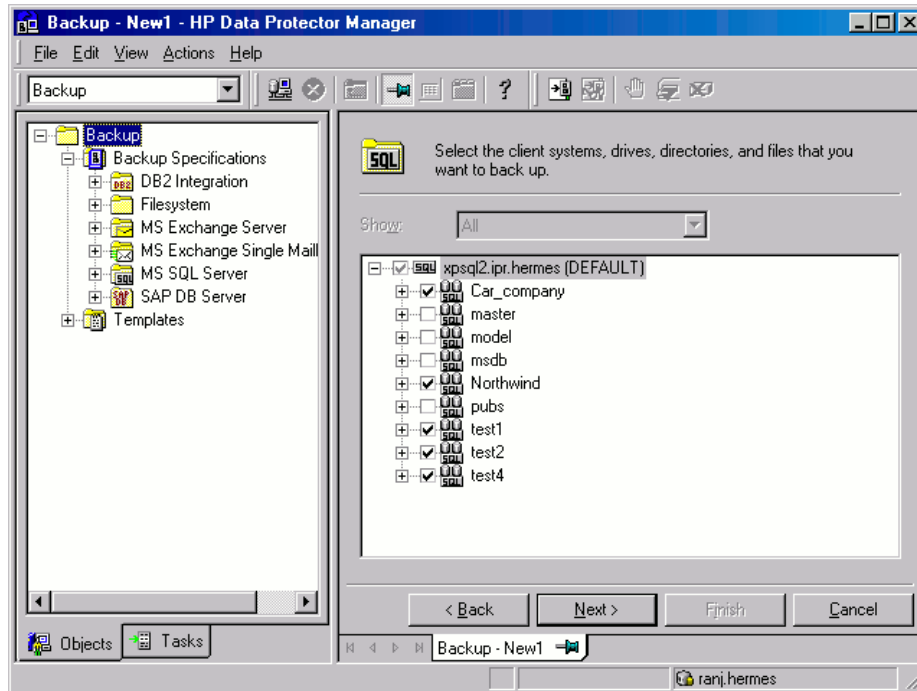
5. In **Application database**, specify the name of the SQL Server instance.

**Windows Server 2008 only:** If you intend to use the **Integrated authentication** option, specify the **User and group/domain** options. For information on the options, press **F1**.

Click **Next**.

6. If the client is not configured, the **Configure MS SQL Server** dialog box appears. Configure it as described in [“Configuring SQL Server instances”](#) (page 156).
  7. Select the databases to be backed up.
- 
- ❗ **IMPORTANT:** To enable instant recovery, create different backup specifications for user and system databases.
- 

**Figure 68** Selecting user databases



Click **Next**.

8. Select the devices. Click **Properties** to set the media pool and preallocation policy. The device concurrency is set to 1 and cannot be changed. For more information on options, press **F1**.  
To create additional backup copies (mirrors), specify the desired number by clicking **Add mirror/Remove mirror**. Select separate devices for each mirror. The minimum number of devices for mirroring equals the number of devices used for backup.  
For more information on object mirroring, see the online Help.
- 

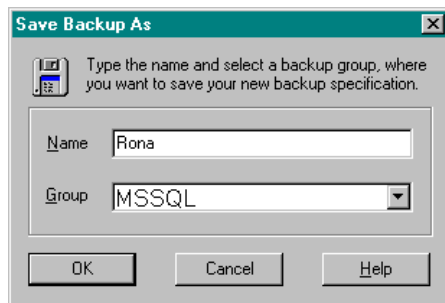
**NOTE:** Object mirroring is not supported for ZDB to disk.

---

Click **Next**.

9. Select backup options.  
For information on **Backup Specification Options** and **Common Application Options**, see the online Help.  
For information on **Application Specific Option**, see [“SQL Server-specific backup options”](#) (page 164).  
Click **Next**.
10. Optionally, schedule the backup. For information on scheduler, press **F1**.  
Note that only **Full** backup is performed.
11. Save the backup specification, specifying a name and backup specification group. You start the backup specification by clicking **Start Backup**.

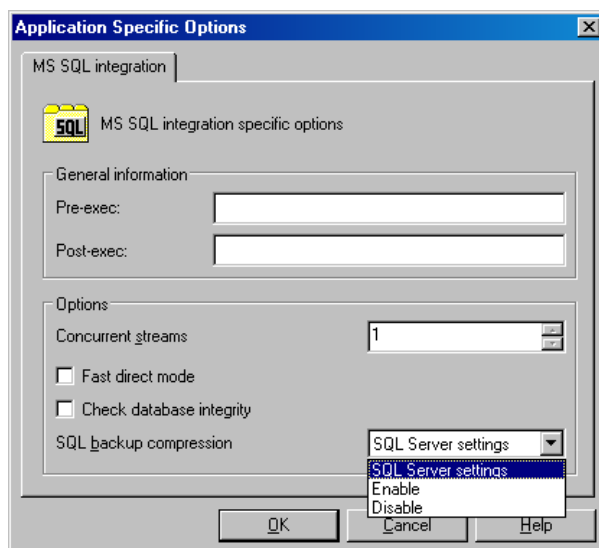
**Figure 69 Saving a backup specification**



## SQL Server-specific backup options

SQL Server-specific backup options are specified by clicking the **Advanced** tab in the **Application Specific Options** group box.

**Figure 70 Application-specific options**



**Table 16 SQL Server backup options**

|                                 |                                                                                                                                                                                                                                                |                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------|
| <b>Pre-exec</b>                 | Specifies a command with arguments or a script started by <code>sql_bar.exe</code> on SQL Server before backup. Resides in the <code>Data_Protector_home\bin</code> directory. Only the filename must be provided in the backup specification. |                                                                                |
| <b>Post-exec</b>                | Specifies a command with arguments or a script started by <code>sql_bar.exe</code> on SQL Server after backup. Resides in the <code>Data_Protector_home\bin</code> directory. Only the filename must be provided in the backup specification.  |                                                                                |
| <b>Concurrent streams</b>       | Sets the number of concurrent streams used to back up SQL Server databases from the replica to tape. Applicable for ZDB-to-tape and ZDB-to-disk+tape sessions.                                                                                 |                                                                                |
| <b>Fast direct mode</b>         | Ignored for ZDB sessions.                                                                                                                                                                                                                      |                                                                                |
| <b>Check database integrity</b> | Performs data integrity validation before backup. If the check fails, the session completes with warnings.                                                                                                                                     |                                                                                |
| <b>SQL backup compression</b>   | Specify how Data Protector should handle the Microsoft SQL Server backup compression.                                                                                                                                                          |                                                                                |
|                                 | <b>SQL Server settings</b> (default)                                                                                                                                                                                                           | Handles the backup compression according to the Microsoft SQL Server settings. |

**Table 16 SQL Server backup options** *(continued)*

|  |                |                                                                                                               |
|--|----------------|---------------------------------------------------------------------------------------------------------------|
|  | <b>Enable</b>  | Executes the backup compression regardless of the Microsoft SQL Server settings.                              |
|  | <b>Disable</b> | Specifies that the backup compression should not be executed regardless of the Microsoft SQL Server settings. |

**NOTE:** Do not use double quotes (" ") in object-specific pre-exec and post-exec commands.

## Scheduling backups

You can run unattended ZDB at specific times or periodically. For details on scheduling, see the online Help index: "scheduled backups".

**NOTE:** You cannot run ZDB to disk or ZDB to disk+tape if **Track the replica for instant recovery** is not selected in the backup specification.

## Scheduling example

To schedule a database ZDB at 8:00, 13:00, and 18:00 during week days:

1. In the **Schedule** property page, select the starting date in the calendar and click **Add** to open the **Schedule Backup** dialog box.
2. Under **Recurring**, select **Weekly**. Under **Time options**, select **8:00**. Under **Recurring Options**, select **Mon, Tue, Wed, Thu, and Fri**.  
Click **OK**.
3. Repeat [Step 1](#) and [Step 2](#) to schedule backups at 13:00 and 18:00.
4. Click **Apply** to save the changes.

**NOTE:** For ZDB sessions, the backup type is set to **Full**.

## Starting backup sessions

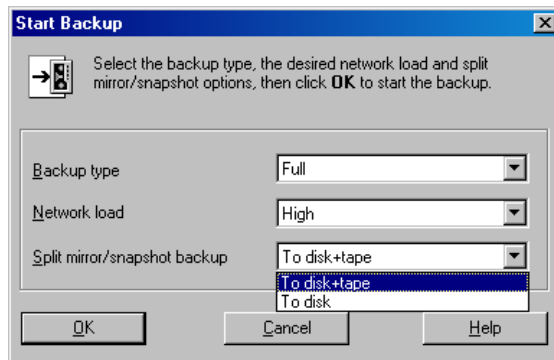
Interactive backups are run on demand. They are useful for urgent backups or restarting failed backups.

## Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, and then **MS SQL Server**. Right-click the backup specification you want to use and select **Start Backup**.

3. Select **Network load**. For information on network load, click **Help**. Click **OK**.  
For ZDB sessions, the backup type is set to **Full**.  
For ZDB to disk or ZDB to disk+tape, specify the **Split mirror/snapshot backup** option.

**Figure 71 Starting interactive backups**



## Using the Data Protector CLI

To start ZDB to tape or ZDB to disk+tape, run:

```
omnib -mssql_list ListName
```

To start ZDB to disk, run:

```
omnib -mssql_list ListName -disk_only
```

where *ListName* is the name of the backup specification. For more information on omnib, see its man page.

## Restore

Data Protector offers restore from backup media to the application system on LAN (standard restore), where you can select various restore options depending on your restore scenario, and instant recovery. For more information, see the following sections.

### Before you begin

- Verify that the databases to be restored are not being in use.

---

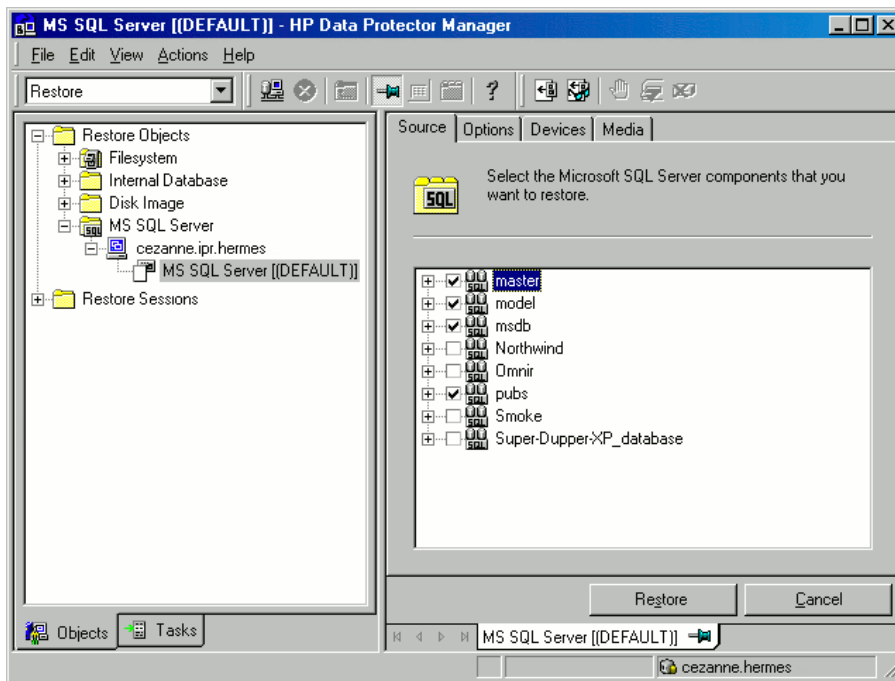
**NOTE:** There is no need to create an empty database before restore, because the database and its files are generated automatically by SQL Server.

---

Proceed as follows using the Data Protector Manager:

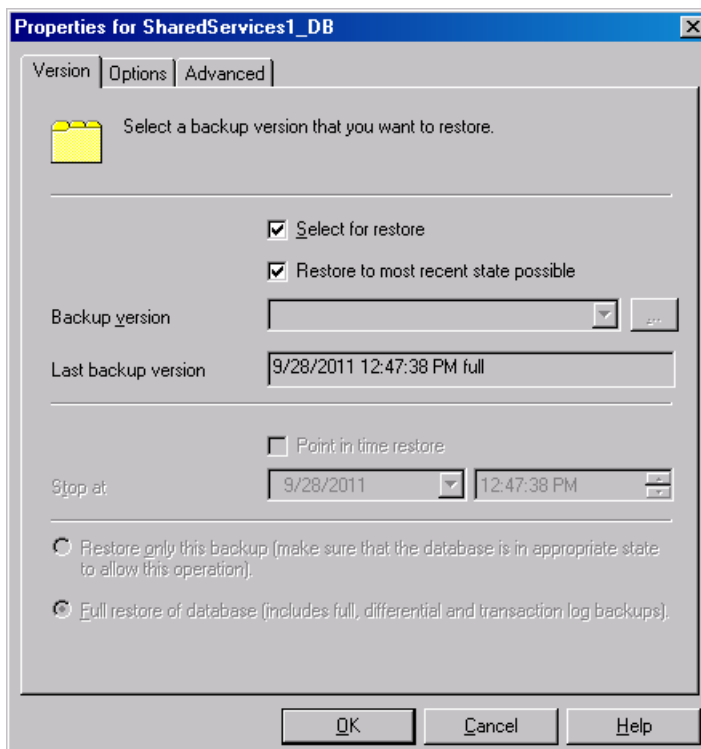
1. In the Context List, click **Restore**.
2. In the Scoping Pane, expand **Restore Objects, MS SQL Server**, and then select the client (backup system) from which you want to restore. A list of backed up objects is displayed in the Results Area.
3. Select the SQL Server objects that you want to restore. See [“Selecting backup objects for restore” \(page 167\)](#).

**Figure 72 Selecting backup objects for restore**



To select backup object-specific options, right-click the object and select **Properties**.

**Figure 73 Selecting object-specific options**



Select the version (backup date) which you want to use for restore or select the option **Restore to most recent state possible**. The latter always restores the chain of backups as if the **Full restore of database** option is selected. It includes the most recent full, differential, and transaction log backups. Select other restore options as appropriate. Note that some options are not available for restore of data files. See ["Restore options"](#) (page 168) for details. Click **OK**.

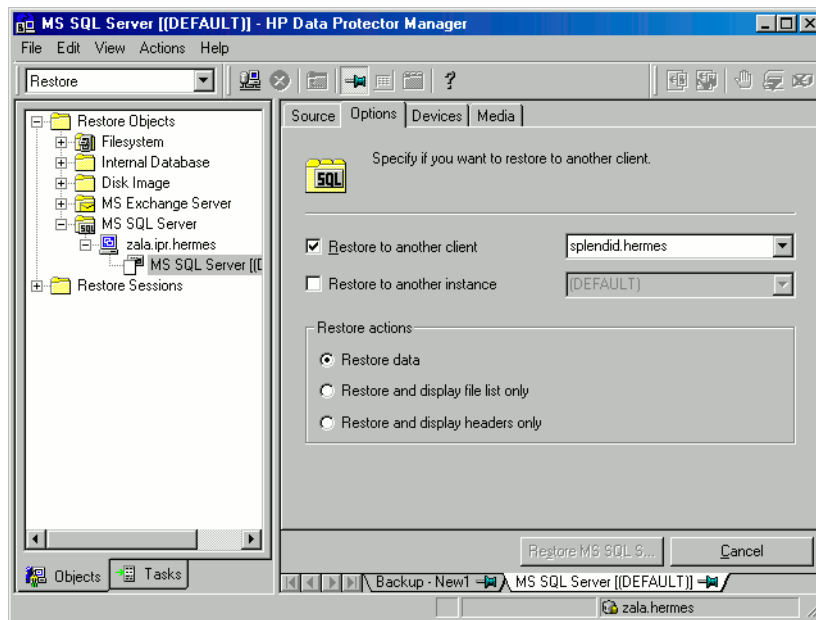
4. If you want to restore your data to another client or instance, specify new locations for the databases in the **Options** property page.

**NOTE:** When you click **Options**, the cell is browsed for running SQL Server instances that can become target instances for restore. If no instances are found, **Restore to another instance** is disabled and the message **There are no instances on this client system** is displayed.

Select one of the following **Restore actions**:

- **Restore data.** Select to restore the whole database. This option is selected by default.
- **Restore and display file list only.** Select if you do not know the original filenames. In this case, the files backed up in a particular session are displayed.
- **Restore and display headers only.** Select if you need specific details about backup. SQL Server header information is displayed.

**Figure 74 Restore options**



5. In the **Devices** page, select the devices to be used for the restore.  
For more information of how to select devices for a restore, see the online Help index: "restore, selecting devices for".
6. Click **Restore MS SQL Server** and then **Next** to select **Report level** and **Network load**.  
Click **Finish** to start restore.

## Restore options

**Table 17 Microsoft SQL Server database restore options**

| Option                       | Description                                                                                                                                                                                                                                                                                                                                                                                               |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Backup version</b>        | Specifies the backup session from which the selected objects will be restored.                                                                                                                                                                                                                                                                                                                            |
| <b>Point-in-time restore</b> | <p>This option is only available for database objects.</p> <p>Specifies a point in time to which the database state will be restored (you also need to select <b>Backup version</b> and set <b>Stop at</b>). After recovery, the database is in the state it was at the specified date and time.</p> <p>Only transaction logs written before the specified date and time are applied to the database.</p> |



**Table 17 Microsoft SQL Server database restore options** *(continued)*

| Option                                                      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Stop at</b>                                              | <p>This option is only available for database objects.</p> <p>Specifies the exact time when the rollforward of transactions will be stopped. Therefore, to enable database recovery to a particular point in time, backup you restore from must be a transaction log backup.</p> <p>You cannot use this option with <b>NORECOVERY</b> or <b>STANDBY</b>. If you specify <b>Stop at</b> time that is after the end of <b>RESTORE LOG</b> operation, the database is left in a non-recovered state (as if <b>RESTORE LOG</b> is run with <b>NORECOVERY</b>).</p>                                                                                 |
| <b>Restore only this backup</b>                             | <p>If you restored a database version and left it in a non-operational or standby state, you can subsequently restore differential or transaction log backups one by one, leaving each version non-operational to restore additional backups.</p>                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Full restore of the database</b>                         | <p>All necessary versions are restored, including the latest full backup, the latest differential backup (if one exists), and all transaction log backups from the last differential up to the selected version.</p>                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Force restore over the existing database</b>             | <p>Select this option if a database with the same name but a different internal structure already exists at the target Microsoft SQL Server instance.</p> <p>If this option is not selected, the Microsoft SQL Server does not let you overwrite the existing database - the restore will fail.</p> <p>If you are restoring a data file from the PRIMARY group to an existing database, you must specify the option at the data file level.</p> <p>When using this option, ensure that the most recent logs are backed up before the restore.</p>                                                                                              |
| <b>Put database in single user mode - log off all users</b> | <p>Disconnects all users that are connected to the target Microsoft SQL Server database and puts the database in the single user mode. Note that if the database is not in the simple recovery mode, the <b>Force restore over the existing database</b> option should also be selected.</p>                                                                                                                                                                                                                                                                                                                                                   |
| <b>Recovery completion state</b>                            | <p>Enables selecting the database state after recovery. You may select from:</p> <ul style="list-style-type: none"> <li>• Leaving the database operational. Once the last transaction log is restored and the recovery completed, the database becomes operational.</li> <li>• Leaving the database non-operational after the last transaction log is restored. You may restore additional transaction logs one by one.</li> <li>• Leaving the database in read-only mode. You may restore additional transaction logs before the database is set to read-write mode.</li> </ul> <p>This selection is only available for database objects.</p> |
| <b>Restore database with a new name</b>                     | <p>This option is only available for database objects.</p> <p>Restores the database under a different name. Specify the database logical filename and the destination filename (suboptions of <b>Restore files to new locations</b>).</p>                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Restore files to new locations</b>                       | <p>Restores files to a new location. Specify the database logical filename and a destination target filename for the specified logical filename. Use this option if you restore data to another server, instance, or make a database copy on the same server.</p>                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Restore to most recent state possible</b>                | <p>Restores the entire backup chain (includes the most recent full, differential, and transaction log backups).</p> <p>This option is selected by default.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |



**TIP:** To allow different restore scenarios, you can combine general restore options, such as **Restore database to another Microsoft SQL Server** and **Restore using a different device**, with object-specific restore options, such as **Point-in-time restore**, **Recovery completion state**, **Force restore over the existing database**.

## Restoring to another SQL Server instance or/and another SQL Server

### Prerequisites

- Both SQL Servers must have the same local settings (code page and sort order). This information is displayed in the session monitor for each backup.
- The target SQL Server must be configured and reside in the same Data Protector cell as the original SQL Server. For configuration procedure, see [“Creating ZDB specifications” \(page 160\)](#).
  1. Select the databases you want to restore and their versions.
  2. Select the following:
    - To restore to another SQL Server client, select **Restore to another client** and the target client from the drop-down list.
    - To restore to another SQL Server instance, select **Restore to another instance**. If there are no instances in the drop-down list, enter the instance name yourself.
    - Ensure that the specified SQL Server instance exists on the target client. Otherwise, restore fails.
  3. Specify new database locations.
  4. Start restore. See [“Restore” \(page 166\)](#).

### Instant recovery

See the *HP Data Protector Zero Downtime Backup Concepts Guide* and *HP Data Protector Zero Downtime Backup Administrator's Guide* for general information on instant recovery.

### Prerequisites

- If you restore user databases, put the databases offline:
  1. Start SQL Server Enterprise Manager.
  2. Selecting the database and click **Action**.
  3. Select all tasks and take them offline.
- If you restore system databases, put SQL Server offline by starting SQL Server Enterprise Manager, right-clicking SQL Server, and clicking **Stop**.


Perform instant recovery using the Data Protector Manager:

1. In the Context List, click **Instant Recovery**.
2. Expand **MS SQL Server** and select the backup session (replica) from which you want to restore. By default, the database will be recovered until the last backed up transaction.

3. To recover user databases to a specific point in time:
  - a. In the **Source** property page, under **Restore Objects**, right-click a database and click **Properties**.
  - b. In the **Backup version** drop-down list, select the required replica. The latest version is selected by default.

Select **Point in time restore**. From the **Stop at** drop-down list, select the point in time to which the transactions should be applied, and click **OK**. If no transaction logs are available, this option is disabled.

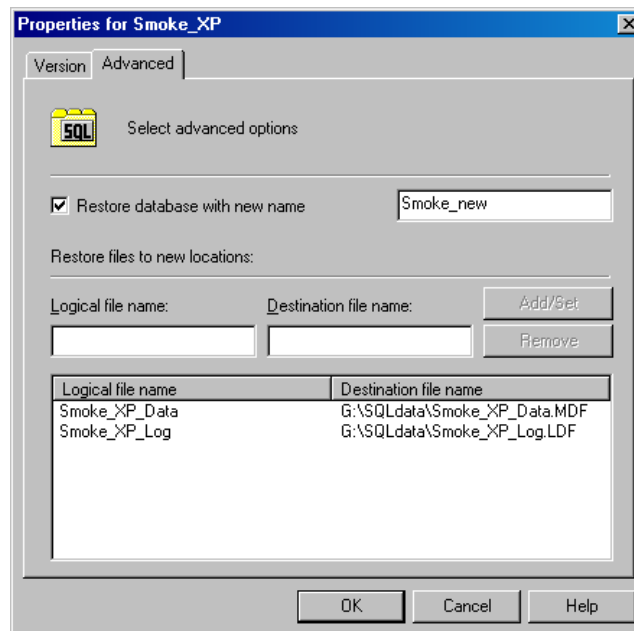
**Figure 75 Point-in-Time restore**



To restore the database under a different name, click **Advanced** and select **Restore database with new name**. See ["Restoring database with a new name" \(page 172\)](#).

❗ **IMPORTANT:** If the logical filename and physical filename are not listed, add them to the list. Specify the same names as used for ZDB; otherwise, instant recovery fails.

**Figure 76 Restoring database with a new name**



4. Click **Restore MS SQL Server**.

If you restore system databases, SQL Server displays errors because its services are offline. Therefore, when restore completes, start SQL Server manually using SQL Server Enterprise Manager.

## Monitoring sessions

You can monitor currently running or view previous sessions in the Data Protector GUI. When you run an interactive session, the monitor window shows you the session progress. Closing the GUI does not affect the session.

You can also monitor sessions using the **Monitor** context from any Data Protector client with the **User Interface** component installed.

For information on monitoring sessions, see the online Help index: “viewing currently running sessions” and “viewing finished sessions”.

## Troubleshooting

This section lists general checks and verifications, plus problems you might encounter when using the Data Protector SQL Server integration. Start at “[Problems](#)” (page 173). If you cannot find a solution there, perform general checks and verifications.

For general Data Protector troubleshooting information, see the *HP Data Protector Troubleshooting Guide*.

For general ZDB, restore, and instant recovery related troubleshooting, see the *HP Data Protector Zero Downtime Backup Administrator's Guide*.

## Before you begin

- Ensure that the latest official Data Protector patches are installed. For details of how to verify this, see the online Help index: “patches”.
- See the *HP Data Protector Product Announcements, Software Notes, and References* for general Data Protector limitations, as well as recognized issues and workarounds.
- See <http://www.hp.com/support/manuals> for an up-to-date list of supported versions, platforms, and other information.

## Checks and verifications

If your configuration, backup, or restore failed:

- Check that SQL Server services are running.
- Examine system errors reported in *Data\_Protector\_home\log\debug.log* on the SQL Server client.  
Additionally, check *errorlog* and *VDI.log* files in the *MSSQL\log* directory.
- Make a test filesystem backup and restore of the problematic client. For information, see the online Help.
- Check that every SQL Server used with Data Protector has the MS SQL Integration component installed.
- Connect to SQL Server via SQL Server Enterprise Manager using the same login ID as you specified in the Data Protector **Configuration** dialog box.
- Perform a database backup using SQL Server Enterprise Manager. If the backup fails, fix any SQL Server problems, and then perform a backup using Data Protector.

Additionally, if your backup failed:

- Verify the configuration file to check if the Cell Manager is correctly set on SQL Server.
- If you do not see the SQL Server instance as the application database when creating a backup specification, enter the instance name yourself. When "not-named instance" is not displayed, insert the *DEFAULT* string.
- If Data Protector reports that the integration is properly configured, verify that the SQL Server user has appropriate rights to access the required databases.

During master database restore, the following error occurs when executing an SQL statement:

```
Error has occurred while executing an SQL statement.
Error message: 'SQLSTATE: [42000] CODE: (3108) MESSAGE: [Microsoft]
[ODBC SQL Server Driver] [SQL Server]To restore the master database,
the server must be running in single user mode. For information on
starting in single user mode, see "How to: Start an Instance of SQL
Server (sqlservr.exe)" in Books Online.
```

Note that this behavior is expected when the master database is not restored in single user mode.

## Problems

### Problem

#### The integration is properly configured but the database backup fails after a timeout

- With an error similar to:  

```
[Warning] From: OB2BAR@computer.company.com "SQLSRV"
Time: 7/29/2011 8:19:22 PM
Error has occurred while executing SQL statement.
[Microsoft][ODBC SQL Server Driver][SQL Server]Backup or restore
operation terminating abnormally.'
[Critical] From: OB2BAR@computer.company.com "SQLSRV"
Time: 7/29/11 8:19:24 PM
Received ABORT request from SM => aborting
```
- SQL Server error log contains an entry similar to:  

```
2011-07-29 20:19:21.62 kernel
BackupVirtualDeviceSet::Initialize: Open failure on backup
```

```
device 'Data_Protector_master'.
Operating system error -2147024891(Access is denied.).
```

- SQL Server VDI.LOG file contains an entry similar to:  
2011/07/30 13:19:31 pid(2112)  
Error at BuildSecurityAttributes: SetSecurityDescriptorDacl  
Status Code: 1338, x53A Explanation: The security descriptor  
structure is invalid.

SQL Server service and Data Protector Inet are running under different accounts. The integration cannot access SQL Server due to security problems.

#### Action

Restart the Data Protector Inet service under the same account as the SQL Server service is running.

#### Problem

##### **Backup fails if the appropriate drive letter on the backup system does not exist**

Backup fails with an error, similar to:

```
[Major] From: SSEA@computer1.com " " Time: 02-Feb-11 14:07:54
Filesystem \\.\Volume{ef58fe0e-b2b8-11db-aa08-000802804af6} could not be mounted
to Q:\Program Files\Microsoft SQL Server\MSSQL.1\MSSQL\DATA\
([87] The parameter is incorrect.).
```

Backup fails because SSE agent tries to mount the filesystem to the drive letter which does not exist on the backup system. The drive letter must be the same as on the application system. SSE or SMI-S agents always mount the filesystems to the same drive letters on the backup systems as if the ZDB\_PRESERVE\_MOUNTPOINTS omnirc variable is set to 1.

#### Action

For source volumes to be successfully replicated, create the same drive letters on the backup system as they are used on the application system for mounting the source volumes.

#### Problem

##### **Database is left in unrecovered state after “Invalid value specified for STOPAT parameter” is reported**

The database remains in an unrecovered state as if the RESTORE LOG operation was run with **Leave the database non-operational**.

#### Action

Recover the database to the latest point in time using SQL Query Analyzer:

```
RESTORE DATABASE database_name WITH RECOVERY
```

After the recovery, additional transaction logs cannot be applied.

#### Problem

##### **Transaction logs cannot be restored from tape**

The recovery completed successfully and the database was put in `norecovery` state, but transaction logs cannot be restored from tape.

#### Action

Recover the database to the state of ZDB to disk using the SQL Query Analyzer:

```
RESTORE DATABASE database name WITH RECOVERY
```

After the recovery, additional transaction logs cannot be applied.

## Problem

### Instant recovery of SQL Server databases fails

If the SQL Server service is offline prior to the instant recovery, instant recovery of the SQL Server databases fails.

The following errors are displayed:

```
[Critical] From: computer@company.com "(DEFAULT)" Time: 4/9/2011 7:01:42 PM
Microsoft SQL Server reported the following error during login:
The object was not open.
[Warning] From: computer@company.com "(DEFAULT)" Time: 4/9/2011 7:01:42 PM[152:9208] Data Protector is probably
not configured for use with SQL Server on this host.
```

## Action

Perform one of the following:

- Set the Data Protector `omnirc` variable `OB2_SQLRESTORE_STARTSRV`, which starts the SQL Server service prior to the recovery of SQL databases, to 1.

During the master database restore, the following error is displayed:

```
RESTORE master with SNAPSHOT is not supported.
```

Note that this behavior is expected. No further steps are needed after the instant recovery.

- Restart SQL Server instance services after instant recovery completes. If restarting services does not automatically start the recovery of all system databases, start the SQL Server instance in single user mode and manually start the recovery of the master database. Follow the same procedure for other system databases. At the end, restart SQL Server instance services.

## Problem

### Restore to another client in the Data Protector cell not configured for use with SQL Server fails

## Action

Configure the SQL integration on this client (see ["Configuring the integration"](#) (page 154)).

## Problem

### Database is left in unrecovered state after restore completed successfully

If you set the time for `Stop` at beyond the end of the `RESTORE LOG` operation, the database remains in the unrecovered state as if the `RESTORE LOG` operation was run with `Leave the database non-operational`.

## Action

Recover the database to the latest point in time by using the SQL Query Analyzer:

```
RESTORE DATABASE database_name WITH RECOVERY
```

After the recovery, additional transaction logs cannot be applied.

## Problem

### Instant recovery of a Microsoft SQL Server database configured on a Microsoft Cluster Server system fails with "The physical filename may be incorrect"

Instant recovery of the Microsoft SQL Server data in an HP Business Copy P9000 XP configuration on a Microsoft Cluster Server system fails with the following error:

```
[Microsoft] [ODBC SQL Server Driver] [SQL Server]Device activation error.
The physical file name '<Data/Log filename>' may be incorrect.
```

## Action

Perform the following steps:

1. Using Microsoft SQL Server Enterprise Manager, detach the Microsoft SQL Server database that you want to recover.
2. Using Cluster Administrator, take the Microsoft SQL Server Disk resource offline.
3. On the application system, configure the `ZDB_TAKE_CLUSRES_ONLINE` omnirc variable. For details, see [“ZDB integrations omnirc variables” \(page 271\)](#).
4. Start instant recovery.
5. When the message `Please, take MS SQL cluster resources online` appears in the Data Protector GUI, bring the Microsoft SQL Server cluster resources online using Cluster Administrator.

## Problem

### Restoring a Microsoft SQL Server 2005 instance to an alternate location when full-text indexing is enabled fails

When the Use full-text indexing option is enabled for a particular database in a Microsoft SQL Server 2005 instance, the restore session does not complete successfully, since restore of the full-text catalog of the SQL database fails. The session report contains warning messages about the full-text catalog file being used by the affected database.

## Action

To solve the problem:

1. In the HP Data Protector Manager, switch to the **Restore** context.
2. In the Scoping Pane, expand **Restore Objects** and then **MS SQL Server**. Select name of the Microsoft SQL Server for which you want to perform restore.
3. In the Results Area, double-click the bar name corresponding to the particular Microsoft SQL Server instance. A list of backed up objects gets displayed.
4. Select the desired Microsoft SQL Server database, right-click it, and click **Properties**.
5. In the Properties window, click the **Advanced** tab.
6. Select the **Restore database with new name** option, and enter the new database name in the text box.
7. For all logical file names that are already present on the list, update contents of the Destination file name column accordingly.
8. Add the full-text catalog to the list.

In the Logical file name text box, enter the string `sysft_Full-Text_Catalog_Name`. In the Destination file name text box, enter the corresponding physical location.

---

**NOTE:** The full-text catalog is always restored to its original location, regardless of the specified physical location.

---

9. Click **Add/Set**.
10. In the Version and Options property pages, specify the appropriate options. For details, see [“” \(page 166\)](#).
11. Click **OK** to close the Properties window.
12. In the Options, Devices, and Media property pages, specify the appropriate options. For details, see [“” \(page 166\)](#).
13. Click **Restore** and then **Next** to select the Report level and Network load.
14. Click **Finish** to start the restore session.

## Problem

### Database restore fails

The restore session aborts with a major error similar to:



Error has occurred while executing a SQL statement. Error message:  
'SQLSTATE: [42000] CODE: (3159) MESSAGE: [Microsoft] [ODBC SQL Server Driver] [SQL Server]The tail of the log for the database "test2" has not been backed up. Use BACKUP LOG WITH NORECOVERY to backup the log if it contains work you do not want to lose. Use the WITH REPLACE or WITH STOPAT clause of the RESTORE statement to just overwrite the contents of the log. SQLSTATE: [42000] CODE: (3013) MESSAGE: [Microsoft] [ODBC SQL Server Driver] [SQL Server]RESTORE DATABASE is terminating abnormally.'

#### Action

To solve the problem, perform one of the following before restarting the restore session:

- Select the restore option **Enable tail log backup** (recommended).
- Perform a transaction log backup to obtain the most recent transaction logs.

#### Problem

##### **Restore of a database to another client or instance with tail log backup enabled fails if the target database does not exist**

If Data Protector backup session for a non-existing Microsoft SQL Server database is started, the session fails. When restoring a database to another client or instance the tail log backup is performed on the target database.

As tail log backup is a transaction log backup, in the described circumstances the tail log backup session fails. Consequently, the restore sessions fails too.

#### Action

Disable the restore option **Enable tail log backup** and restart the restore session.

#### Problem

##### **Restore of a database in a log shipping configuration with the tail log backup enabled fails**

In a Microsoft SQL Server log shipping configuration, Data Protector performs differential database backup instead of transaction log backup when the latter is run. The automatic backup type switch takes place also with tail log backup. In these circumstances, the database backup chain does not contain most recent transactions from the tail of the log. If tail of the log of the target database has not been backed up yet, Microsoft SQL Server does not allow restoring over this database.

#### Action

Perform one of the following and restart the restore session:

- Disable Microsoft SQL Server log shipping.
- Enable the option **Force restore over existing database** for all involved databases.



---

**CAUTION:** Tails of the logs of all involved databases will be lost.

---

# 4 Data Protector Microsoft Exchange Server 2003 ZDB integration

## Introduction

This chapter explains how to configure and use the Data Protector Microsoft Exchange Server ZDB integration. It describes the concepts and methods you need to understand to back up and restore the Microsoft Exchange Server (**Exchange Server**) database objects.

During backup, the database is only stopped for the time it takes to create a **replica** (split the mirror disks or create snapshots). If ZDB to tape is performed, backup is subsequently performed offline on the backup system.

The following disk arrays and array configurations are supported:

| Supported array                                  | Supported configurations                                   |
|--------------------------------------------------|------------------------------------------------------------|
| HP P9000 XP Disk Array Family (P9000 XP Array)   | HP BC P9000 XP, HP CA P9000 XP, combined HP CA+BC P9000 XP |
| HP P6000 EVA Disk Array Family (P6000 EVA Array) | HP BC P6000 EVA, combined HP CA+BC P6000 EVA               |

All ZDB types (ZDB to tape, ZDB to disk, and ZDB to disk+tape) are supported by this integration. For ZDB types description, see the *HP Data Protector Zero Downtime Backup Concepts Guide* and the online Help.

Using Data Protector, you can restore Exchange Server data:

- From backup media to the application system on LAN (standard restore).
- Using instant recovery.

The following table gives an overview of Exchange Server recovery methods:

|                         | Rollforward recovery                                                                                                                                                                            | Point-in-time recovery                                                                                              |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------|
| <b>ZDB to tape</b>      | Standard restore (database) + standard restore (transaction logs)                                                                                                                               | Standard restore (database)                                                                                         |
| <b>ZDB to disk</b>      | Instant recovery (database) + standard restore (transaction logs)                                                                                                                               | Instant recovery (database)                                                                                         |
| <b>ZDB to disk+tape</b> | <ul style="list-style-type: none"><li>• Instant recovery (database) + standard restore (transaction logs)</li><li>• Standard restore (database) + standard restore (transaction logs)</li></ul> | <ul style="list-style-type: none"><li>• Instant recovery (database)</li><li>• Standard restore (database)</li></ul> |

For a description of ZDB and instant recovery concepts, see the *HP Data Protector Zero Downtime Backup Concepts Guide*.

## Integration concepts

The Exchange ZDB integration backs up Microsoft Information Store (MIS), Key Management Service (KMS), and Site Replication Service (SRS), which are replicated within a disk array.

Data Protector integrates with Exchange Server through:

- The Data Protector `ese_bar.exe` executable, installed on Exchange Server and used for retrieving locations of files and folders pertaining to a storage group/store.
- The `omniEx2000.exe` command, responsible for mounting and dismounting databases and purging transaction logs after a transaction logs backup.

An Exchange ZDB specification is created using the `omnicreated1` command. Based on this specification, you can perform a filesystem ZDB using the Data Protector GUI. During backup, `ese_bar.exe` resolves database objects and retrieves a list of files and folders to be backed up. The following happens:

- The command specified in the option `Stop/quiesce` the application command line is executed, which dismounts the databases and verifies their consistency using `eseutil.exe`. See “Checking Exchange files for consistency” (page 187).
- When the replica is created, the command specified in the option `Restart` the application command line is executed, which mounts dismounted databases.
- If ZDB to tape is performed, Data Protector backs up Exchange databases on the backup system. Transaction logs are not deleted and can be backed up in a separate session using non-ZDB procedure.

ZDB process depends on whether you replicate your data in a split mirror or snapshot replication and on the selected ZDB type. Restore process depends on the restore type - standard restore or instant recovery. See the *HP Data Protector Zero Downtime Backup Concepts Guide* for a detailed description of replication techniques, as well as ZDB and restore processes.

## Configuring the integration

### Prerequisites

- You need to have the Data Protector online extension license-to-use (LTU) for Windows to be able to use the Data Protector Microsoft Exchange Server ZDB integration.  
For more information, see the *HP Data Protector Installation and Licensing Guide*.
- You need to have appropriate Data Protector zero downtime backup extension and instant recovery extension licenses-to-use (LTU).
- Ensure that you correctly installed and configured Exchange Server.
  - For supported versions, platforms, devices, and other information, see the latest support matrices at <http://www.hp.com/support/manuals>.
  - For information on installing, configuring, and using Exchange Server, see the Exchange Server documentation.
- Ensure that you correctly installed Data Protector. For information on installing Data Protector in various architectures and installing a Data Protector disk array integration (P9000 XP Array or P6000 EVA Array) with Exchange Server, see the *HP Data Protector Installation and Licensing Guide*.  
Every Exchange Server system to be used with Data Protector must have the MS Exchange Integration component installed.
- Install Exchange Server on the application system. All parts of the Exchange database (Information Store (MIS), Key Management Service (KMS), and Site Replication Service (SRS)) must be installed on the disk array source volumes.
- Zero downtime backup of Exchange Server 2003 databases residing on a disk array of the HP P9000 XP Disk Array Family requires Exchange Server 2003 SP2 or higher.

### Before you begin

- Configure devices and media for use with Data Protector. For instructions, see the online Help index: “configuring devices” and “creating media pools”.
- To test whether Exchange Server and Cell Manager communicate properly, configure and run a Data Protector filesystem ZDB and restore. For instructions, see the online Help.

- Ensure that the database files and the log files are located on the same drive letter or mount point.
  - Before performing transaction logs backups, disable **circular logging** for all storage groups. If the application is cluster-aware, disable circular logging on all cluster nodes.
  - Add the `Exchange_home\bin` directory to the Windows **Path** environment variable:
    1. In the Windows Explorer, right-click **My Computer** and click **Properties**.
    2. In the **Properties** dialog box, click **Advanced** and then **Environment Variables**.
    3. Select **Path** in the **System Variables** list and click **Edit**.
    4. Add `Exchange_home\bin` in the **Variable Value** text box and click **OK**.

If the integration is cluster-aware, perform this procedure on all cluster nodes.
  - In cluster environments, prior to configuring a ZDB specification, create a new physical resource (virtual disk for P6000 EVA Array or LDEV for P9000 XP Array) for the cluster and move the Exchange Server database and transaction log files to it.
- For information on moving the Exchange Server database and transaction log files, see the following documents available at <http://support.microsoft.com>:
- *How to Move Exchange Databases and Logs in Exchange Server 2003 (821915)*
  - *HOW TO: Add New Mailbox Stores in Exchange Server 2003 (821748)*

## Backup

To run ZDB of an existing Exchange Server ZDB specification:

- Schedule a backup using the Data Protector Scheduler.
- Start an interactive backup using the Data Protector GUI or CLI.

### Limitations

- Backup preview is not supported.

### Considerations

- You can perform transaction logs backups only if circular logging is disabled for the involved Exchange Server.  
 Circular logging is a Microsoft Exchange mode, where transaction logs are automatically overwritten when the data they contain is committed to the database.  
 If enabled, this option reduces disk storage space requirements, but does not allow you to perform transaction logs backups.
- Exchange Server on the application system must be running. Otherwise, data consistency is not guaranteed.
- You cannot start ZDB, restore, or instant recovery using the same source volume on the application system at the same time. A session must be started only after the preceding session using the same source volume on the application system finishes; otherwise, the session fails.
- Transaction logs must be backed up using the common filesystem backup functionality.
- When backing up log files, select either **Log all** or **Log files** logging level in the **Advanced Filesystem Options** dialog box. This enables purging backed up log files from disk. For more information, see the online Help index: "filesystem options".

## Configuring Exchange Server ZDB

To configure a ZDB:

1. Configure devices and media for backup.
2. Create a Data Protector Exchange Server ZDB specification.

## Creating ZDB specifications

Create a ZDB specification using the `omnicreatedl` command. After that, perform a filesystem ZDB using the Data Protector GUI.

The `omnicreatedl` command creates a ZDB specification with included scripts `Stop/quietse` the application command line and `Restart the application command line` (`omniEx2000.exe`) for mounting/dismounting backed up databases and checking their consistency.

Additionally, `omnicreatedl` creates a transaction logs backup specification for each storage group specified in a ZDB specification (circular logging disabled). See [Figure 91 \(page 198\)](#). The transaction logs backup specification includes the post-exec script for purging backed up log files.



**TIP:** A transaction logs backup specification can either be triggered by post-exec (configured on the backup specification level) in the backup specification for database files backup (recommended), or started manually after the backup specification for database files backup is started. For more information, see the online Help index: "pre- and post-exec commands".

The following is the synopsis of the `omnicreatedl` command:

```
omnicreatedl -ex2000 -datalistName [-deviceName] {
P9000_DISK_ARRAY_XP_OPTIONS | P6000_ENTERPRISE_VIRTUAL_ARRAY_OPTIONS }
EXCHANGE_OPTIONS [-force] [-virtualSrv Name]
```

For a detailed synopsis of `P9000_DISK_ARRAY_XP_OPTIONS`, `P6000_ENTERPRISE_VIRTUAL_ARRAY_OPTIONS`, and `EXCHANGE_OPTIONS`, see the `omnicreatedl` man page.

Descriptions of the parameters are presented in the tables below:



**IMPORTANT:** If parameters contain spaces, use double quotes when specifying them in the `omnicreatedl` command. For example, `-storage_group "First Storage Group"`.

**Table 18 General options**

| Parameter                     | Description                                                                                                                                                                                                                                                                                                                                          |
|-------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>-ex2000</code>          | Instructs the <code>omnicreatedl</code> command to create a Microsoft Exchange Server ZDB backup specification and Microsoft Exchange Server transaction logs backup specification for all specified storage groups.                                                                                                                                 |
| <code>-datalist Name</code>   | Specifies the name of the Exchange Server ZDB backup specification file (datalist) for the Microsoft Exchange Server ZDB. The datalist is created on the Cell Manager in the <code>Data_Protector_home\config\server\datalists</code> directory (on Windows systems) or the <code>/etc/opt/omni/server/datalists</code> directory (on UNIX systems). |
| <code>-device Name</code>     | Specifies the backup device to be used for the backup. If this option is not specified, the backup device must be specified using the Data Protector GUI.                                                                                                                                                                                            |
| <code>-force</code>           | Forces overwriting of an existing ZDB backup specification file with the same name.                                                                                                                                                                                                                                                                  |
| <code>-virtualSrv Name</code> | The name of the virtual server system on which Microsoft Exchange Server is running. This option is mandatory in cluster configurations.                                                                                                                                                                                                             |

**Table 19 P9000 XP Array options**

| Parameter                                                                                     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-----------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>-split_mirror -sse</code>                                                               | Instructs the <code>omnicreatedl</code> command to create an HP P9000 XP Disk Array Family Microsoft Exchange Server ZDB backup specification.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <code>( -local App_sys Bck_sys   -remote App_sys Bck_sys   -combined App_sys Bck_sys )</code> | <p>Specifies one of the three HP P9000 XP Disk Array Family configurations:</p> <ul style="list-style-type: none"> <li>• <code>-local</code> selects the HP Business Copy (BC) P9000 XP configuration.</li> <li>• <code>-remote</code> selects the HP Continuous Access (CA) P9000 XP configuration.</li> <li>• <code>-combined</code> selects the combined (HP CA+BC P9000 XP configuration.</li> </ul> <p>With any configuration, <code>App_sys</code> specifies the application system and <code>Bck_sys</code> specifies the backup system.</p> <p>In a cluster environment, specify the virtual server host name instead of the physical node host name.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <code>-mirrors MU_numbers</code>                                                              | <p>This option is optional. It is only considered when the HP Business Copy (BC) P9000 XP configuration is chosen.</p> <p>Specify the mirror unit (MU) number(s) of a replica or a replica set from which the Data Protector HP StorageWorks P9000 XP Agent, according to the replica set rotation, selects the replica to be used in the zero downtime backup session. The replica selection rule is described in the <i>HP Data Protector Zero Downtime Backup Concepts Guide</i>. The maximum number of replicas that can be created for the same source volumes is different for mirror copies and snapshots. Both limitations are imposed by the HP P9000 XP Disk Array Family storage system.</p> <p>You can specify one or more non-negative integer numbers, one or more ascending ranges of such numbers, or any combination of both. Use a comma as the separator character. Examples:</p> <p>5</p> <p>7-9</p> <p>4,0,2-3</p> <p>When a sequence is specified, it does not define the order in which the replicas are used. If this option is not specified, the MU number 0 is used.</p> |
| <code>-instant_restore</code>                                                                 | <p>This option is optional.</p> <p>When this option is specified, the <code>omnicreatedl</code> command automatically sets the <code>-keep_version</code> option.</p> <p>Specify the <code>-instant_restore</code> option to enable ZDB to disk or ZDB to disk+tape and instant recovery from the replica. If this option is not specified, it is not possible to perform a ZDB to disk or a ZDB to disk+tape and instant recovery from the replica. However, this option does not influence the replica set rotation.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <code>-keep_version</code>                                                                    | <p>This option is optional.</p> <p>If configuring a ZDB to tape, specify select this option to keep the replica on the disk array after the zero downtime backup session. The replica becomes part of a replica set (specify a value for the option <code>-mirrors</code>). Unless the additional option <code>-instant_restore</code> is specified, the replica is not available for instant recovery.</p> <p>If this option is not specified, the replica is removed at the end of the session. In this case, it is also not possible to specify the <code>-leave_enabled_bs</code> option.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <code>-leave_enabled_bs</code>                                                                | <p>This option is optional.</p> <p>To specify this option, the <code>-keep_version</code> option has to be specified.</p> <p>By default, Data Protector dismounts the filesystems on the backup system after each ZDB session.</p> <p>If this option is specified, the filesystems remain mounted after the backup. Thus, you can use the backup system for some data warehouse activity afterwards, but not for instant recovery.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <code>[ -split   -establish ]</code>                                                          | Specifying one of these two options is optional. If none is specified, the <code>-establish</code> option is set by default.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

**Table 19 P9000 XP Array options** (*continued*)

| Parameter | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|           | <p>If the <code>-split</code> option is specified, the volumes of the replica selected for the current ZDB session are prepared for the zero downtime backup at the start of the current ZDB session: mirrors are resynchronized with the P-VOLs, and volumes to be used for snapshot storage are made empty.</p> <p>If the <code>-establish</code> option is specified, if the volumes of the replica to be used in the next ZDB session are not ready for ZDB, they are prepared for ZDB at the end of the current ZDB session.</p> |

**Table 20 P6000 EVA Array options**

| Parameter                                              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>-snapshot</code>                                 | Instructs the <code>omnicreatedl</code> command to create an HP P6000 EVA Disk Array Family Microsoft Exchange Server ZDB backup specification.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <code>-smis app_sys bck_sys</code>                     | Specifies the application system <code>app_sys</code> and the backup system <code>bck_sys</code> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <code>-instant_recovery</code>                         | <p>This parameter is optional.</p> <p>Select to perform a ZDB to disk or ZDB to disk+tape and leave the replica on a disk array for instant recovery. Also, specify <code>-snapshots number</code>. The options <code>-snapshot_type clone</code> and <code>-snapshot_policy strict</code> are automatically selected.</p> <p>If this option is not set, you cannot perform instant recovery from the replica created or reused in this session.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <code>-snapshots number</code>                         | <p>This parameter is optional.</p> <p>During ZDB, Data Protector creates a new replica and leaves it on the array until the specified <code>-snapshots number</code> is reached. After that, the oldest replica is reused.</p> <p>Default: 1.</p> <p>The maximum number for vsnaps and standard snapshots is 7. Data Protector does not limit the number of replicas rotated, but the session fails if the limit is exceeded.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <code>-snapshot_type {standard   vsnap   clone}</code> | <p>Instructs Data Protector to create a particular type of P6000 EVA Array snapshots:</p> <ul style="list-style-type: none"> <li><code>-standard</code> creates snapshots with the pre-allocation of disk space.</li> <li><code>-vsnap</code> creates snapshots without the pre-allocation of disk space.</li> <li><code>-clone</code> creates a clone of a source volume (original virtual disk).</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <code>-snapshot_policy {strict   loose}</code>         | <p>Specifies a snapshot policy depending on the type of already existing snapshots for the same source volume.</p> <p>When <code>-strict</code> is selected, Data Protector attempts to create snapshots as specified by <code>-snapshot_type</code>. If the source volumes used in the session have existing snapshots of a different type, the selected type is not created (the session is aborted).</p> <p>When <code>-loose</code> is selected, Data Protector creates snapshots of a different type than specified by <code>-snapshot_type</code> (if this helps to complete a session successfully).</p> <p>For example, if you select standard snapshots, but Data Protector detects that standard snapshots cannot be created because vsnaps/snapclones of the source volumes already exist in a replica set, it creates either vsnaps or snapclones instead of standard snapshots.</p> <p>Note that Data Protector can use only one type of snapshots in a backup session. For example, if the source volumes used in a session have existing standard snapshots/vsnaps, the session is aborted.</p> |



**Table 20 P6000 EVA Array options** *(continued)*

| Parameter                                                                               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>-wait_clonecopy number</code>                                                     | Available if <code>-snapshot_type clone</code> is selected.<br>Specify to prevent degradation of the application data access times by delaying moving data to tape until the cloning process completes (ZDB to tape, ZDB to disk+tape). Set the maximum waiting time. When the specified time is reached, backup to tape starts (even if cloning is not finished).                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <code>-replica_conf {local   combined}</code>                                           | Specifies P6000 EVA Array configurations:<br><code>local</code> selects HP Business Copy (BC) P6000 EVA.<br><code>combined</code> selects HP Continuous Access + Business Copy (CA+BC) P6000 EVA.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <code>-ca_failover_option {follow_replica_direction   maintain_replica_location}</code> | Available if <code>combined</code> configuration is selected.<br>Specify to control the replication direction after a failover.<br>Select <code>follow_replica_direction</code> to follow the replication direction and create replicas on the array remote to current source. A failover reverses the replication direction and the replicas are created on the array that was originally a source P6000 EVA Array.<br>Select <code>maintain_replica_location</code> to maintain the replica location and create replicas on the array remote to home. After a failover, replicas continue to be created on the destination array that has also become a source P6000 EVA Array.<br>When <code>-ca_failover_option</code> is selected, <code>follow_replica_direction</code> is set as default. |

**Table 21 Exchange Server options**

| Parameter                                      | Description                                                                                                                                                                                                                                                                 |
|------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>-annotation { MIS   SRS   KMS }</code>   | Specifies Exchange Server annotations: Microsoft Information Store (MIS), Site Replication Service (SRS), and Key Management Service (KMS).<br>Default: MIS.                                                                                                                |
| <code>-all_storage_groups</code>               | Creates a backup specification for all databases relating to MIS (specified by <code>-annotation MIS</code> ).                                                                                                                                                              |
| <code>-storage_group Storage_Group_Name</code> | Creates a backup specification for all stores relating to the specified. Multiple declarations of <code>-storage_group</code> create a backup specification for the selected storage groups.                                                                                |
| <code>-store Store</code>                      | Creates a backup specification only for specified store(s) inside the storage group. To create a backup specification for many stores, specify a list of stores after the <code>-store</code> parameter.<br>You can obtain store names using Exchange System Administrator. |

For more information, see the `omnicreated1` man page.

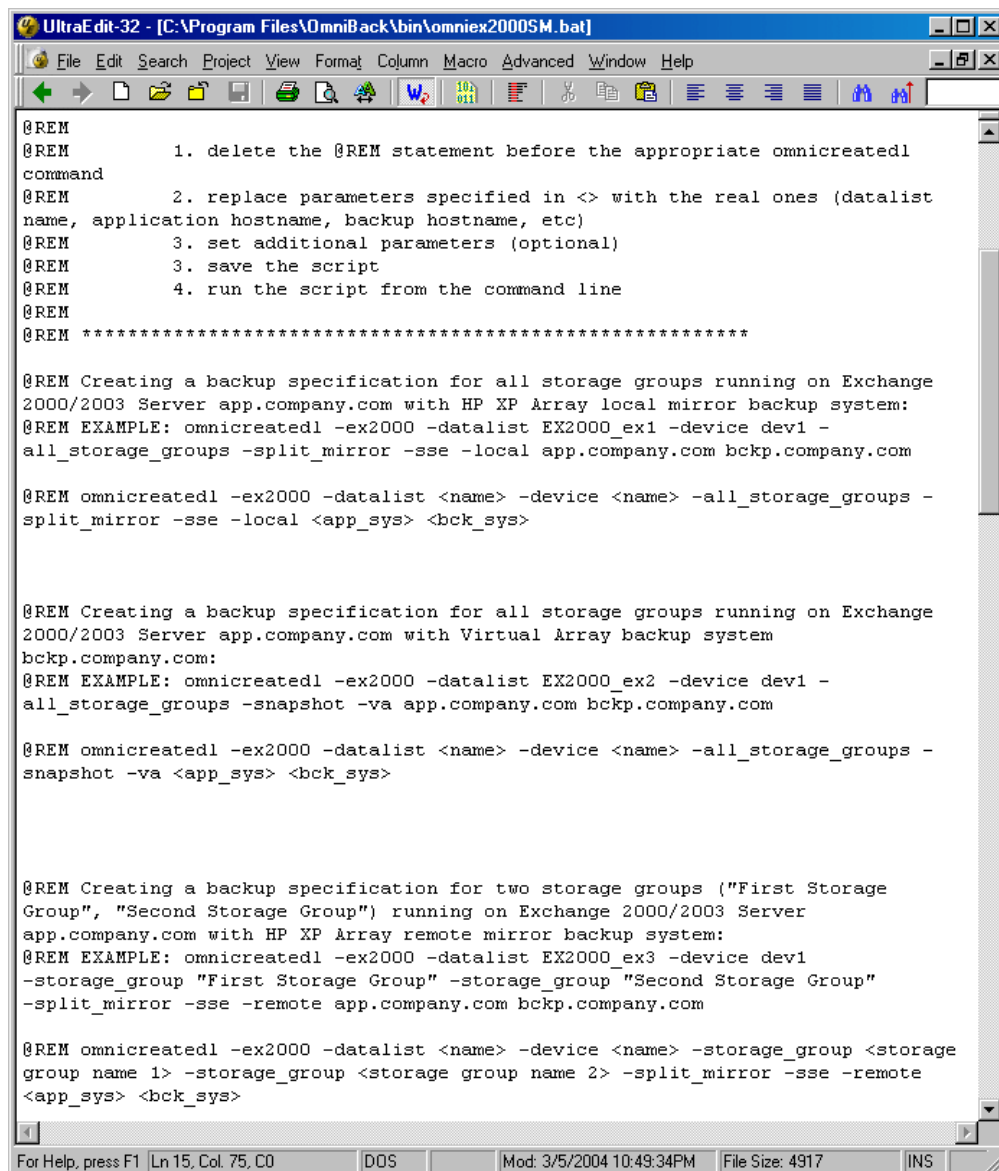


**TIP:** To create an Exchange ZDB specification and a transaction logs backup specification, you can also use the `omniex2000SM.bat` file, located in the `Data_Protector_home\bin` directory on the application system. The script provides templates and examples of `omnicreated1` usage. You can modify `omniex2000SM.bat` using any text editor. Uncomment (delete `@REM` before the command) the line with the appropriate command and edit parameters. See [Figure 77 \(page 185\)](#).

To create an Exchange ZDB specification and a transaction logs backup specification using `omniex2000SM.bat`, make the necessary modifications and run:  
`Data_Protector_home\bin\omniex2000SM.bat`.



Figure 77 omniex2000SM.bat file



```
@REM
@REM 1. delete the @REM statement before the appropriate omnicreatedl
command
@REM 2. replace parameters specified in <> with the real ones (datalist
name, application hostname, backup hostname, etc)
@REM 3. set additional parameters (optional)
@REM 3. save the script
@REM 4. run the script from the command line
@REM
@REM *****

@REM Creating a backup specification for all storage groups running on Exchange
2000/2003 Server app.company.com with HP XP Array local mirror backup system:
@REM EXAMPLE: omnicreatedl -ex2000 -datalist EX2000_ex1 -device dev1 -
all_storage_groups -split_mirror -sse -local app.company.com bckp.company.com

@REM omnicreatedl -ex2000 -datalist <name> -device <name> -all_storage_groups -
split_mirror -sse -local <app_sys> <bck_sys>

@REM Creating a backup specification for all storage groups running on Exchange
2000/2003 Server app.company.com with Virtual Array backup system
bckp.company.com:
@REM EXAMPLE: omnicreatedl -ex2000 -datalist EX2000_ex2 -device dev1 -
all_storage_groups -snapshot -va app.company.com bckp.company.com

@REM omnicreatedl -ex2000 -datalist <name> -device <name> -all_storage_groups -
snapshot -va <app_sys> <bck_sys>

@REM Creating a backup specification for two storage groups ("First Storage
Group", "Second Storage Group") running on Exchange 2000/2003 Server
app.company.com with HP XP Array remote mirror backup system:
@REM EXAMPLE: omnicreatedl -ex2000 -datalist EX2000_ex3 -device dev1
-storage_group "First Storage Group" -storage_group "Second Storage Group"
-split_mirror -sse -remote app.company.com bckp.company.com

@REM omnicreatedl -ex2000 -datalist <name> -device <name> -storage_group <storage
group name 1> -storage_group <storage group name 2> -split_mirror -sse -remote
<app_sys> <bck_sys>
```

## Examples of using omnicreatedl

### Example 1 – P9000 XP Array:

To create a ZDB specification BS1, using device dev1, for all storage groups running on Exchange Server on the application system computer\_app.company.com and backup system computer\_bck.company.com using the HP BC P9000 XP configuration, run:

```
omnicreatedl -ex2000 -datalist BS1 -device dev1 -all_storage_groups
-split_mirror -sse -local computer_app.company.com
computer_bck.company.com
```

### Example 2 – P9000 XP Array:

To create a ZDB specification BS2, using device dev1, for two storage groups (First Storage Group and Second Storage Group) running on Exchange Server on the application system computer\_app.company.com and backup system computer\_bck.company.com using the HP CA P9000 XP configuration, run:

```
omnicreatedl -ex2000 -datalist BS2 -device dev1 -storage_group "First
Storage Group" -storage_group "Second Storage Group" -split_mirror -sse
-remote computer_app.company.com computer_bck.company.com
```

### **Example 3 – P9000 XP Array:**

To create a ZDB specification BS3, using device dev1, for two stores (Public and Mailbox, part of First Storage Group) running on Exchange Server on the application system computer\_app.company.com and backup system computer\_bck.company.com using the HP CA+BC P9000 XP configuration, run:

```
omnicreatedl -ex2000 -datalist BS3 -device dev1 -storage_group "First Storage Group" -store "Public" "Mailbox" -split_mirror -sse -combined computer_app.company.com computer_bck.company.com
```

### **Example 4 – P6000 EVA Array:**

To create a ZDB-to-tape specification BS1, using the backup device dev1, for all storage groups running on Exchange Server on the application system computer1.company.com and backup system computer2.company.com, run:

```
omnicreatedl -ex2000 -datalist BS1 -device dev1 -snapshot -smis computer_app.company.com computer_bck.company.com -snapshot_type vsnap -snapshot_policy strict -storage_group "First Storage Group"
```

The omnicreatedl creates a transaction logs backup specification file First Storage Group (LOGS) computer1.company.com for First Storage Group log files backup (if it does not already exist). Data Protector attempts to create vsnaps, and if they cannot be created, the session aborts.

### **Example 5 – P6000 EVA Array:**

To create a ZDB-to-disk specification Exchange\_example to back up Site Replication Service using the backup device dev1, using the replica set with 5 replicas, run:

```
omnicreatedl -ex2000 -datalist Exchange_example -device dev1 -snapshot -smis computer1.company.com computer2.company.com -instant_recovery -snapshots 5 -annotation SRS
```

The omnicreatedl creates a transaction logs backup specification file SRS (LOGS) computer1.company.com for Site Replication Service log files backup (circular logging disabled). When omnib or Data Protector GUI is used to start backup, select ZDB to disk.

### **Example 6 – P6000 EVA Array:**

To create a ZDB-to-disk+tape specification Exchange\_example to back up Site Replication Service to the backup device dev1, using the replica set with 3 replicas and delay backup to tape for 50 minutes, run:

```
omnicreatedl -ex2000 -datalist Exchange_example -device dev1 -snapshot -smis computer1.company.com computer2.company.com -instant_recovery -snapshots 3 -wait_cloncopy 50 -annotation SRS
```

The omnicreatedl creates a transaction logs backup specification file SRS (LOGS) computer1.company.com for Site Replication Service log files backup (circular logging disabled). When omnib or Data Protector GUI is used to start backup, select ZDB to disk+tape.

## **Modifying ZDB specifications**

After you created an Exchange ZDB specification and a transaction logs backup specification using omnicreatedl, you can modify them using the Data Protector GUI.

---

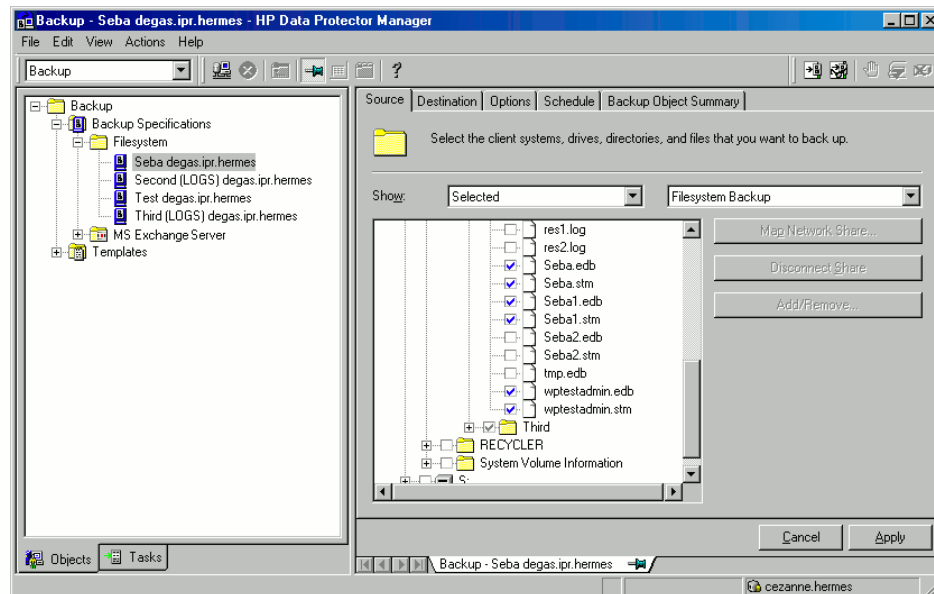
**NOTE:** To add/remove storage groups from a ZDB specification, create a new backup specification rather than modify the existing one. Changing the saved ZDB specification manually may alter the options Stop/quiesce the application command line and Restart the application command line, which are automatically defined when a new backup specification is created using omnicreatedl.

---

Proceed as follows using the Data Protector Manager:

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, then **Filesystem**, and click a ZDB specification. See Figure 78 (page 187).

**Figure 78 Exchange ZDB specification**



3. Modify backup devices, set ZDB options, schedule, and start backup. For a particular database, back up .edb and .stm files.

To create additional backup copies (mirrors), specify the desired number by clicking **Add mirror/Remove mirror** under the **Destination** tag. Select separate devices for backup and each mirror.

For information on object mirroring, see the online Help index: "object mirroring".

---

**NOTE:** Object mirroring is not supported for ZDB to disk.

---

## Checking Exchange files for consistency

Exchange Server allows you to perform a page level integrity check of the database files. It is run on the Exchange offline database during ZDB using `eseutil` (Exchange 2003) utilities. For more information, see *How To: Use the Eseutil Utility to Detect File Header Damage in Exchange 2003 (825088)* available at <http://support.microsoft.com>.

- 
- ❗ **IMPORTANT:** Page level integrity check is not supported on P6000 EVA Array.
- 

Proceed as follows:

1. Copy the utility to the backup system.

The `eseutil` resides in `Exchange_Server_home\bin` once the Exchange Server 2003 is installed.

---

**NOTE:** On P9000 XP Array, do not mirror the location where the utility resides, to prevent the file from being overwritten when the mirror is synchronized.

---

2. Write a script that runs `eseutil.exe` and save it, for example, as `integritycheck.bat`. Basically, the command looks like: `Path_to_eseutil\eseutil.exe /k Exchange_Server_database_file`. The `eseutil.exe` utility with the `/k` option tests the checksums on the Exchange Server database.

3. Modify the ZDB specification to include `integritycheck.bat` as post-exec. Execute the script on the backup system. For more information, see the online Help index: "pre- and post-exec commands".
4. In the backup specification, select `Leave the backup system enabled`. This ensures that the database is available on the backup system when the integrity check is started.

After backup, `esefile` or `esutil` perform page level integrity check against database files specified in `integritycheck.bat`.

## Scheduling backups

You can run unattended ZDB at specific times or periodically. For details on scheduling, see the online Help index: "scheduled backups".

---

**NOTE:** You cannot run ZDB to disk or ZDB to disk+tape if `-instant_restore` (P9000 XP Array) or `-instant_recovery` (P6000 EVA Array) is not selected in the backup specification.

---

### Scheduling example

To schedule a database ZDB at 8:00, 13:00, and 18:00 during weekdays:

1. In the **Schedule** property page, select the starting date in the calendar and click **Add** to open the **Schedule Backup** dialog box.
2. Under **Recurring**, select **Weekly**. Under **Time options**, select **8:00**. Under **Recurring Options**, select **Mon, Tue, Wed, Thu, and Fri**.

For ZDB sessions, the backup type is automatically set to **Full**.

For ZDB to disk or ZDB to disk+tape, specify the **Split mirror/snapshot backup** option.

Click **OK**.

3. Repeat [Step 1](#) and [Step 2](#) to schedule backups at 13:00 and 18:00.
4. Click **Apply** to save the changes.

## Starting backup sessions

Interactive backups are run on demand. They are useful for urgent backups or restarting failed backups.

---

**NOTE:** You cannot run ZDB to disk or ZDB to disk+tape if `-instant_restore` (P9000 XP Array) or `-instant_recovery` (P6000 EVA Array) is not selected in the backup specification.

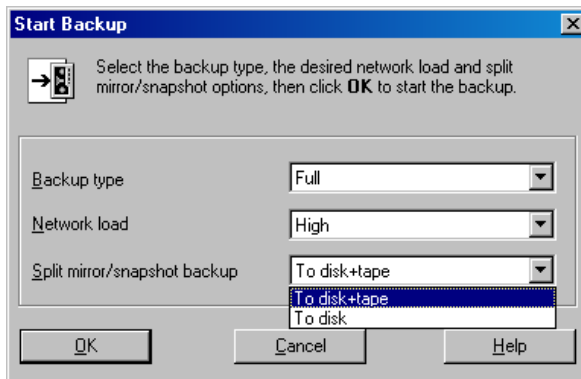
---

### Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **Filesystem**. Right-click the backup specification you want to use and select **Start Backup**.

3. Select **Network load**. For information on network load, click **Help**. Click **OK**.  
For ZDB sessions, the backup type is automatically set to **Full**.  
For ZDB to disk or ZDB to disk+tape, specify the **Split mirror/snapshot backup** option.

**Figure 79 Starting interactive backups**



Click **OK**.

## Using the Data Protector CLI

To start ZDB to tape or ZDB to disk+tape, run:

```
omnib -datalist Name
```

To start ZDB to disk, run:

```
omnib -datalist Name -disk_only
```

where *Name* is the backup specification name. For more information on `omnib`, see its man page.

## Restore

Data Protector offers restore from backup media to the application system on LAN (standard restore), where you can select among various restore options depending on your restore scenario, and instant recovery. See the below sections for more information.

### Considerations

- Instant recovery restores data from a replica on the backup system to the source volumes on the application system. Therefore, you cannot *selectively* restore objects (storage groups, stores, or Exchange Server) that do not reside on separate source volumes.
- Transaction logs must be backed up to perform rollforward recovery.

---

**NOTE:** In case of a disaster, Exchange Server must be installed and configured with the same database names and locations as before the disaster.

---

## Standard restore

You can restore Exchange objects to any Exchange Server with the same configuration as the original system within the Data Protector cell.

You can recover Exchange databases using:

- Point-in-time recovery  
Only data files (.edb and .stm) are restored. Databases are restored to their state at the backup time, and all data modified after that is lost.
- Rollforward recovery

Restores Exchange database files and transaction logs, and then replays the transaction logs. Databases are recovered to their last consistent state.



**TIP:** You can improve data transfer rate by configuring a backup device on the application system. For details, see the online Help index: “configuring, backup devices”. You can also perform a restore using another device. See the online Help index: “select, devices for restore”.

You can perform a filesystem restore of Exchange databases (.edb and .stm files) and transaction logs (.log files); however, to recover the database, you need to perform some additional steps. For details, see *Offline Backup and Restoration Procedures for Exchange (296788)* at <http://support.microsoft.com>.

For a detailed procedure on performing filesystem restore, see the online Help index: “standard restore procedure”. The procedures below provide only a general description of the restore process.

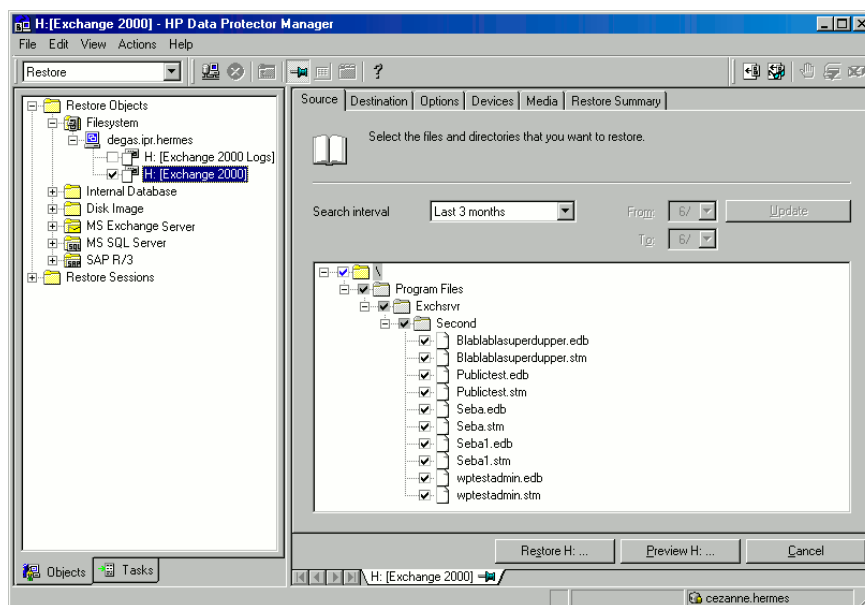
## Point-in-time recovery

- ① **IMPORTANT:** The procedure described below gives instructions on how to copy .edb and .stm files to the appropriate database location and is only part of the procedure described in *Offline Backup and Restoration Procedures for Exchange (296788)*.

Proceed as follows using the Data Protector Manager:

1. In the Context List, select **Restore**.
2. In the Scoping Pane, expand **Restore Objects, Filesystem**, and then select the client (backup system). Mark the drive letter on which the database resides.

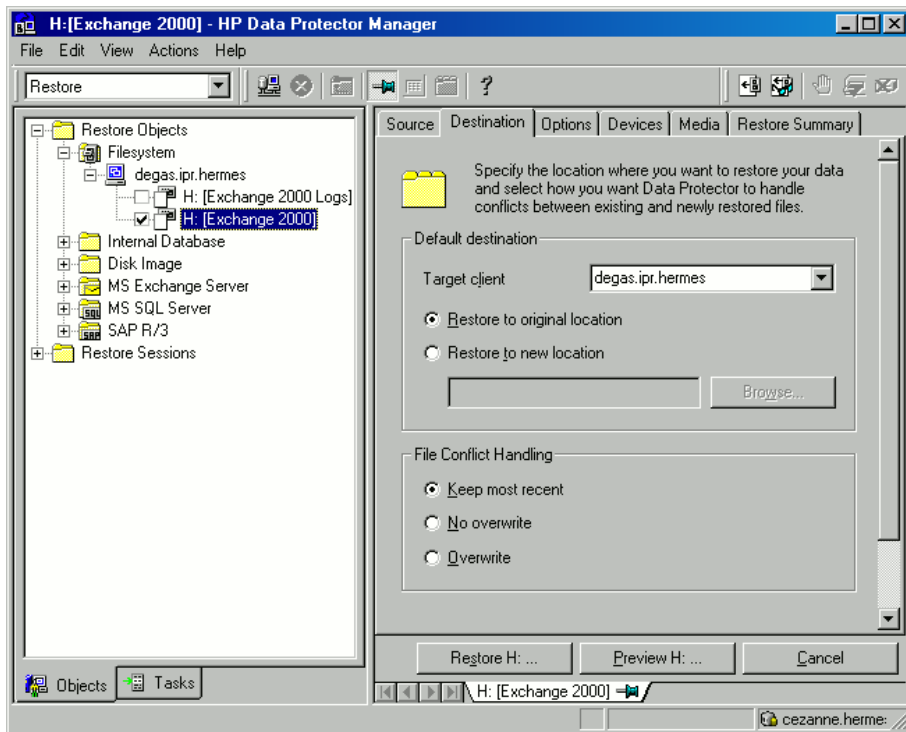
**Figure 80 Restoring Exchange database**



Exchange database consists of two files: *name.edb* and *name.stm*. For a particular database, select both files. You can restore a storage group by selecting the storage group folder when the entire storage group was backed up and all databases reside in the same directory.

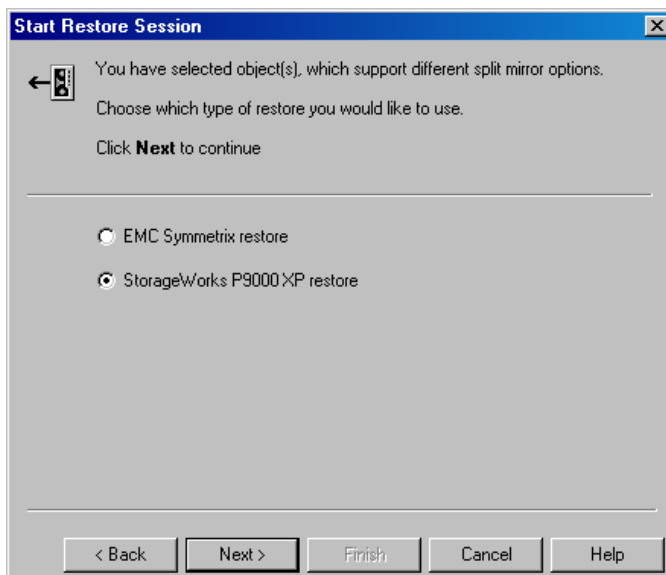
3. Under the **Destination** tab, select the application system as the **Target client**.

**Figure 81 Selecting the application system**



4. Set restore options. For information, see the online Help index: “restore, options”.
5. In the **Devices** page, select the devices to be used for the restore.  
For more information on how to select devices for a restore, see the online Help index: “restore, selecting devices for”.
6. Click **Restore**. The **Start Restore Session** dialog box is displayed.
7. Specify **Report level** and **Network load**.
8. **P9000 XP Array, EMC:**  
Select **StorageWorks P9000 XP restore**. Click **Next**.

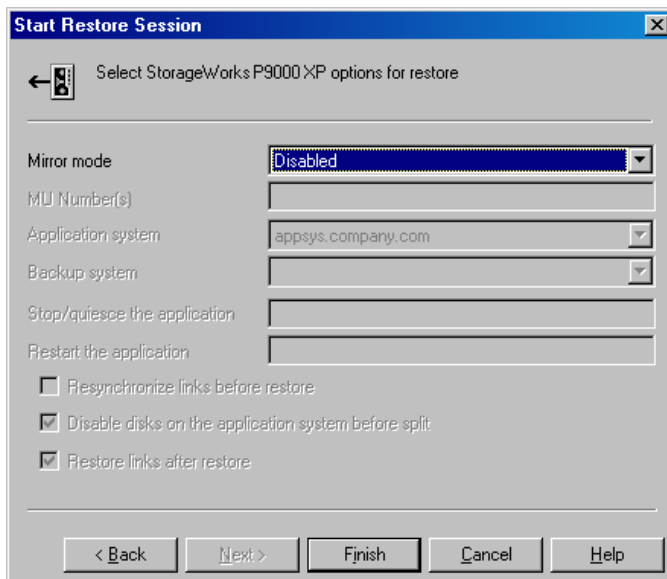
**Figure 82 Selecting StorageWorks P9000 XP restore**



9. **P9000 XP Array:**

In the **Start Restore Session** dialog box, select **Disabled** in the **Mirror mode** drop-down list. This sets a direct restore from the backup media to the application system on LAN.

**Figure 83 P9000 XP Array restore option**



10. Click **Finish** to start restore.

## Rollforward recovery

- ❗ **IMPORTANT:** The procedure described below gives instructions on how to copy .edb, .stm, and .log files to the appropriate database location and only part of the procedure described in the *Offline Backup and Restoration Procedures for Exchange (296788)*.

When following the Microsoft procedure, use the procedure below to copy .edb, .stm and .log files to the appropriate database location.

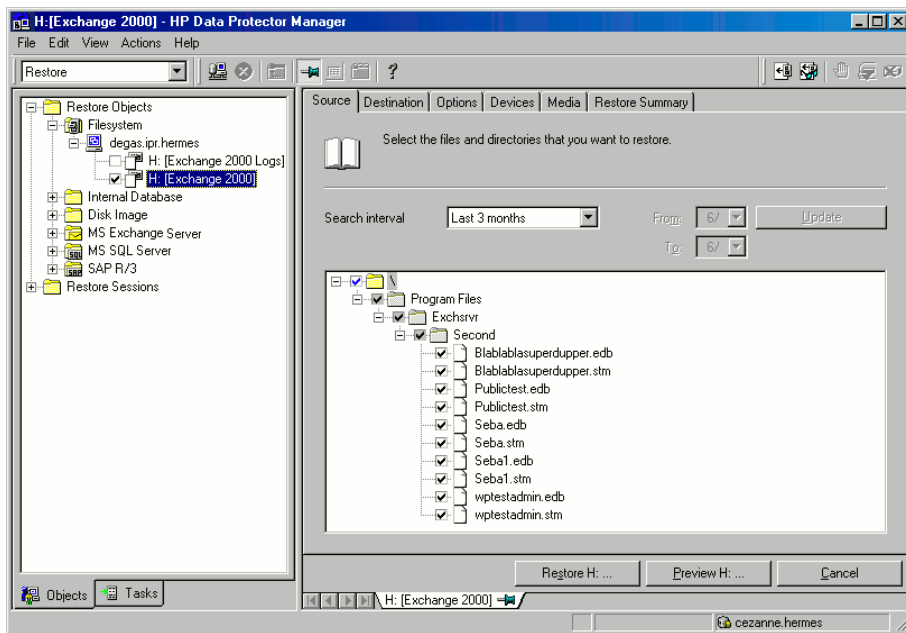
The procedure below needs to be utilized twice: for the database and for transaction logs restore.

Proceed as follows using the Data Protector Manager:

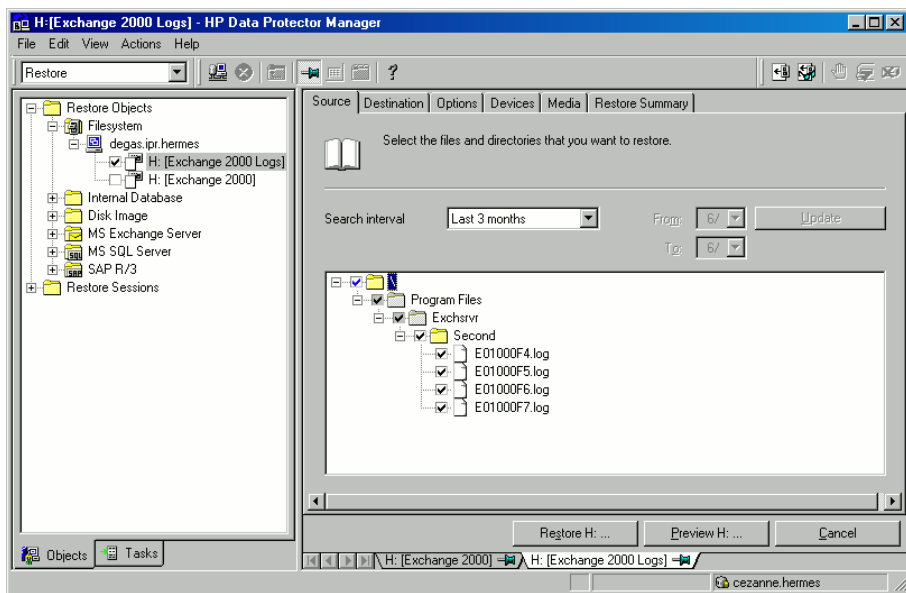
1. In the Context List, select **Restore**.
2. Proceed as follows:
  - To restore the database, expand **Restore Objects, Filesystem**, and the name of the backed up server. Then select the **[Exchange 2000]** object as shown in [Figure 84 \(page 193\)](#).
  - To restore transaction logs, expand **Restore Objects, Filesystem**, and the name of the backed up server. Then select the **[Exchange 2000 Logs]** object as shown in [Figure 85 \(page 193\)](#).



**Figure 84 Restoring Exchange database**



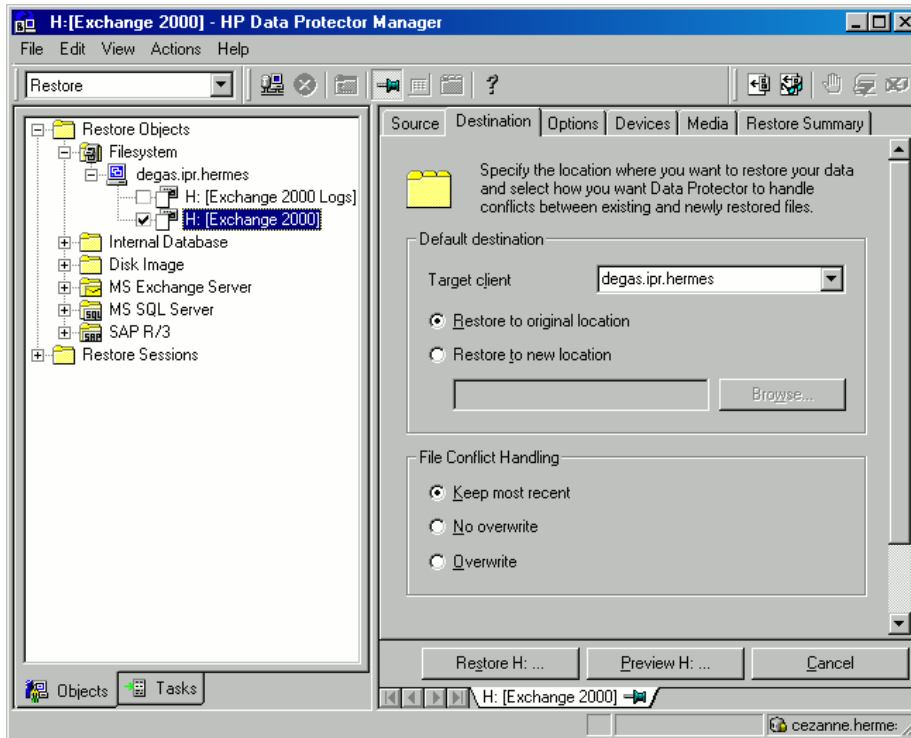
**Figure 85 Restoring log files**



3. Make the following selections in the Results Area:
- To restore the Exchange database, select .edb and .stm files. For a particular database, select both files. You can restore a storage group by selecting the storage group folder when the entire storage group was backed up and all databases reside in the same directory.
  - To restore transaction logs, select .log files.

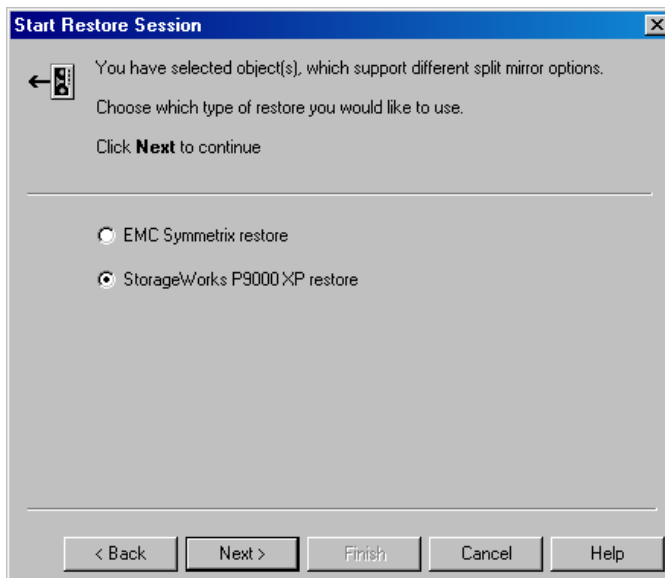
4. Select the application system as the **Target client** under the **Destination** tab.

**Figure 86 Selecting the application system**



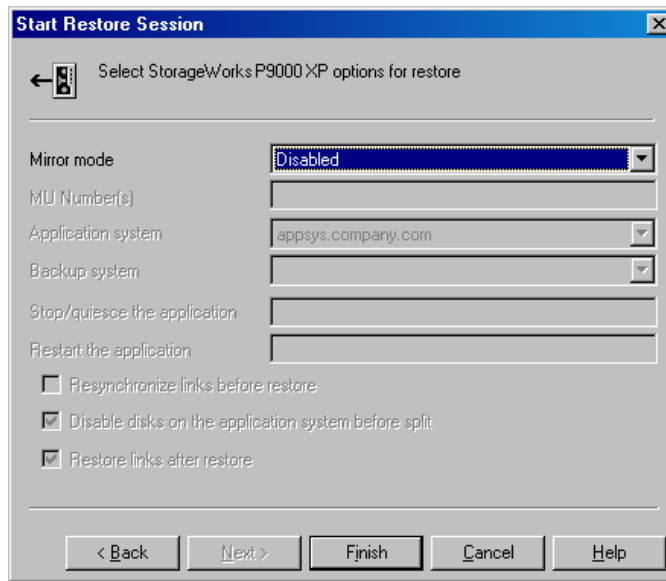
5. Set restore options. For information, see the online Help index: "restore, options".
6. Click **Restore**. The **Start Restore Session** dialog box is displayed.
7. Specify **Report level** and **Network load**. Click **Next**.
8. **P9000 XP Array, EMC:**  
Select **StorageWorks P9000 XP restore**. Click **Next**.

**Figure 87 Selecting StorageWorks P9000 XP restore**



9. **P9000 XP Array:**  
In the **Start Restore Session** dialog box, select **Disabled** in the **Mirror mode** drop-down list. This sets a direct restore from the backup media to the application system on LAN.

**Figure 88 P9000 XP Array restore option**



10. Click **Finish** to start restore.

## Instant recovery

See the *HP Data Protector Zero Downtime Backup Concepts Guide* and *HP Data Protector Zero Downtime Backup Administrator's Guide* for general information on instant recovery.

You can recover Exchange databases using:

- **Point-in-time recovery**  
Databases are restored to their state at the backup time, and all data modified after that is lost.
- **Rollforward recovery**  
Restores Exchange database files and transaction logs, and then replays the transaction logs. Databases are recovered to their last consistent state.

## Point-in-time recovery

For detailed procedure on how to recover the Exchange Server, see *Offline Backup and Restoration Procedures for Exchange (296788)* at <http://support.microsoft.com>.

When following the Microsoft procedure, see the *HP Data Protector Zero Downtime Backup Administrator's Guide* for information on copying backed up .edb and .stm files to the appropriate database using instant recovery.

## Rollforward recovery

For detailed procedure on how to recover the Exchange Server, see *Offline Backup and Restoration Procedures for Exchange (296788)* at <http://support.microsoft.com>.

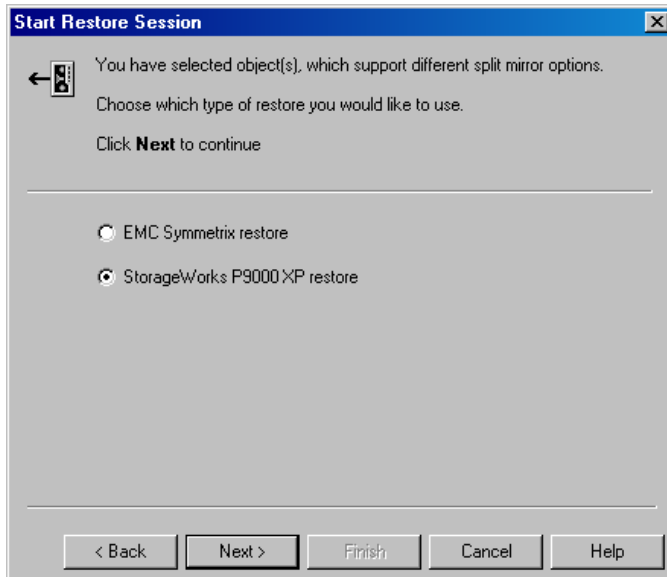
When following the Microsoft procedure, see the *HP Data Protector Zero Downtime Backup Administrator's Guide* for information on copying backed up .edb and .stm files to the appropriate database using instant recovery.

The procedure below gives instructions on how to restore Exchange transaction logs using the Data Protector GUI when following the Microsoft procedure:

1. In the Context List, select **Restore**.
2. In the Scoping Pane, expand **Restore Objects, Filesystem**, and then the name of the backed up server. Select the **[Exchange 2000 Logs]** object as shown in [Figure 85 \(page 193\)](#).

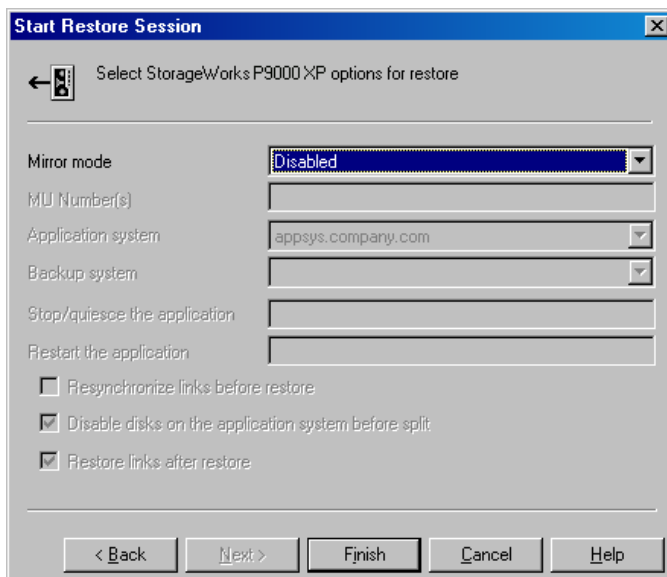
3. Under the **Destination** tab, select the application system as the **Target client as shown in Figure 86 (page 194)**.
4. Set restore options. For information, see the online Help index: "restore, options".
5. Click **Restore**. The **Start Restore Session** dialog box is displayed.
6. Specify **Report level** and **Network load**. Click **Next**.
7. **P9000 XP Array, EMC:**  
Select **StorageWorks P9000 XP restore**. Click **Next**.

**Figure 89 Selecting StorageWorks P9000 XP restore**



8. **P9000 XP Array:**  
In the **Start Restore Session** dialog box, select **Disabled** in the **Mirror mode** drop-down list. This sets a direct restore from the backup media to the application system on LAN.

**Figure 90 P9000 XP Array restore option**



9. Click **Finish** to start restore.

## Troubleshooting

This section lists general checks and verifications, plus problems you might encounter when using the Data Protector Exchange Server integration. Start at “Problems” (page 198). If you cannot find a solution there, perform general checks and verifications.

For general Data Protector troubleshooting information, see the *HP Data Protector Troubleshooting Guide*.

For general ZDB, restore, and instant recovery related troubleshooting, see the *HP Data Protector Zero Downtime Backup Administrator's Guide*.

### Before you begin

- Ensure that the latest official Data Protector patches are installed. On how to verify this, see the online Help index: “patches”.
- See the *HP Data Protector Product Announcements, Software Notes, and References* for general Data Protector limitations, as well as recognized issues and workarounds.
- See <http://www.hp.com/support/manuals> for an up-to-date list of supported versions, platforms, and other information.

### Checks and verifications

If your configuration, backup, or restore failed:

- Check that Exchange Server services (Microsoft Exchange System Attendant and Microsoft Exchange Information Store) are running.
- Using Exchange System Manager, check that all stores to be backed up are mounted and all stores to be restored are dismounted.
- Perform a backup of the Exchange Information Store using Windows Backup. If the backup fails, fix Exchange Server problems first, and then perform a backup using Data Protector.
- Ensure that the Cell Manager is correctly set on Exchange Server by checking the following registry entry:

`HKEY_LOCAL_MACHINE\SOFTWARE\Hewlett-Packard\OpenView\OmniBack II\Site`  
Its name and value must be `CellServer` and “*Cell Manager hostname*”, respectively.

- Examine system errors reported in `Data_Protector_home\log\debug.log` on Exchange Server functioning as a Data Protector client.

Additionally, examine the errors reported in the Windows Event log.

- Check if the following directories exist on the Data Protector Cell Manager:  
`Data_Protector_home\config\server\barlists\msese`  
`Data_Protector_home\config\server\barschedules\msese`
- Make a test filesystem backup and restore of the problematic client. For information, see the online Help.
- Create a backup specification to back up to a null or file device and run the backup. If the backup succeeds, the problem may be related to backup devices. See the *HP Data Protector Troubleshooting Guide* for instructions on troubleshooting devices.
- Try to restart the Microsoft Exchange Server and start the backup again.
- Check that the `Exchange_home\bin` directory is added to the Windows `Path` environment variable. For details, see “Configuring the integration” (page 179).
- If you cannot mount the storage after a successful restore, check that LOGS storage on the same storage group is also restored.

- Define a directory for temporary log files in the **Restore** context. Check if the specified directory exists. If it does not, create it or specify another existing directory.
- To restore to another system, make sure Exchange Server is installed on that system and has the same organization and site names as the restored server.

## Problems

### Problem

#### Exchange integration is unable to initiate clsEx2000 class

When using a disk array of the HP P6000 EVA Disk Array Family, the following error is reported in the debug.log file on the Exchange Server system during the omncreatedl command session or during the execution of the Stop/quiesce the application command line:

Unable to initiate clsEx2000 class. Error: -2147221164

### Action

Register omniex2000.dll by running the following command from *Data\_Protector\_home\bin*: regsvr32 omniex2000.dll

### Problem

#### Omncreatedl cannot create a backup specification for the log file backup

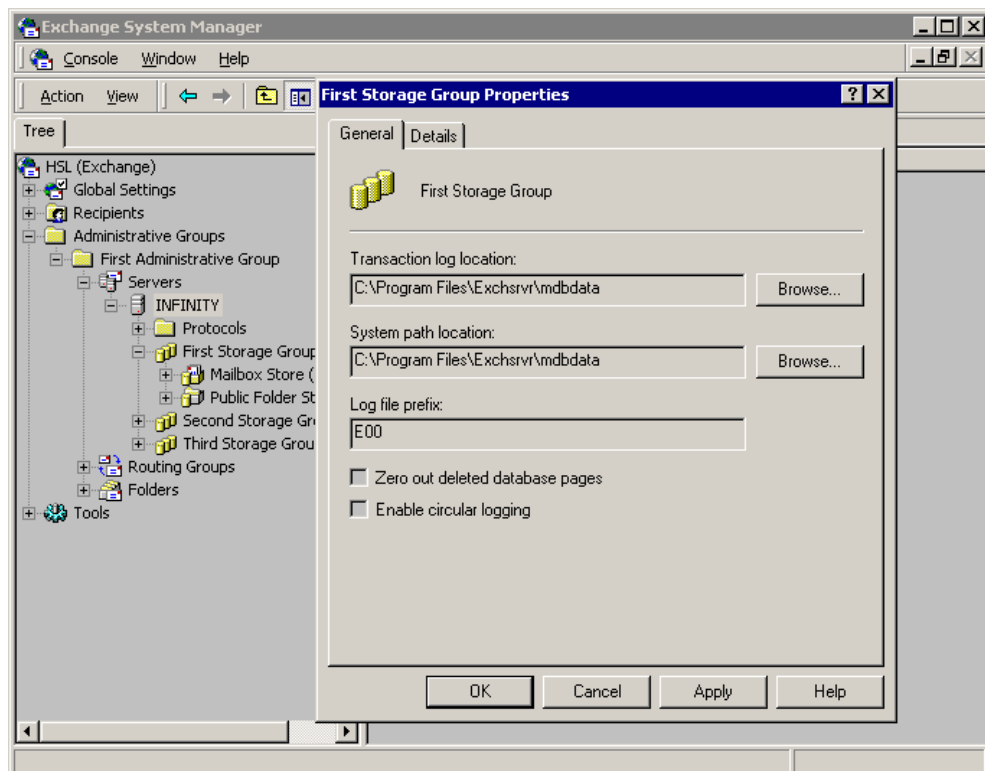
The following error is reported:

Cannot create log datalist for name.

### Action

Check if circular logging is disabled on the Exchange Server for the storage group. If not, disable it.

**Figure 91 Disabling circular logging**



## Problem

### **Omnicreatedl cannot create a backup specification and the session is aborted**

The following error is reported:

```
[ERROR] Could not obtain any data for backup from host app_sys.
Make sure the Exchange server is running and Data Protector Exchange
Integration installed.
```

## Actions

- If the application is cluster-aware, specify the `-virtualSrv` parameter.
- Check the user rights on the application system client. You must have the `Save backup specification` user right to create and save a backup specification.
- Check that `Exchange_home\bin` is listed in the `Path` system variable. For instructions, see [“Creating ZDB specifications” \(page 160\)](#).

## Problem

### **Error starting service**

If the KMS service password is not kept on disk, the following error is reported:

```
[Major] From: OB2BAR_main@computer.company.com "" Time: 10/20/2011 5:17:24
PMError starting service.
```

Reconfigure KMS to store the KMS service password on disk. For information, see *XADM: How to Change the KMS Service Password Startup Location (196129)*. Then restart the backup.

---

# 5 Data Protector Microsoft Exchange Server 2010 ZDB integration

## Introduction

This chapter explains how to configure and use the Data Protector Microsoft Exchange Server 2010 ZDB integration. It describes concepts and methods you need to understand to back up and restore Microsoft Exchange Server 2010 mailbox databases (**databases**).

Both standalone environments and Database Availability Group (**DAG**) environments are supported. The Data Protector Microsoft Exchange Server 2010 integration is based on the Volume Shadow Copy Service (**VSS**) technology. For details on VSS concepts, see the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

### Backup

During backup, databases can be used actively (**online backup**). In DAG environments, you can back up active and/or passive database copies.

As Microsoft Exchange Server, ZDB disk array, and VSS are involved, you can specify different kinds of backup types:

- Microsoft Exchange Server backup types
- VSS backup types
- ZDB backup types

You can select from among the following Microsoft Exchange Server backup types:

- Full
- Copy
- Incremental
- Differential

You can select from among the following ZDB backup types:

- ZDB-to-disk
- ZDB-to-disk+tape
- ZDB-to-tape

You can select from among the following VSS backup types:

- Local or network backup
- VSS transportable

For details on the backup types, see [“Backup types” \(page 206\)](#).

### Restore

You can restore Microsoft Exchange Server databases using standard restore or instant recovery.

During restore, each database can be restored using a different restore method. The following methods are available:

- Repair all passive copies with failed status
- Restore to the latest state
- Restore to a point in time<sup>1</sup>

1. This method is supported only if you have Microsoft Exchange Server 2010 SP1 installed.



- Restore to a new mailbox database<sup>2</sup>
- Restore files to a temporary location

This chapter provides information specific to the Exchange Server 2010 integration. For limitations, see the *HP Data Protector Product Announcements, Software Notes, and References*. For general Data Protector procedures and options, see the online Help.

## Integration concepts

Data Protector integrates with the Microsoft Exchange Server through the Data Protector Microsoft Exchange Server 2010 integration agent, which channels communication between the Data Protector Session Manager and the clients in the Microsoft Exchange Server environment. The agent communicates with the Microsoft Exchange Server through the Microsoft Exchange Management Shell and uses VSS to back up data.

## Supported environments

Data Protector supports Microsoft Exchange Server Database Availability Group environments (**DAG environments**) as well as environments with standalone Microsoft Exchange Server systems (**standalone environments**).


### Standalone environments

In a standalone Microsoft Exchange Server environment, each Microsoft Exchange Server system stands on its own.

In one session, you can back up databases from only one Microsoft Exchange Server system. Data Protector sends backup and restore requests directly to the Microsoft Exchange Server system.

#### Figure 92 Standalone environment (example)

Standalone environment



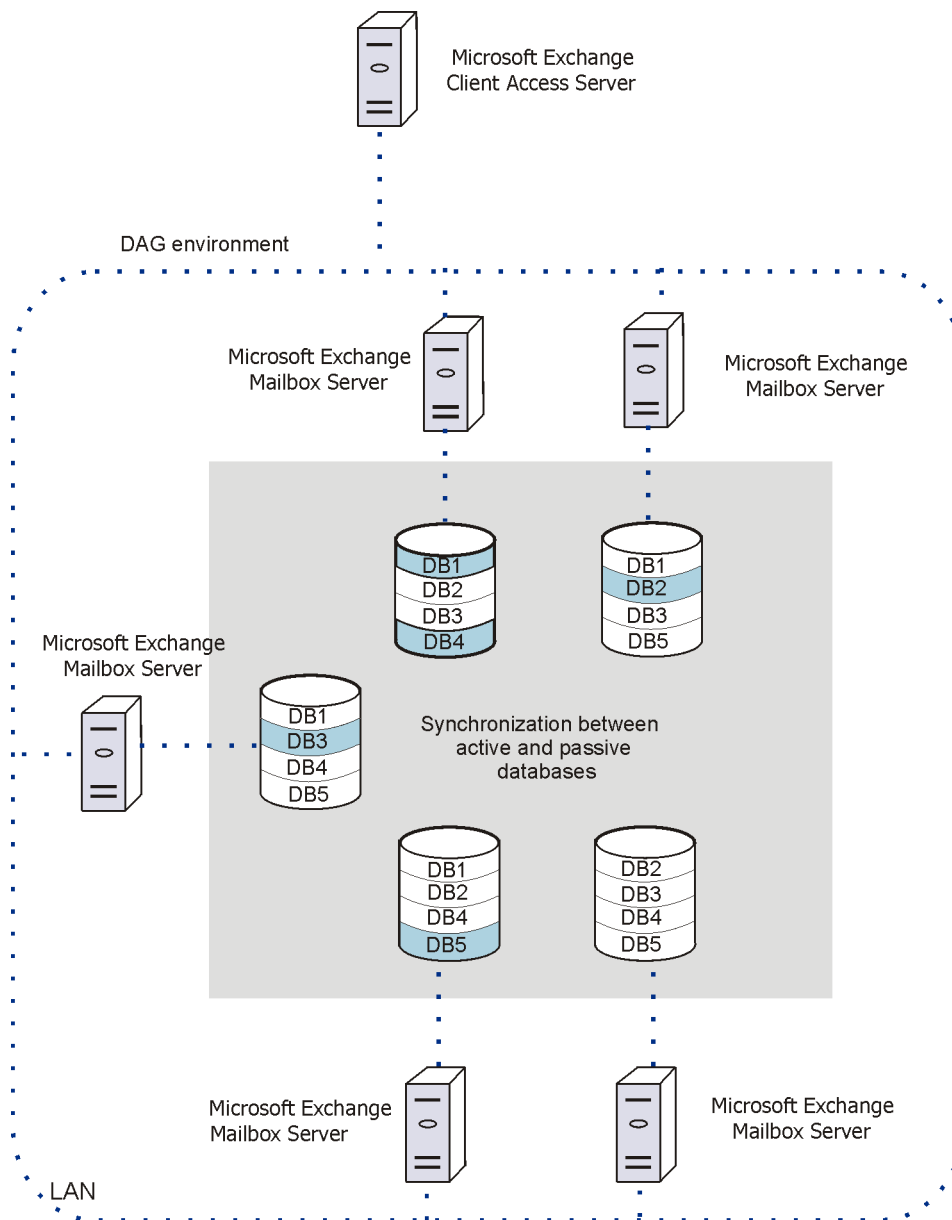
### DAG environments

In a DAG environment, Data Protector communicates with the DAG using one of the Microsoft Exchange Server systems (the one that is currently active in the environment). All backup and restore requests are sent there.

In one session, you can back up active and/or passive database copies from different Microsoft Exchange Server systems that belong to the same DAG.

2. This method offers the possibility of restoring to a recovery database.

**Figure 93 DAG environment (example)**



In [Figure 93 \(page 202\)](#), active databases are shaded in blue.

If a database has multiple passive copies, you can specify which particular passive copy you want to back up, using one of the following backup policies:

- minimize the number of hosts
- lowest activation preference
- highest activation preference
- shortest replay lag time
- longest replay lag time
- longest truncation lag time

You can also specify from which Microsoft Exchange Server systems database copies should not be backed up.

For a brief description of the activation preference number, replay lag time, and the truncation lag time, see "[Microsoft Exchange Server parameters in DAG environments](#)" ([page 203](#)).

**Table 22 Microsoft Exchange Server parameters in DAG environments**


| Parameter                    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Activation preference number | The activation preference number determines which passive copy is activated if multiple passive copies meet the same criteria; the copy assigned the lowest activation preference number is activated.                                                                                                                                                                                                                                                                                                     |
| Replay lag time              | The <code>ReplayLagTime</code> parameter plays a role when synchronizing a passive copy with the active copy. As soon as a log file at the active copy side is filled up, it is copied to the passive copy side. By default, the newly copied log is also applied to the passive copy database files. However, if the passive copy <code>ReplayLagTime</code> parameter is set to a value greater than 0, the log is applied with a lag, creating a lagged database copy.<br>The maximum value is 14 days. |
| Truncation lag time          | The <code>TruncationLagTime</code> parameter specifies how long the Microsoft Exchange Replication service waits before truncating log files that have already been applied to the database files.<br>The maximum value is 14 days.                                                                                                                                                                                                                                                                        |

## Configuring the integration

### Prerequisites

- Ensure that you have correctly installed and configured the Microsoft Exchange Server 2010 environment.
  - For supported versions, platforms, devices, and other information, see the latest support matrices at <http://www.hp.com/support/manuals>.
  - For information on installing, configuring, and using Microsoft Exchange Server 2010, see the Microsoft Exchange Server 2010 documentation.
  - If you intend to use the restore method **Restore to a point in time**, ensure that you have Microsoft Exchange Server 2010 SP1 installed.
- If you intend to run Incremental and Differential backup sessions, ensure that circular logging is disabled.
- If you plan to run instant recovery sessions, it is advisable to keep Microsoft Exchange Server databases on separate storage volumes. Also, keep a database's files (.edb and .log) on separate storage volumes. See [Figure 94 \(page 203\)](#). Such a configuration provides better restore granularity.

Ensure that storage volumes on different Microsoft Exchange Server systems are the same size. Otherwise, you may experience problems during copy-backup instant recovery sessions.

**Figure 94 Where to keep database files**

- Ensure that you have correctly installed Data Protector. On how to install Data Protector in various architectures, see the *HP Data Protector Installation and Licensing Guide*.  
Ensure that the following Data Protector components are installed on all the Microsoft Exchange Server systems:
  - MS Exchange Server 2010 Integration
  - MS Volume Shadow Copy Integration
  - The appropriate Data Protector disk array agent

---

**NOTE:** For VSS transportable backup sessions, the MS Volume Shadow Copy Integration component and the appropriate Data Protector disk array agent must also be installed on the backup systems.

---

In DAG environments, the DAG virtual system (host) must also be imported to the Data Protector Cell. On how to import a client to a Data Protector Cell, see the online Help index: "importing, client systems".

- For limitations, see "Limitations and recommendations" in the *HP Data Protector Product Announcements, Software Notes, and References*.

## Before you begin

- Configure devices and media for use with Data Protector.
- To test whether a Microsoft Exchange Server system and the Cell Manager communicate properly, configure and run a Data Protector filesystem backup and restore on every Microsoft Exchange Server client in your environment.

## Configuring user accounts

Backup and restore sessions are started by the Data Protector Inet service, which by default runs under the Windows local System user account. Consequently, a backup or restore session is performed using the same user account.

However, you can specify that the Data Protector Inet service should use a different Windows domain user account to start a session:

- To perform a backup session under a different user account, specify the **Specify OS user** option (see "[Specifying view type](#)" (page 210)) when creating a backup specification.
- To perform a restore session under a different user account, specify the **User name** and **Group/Domain name** options in the Options page (when performing standard restore, see "[Restore options](#)" (page 228)) or Advanced page (when performing instant recovery, see "[Instant recovery – advanced](#)" (page 237)).

Before you specify a different Windows domain user account, configure the user account as follows:

1. Grant the user appropriate permissions to back up and restore Microsoft Exchange Server databases.
2. Add the user to the Data Protector admin or operator user group. For details on adding users, see the online Help index: "adding users".

3. Save the user and its password to a Windows registry on the Microsoft Exchange Server system on which you plan to start the integration agent (`e2010_bar.exe`). To save the user account, use the Data Protector `omniinetpasswd` or `omnicc` command.

---

**NOTE:** The user account saved in the Windows registry will be used by the Data Protector `Inet` service when needed.

---

For details on setting accounts for the `Inet` service user impersonation, see the online Help index: “`Inet` user impersonation”.

#### Example

To save the user `jane` from the domain `HP` and with the password `mysecret` to a Windows registry, log on to the Microsoft Exchange Server system and, from the `Data_Protector_home\bin` directory, run:

```
omniinetpasswd -add jane@HP mysecret
```

## Backup

When you back up a Microsoft Exchange Server database, the following files are automatically backed up:

- database files (`.edb`)
- transaction logs (`.log`)
- checkpoint files (`.chk`)

However, depending on the Microsoft Exchange Server backup type you select, not all files are always backed up. For details, see “[Microsoft Exchange Server backup types](#)” (page 206).

## Backup types

As Microsoft Exchange Server, ZDB disk array, and VSS are involved, you can specify different kinds of backup types:

- Microsoft Exchange Server backup types
- ZDB backup types
- VSS backup types

### Microsoft Exchange Server backup types

You can select from among the following Microsoft Exchange Server backup types:

**Table 23 Backup types**

|              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Full         | <p>Backs up the database file (.edb), transaction logs (.log), and checkpoint files (.chk), and then truncates the transaction logs.</p> <p><i>DAG environment only:</i> If multiple copies of a database are selected for backup, Data Protector first performs a Full backup of the passive copy that has the fewest logs applied to the database file, and then performs a Copy backup of all the remaining copies, with the active copy being backed up last. The copies are backed up sequentially due to a Microsoft Exchange Server VSS writers limitation.</p> |
| Copy         | <p>Backs up the database file (.edb), transaction logs (.log), and checkpoint files (.chk), without truncating the transaction logs.</p>                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Incremental  | <p>Backs up the transaction logs (.log) that have been created since the last Full or Incremental backup, and then truncates the transaction logs.</p> <p><i>DAG environment only:</i> If multiple copies of a database are selected for backup, Data Protector backs up the transaction logs of only one copy (one of the passive copies is selected).</p>                                                                                                                                                                                                            |
| Differential | <p>Backs up the transaction logs (.log) that have been created since the last Full backup, without truncating the transaction logs.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                |

#### NOTE:

An Incremental or Differential backup of a database cannot be performed:

- If a Full backup has not been performed.
- If an Incremental backup is started just after a Differential backup has been performed, or the reverse.

### ZDB backup types

You can select from among the following ZDB backup types:

- ZDB-to-disk
- ZDB-to-disk+tape
- ZDB-to-tape

**NOTE:** For Microsoft Exchange Server Incremental and Differential backup types, only the ZDB-to-tape backup type is available.

### VSS backup types

You can select from among the following VSS backup types:

- Local or network backup
- VSS transportable

For details, see the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

## Backup parallelism

- During a backup session, copies of different databases are backed up in parallel, however, copies of the same database are not, due to a Microsoft Exchange Server VSS writers limitation.
- If multiple backup sessions that intend to back up the same database are started in parallel, only the session that first locks the database can back up the database; the other sessions cannot. In DAG environments, this also applies if backup sessions intend to back up different copies of the same database; only the session that first locks the database (that is, all its copies) can back up the database copies; the other sessions cannot.

---

**NOTE:** This behavior ensures that the construction of a restore chain is valid. For example, suppose that several Full backup sessions that intend to back up the same database are started in parallel. If all the sessions backed up the database, it might happen that the session given the latest SessionID is not the one that backed up the database last. For details on restore chains, see [“Restore chain” \(page 221\)](#).

---

## Replica rotation in DAG environments

Data Protector Microsoft Exchange Server 2010 integration enables you to perform ZDB sessions in multi-system environments (DAG environments). This brings some changes to the existing replica rotation functionality.

In standalone environments, the replica rotation functionality works as it used to; it limits the number of backups that are kept in the Data Protector IDB database for instant recovery purposes. For example, if the **Number of replicas rotated** option is set to 1, only one backup session is available for instant recovery at a time. If you start another backup session (with the same backup specification), backup storage volumes created in the previous backup session are deleted before new ones are created and the previous session is removed from the Instant Recovery context.

In DAG environments, a database can be backed up from different systems in different sessions. This introduces the changes. For example, suppose you want to back up the database DB1, which is active on `node1.company.com`, and the database DB2, which is active on `node2.company.com`. Suppose the **Number of replicas rotated** option is set to 1 and your backup policy is such that the active copy is always backed up. After the backup, the Data Protector VSSDB database contains the following entries:

**Backup 1** (2011/10/05-1):

- 2011/10/05-1:node1.company.com (containing DB1)
- 2011/10/05-1:node2.company.com (containing DB2)

Let us suppose that a failover occurs and the database DB1 becomes active on `node2.company.com`. We start another backup session. Now both databases are backed up from `node2.company.com`. As a result, the VSSDB database contains the following entries:

**Backup 2** (2011/10/05-2):

- 2011/10/05-1:node1.company.com (containing DB1)
- 2011/10/05-2:node2.company.com (containing DB1 and DB2)

Note that the entry `2011/10/05-1:node2.company.com` is no longer in the VSSDB database as it has been rotated out (the corresponding backup storage volumes have been deleted) due to the replica rotation functionality.

---

**NOTE:** Entries are rotated per system and not per session. As a result, entries that are created in the same session can be rotated out at different points in time (that is, in different sessions).

---

Let us say that another failover occurs and both databases become active on `node1.company.com`. We start another session. Now both databases are backed up from `node1.company.com`. As a result, the VSSDB database contains the following entries:

### Backup 3 (2011/10/05-3):

- 2011/10/05-3:node1.company.com (containing DB1 and DB2)
- 2011/10/05-2:node2.company.com (containing DB1 and DB2)

Note that the entry 2011/10/05-1:node1.company.com has been rotated out as well. Since both parts created in the session 2011/10/05-1 have been rotated out, the session 2011/10/05-1 can no longer be used for instant recovery (it is removed from the Instant Recovery context).

However, here is the problem: if you go to the Instant Recovery context and select the session created in Backup 1 after you have performed Backup 2, the source page shows that both databases can be restored. But this is not true since the entry 2011/10/05-1:node2.company.com (containing DB2) has been rotated out in Backup 2. If you select DB2 for restore and start instant recovery, the session will fail. Therefore, you have to ensure that you restore databases from sessions for which the necessary entries in the VSSDB database still exist.

## Backup considerations

- *Backup strategy:*

Choose one of the following strategies to back up your data:

- Full
- Full, Incremental, Incremental, ...
- Full, Differential, Differential, ...
- Full, Copy, Incremental, ..., Copy, Incremental, ...

---

❗ **IMPORTANT:** An Incremental backup session cannot be followed by a Differential backup session, nor the other way around. You must first run a Full backup session.

---

- *Active copies as opposed to passive copies:*

There is no difference between the active and passive copy, except in the currently active log file (at the active copy side), which is not copied to the passive copy side until the file is filled up (that is, reaches 1 MB). Consequently, if you back up a passive copy, the transactions in the currently active log file are not included.

- *Lagged database copies:*

Backing up a lagged database copy is equivalent to backing up a non-lagged database copy. If you restore from the backup of a lagged database copy, files are not only restored, but logs are also applied to the database file, returning the database to its most recent state. However, restoring the logs and applying them to the database file is time-consuming and, therefore prolongs the restore session. Also note that you need enough disk space to restore all the necessary logs.

On the other hand, restoring from the backup of a lagged database copy enables you to restore the database to a point in time before the backup was taken. Restore the database without performing database recovery and mounting. Then remove unwanted logs, and finally recover and mount the database.

- *Public folders:*

Backup of public folders with activated replication is not supported.

- *Concurrent backup sessions:*

Backup sessions that back up the same database cannot run in parallel.



## Creating backup specifications

Create a backup specification using the Data Protector GUI (**Data Protector Manager**).

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **MS Exchange 2010 Server**, and click **Add Backup**.
3. In the Create New Backup dialog box, specify **Backup type** (VSS backup type). For details, press **F1**. Click **OK**.
4. In **Application system**, select the Microsoft Exchange Server system that you want to back up from. In a DAG environment, select the DAG virtual system or a Microsoft Exchange Server system.

---

**NOTE:** The **Application system** drop-down list contains all clients that have the Data Protector MS Exchange Server 2010 Integration component installed. In a DAG environment, the list contains also the DAG virtual system (host).

The backup session (that is, the integration agent `e2010_bar.exe`) will be started on the client that you specify here. If you select a DAG virtual system, the integration agent is started on the currently active Microsoft Exchange Server node. To find out which node is currently active, see “[Tip](#)” (page 242).

**NOTE:** To back up public folders residing on a Microsoft Exchange Server system that is part of a DAG environment, select the Microsoft Exchange Server system and not the DAG virtual system (host). If you select the DAG virtual system, you can back up only databases that belong to the DAG, of which the public folders database is not part.

---

Depending on the VSS backup type you selected, specify the following:

- If you selected **Local or network backup**, in **Provider**, select **Use hardware provider**.
- If you selected **VSS transportable backup**, specify **Backup system**.

Specify ZDB-specific options.

For details, press **F1** or see the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

---

**NOTE:** In an HP P9000 XP Disk Array Family resync mode environment, the maximum number of replica storage volumes (S-VOL) that can be created for a given storage volume (P-VOL) is limited by the hardware provider configuration — MU range (maximum is 3). To be able to perform Incremental and Differential sessions, the **Number of replicas rotated** option in the backup specification must be set to one less than the MU range. In this way, one replica storage volume is kept free for Incremental and Differential backup sessions.

---

Click **Next**.

5. In you selected the DAG virtual system (host), specify **View Type** to define how Microsoft Exchange Server databases should be organized in the next page (source page):

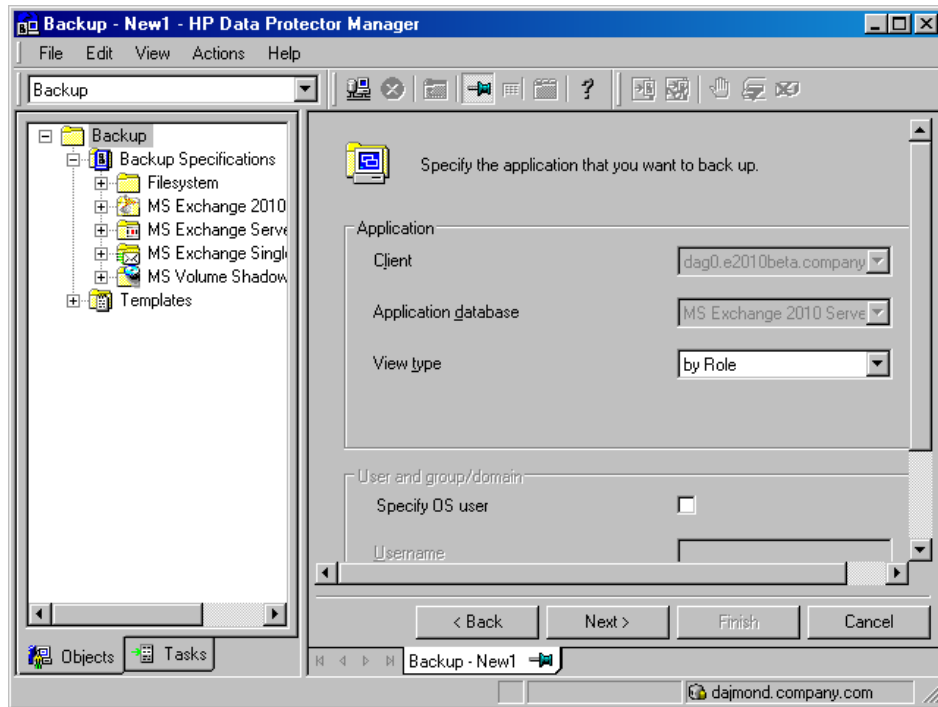
**By Role**

All databases in the DAG are displayed.

**By Client**

All clients in the DAG are displayed, together with all the databases (active or passive) residing on them. Active databases have the label (active) appended at the end. Passive databases have no label.

**Figure 95 Specifying view type**



For details on other options, press **F1**.

Click **Next**.

6. Select which Microsoft Exchange Server databases to back up.

---

**NOTE:** *DAG environment only:*

In a single session, you can back up either of the following:

- multiple databases, but only one copy of each
  - asingle database, but multiple copies of it
-

Figure 96 Selecting databases (DAG environment – by role)

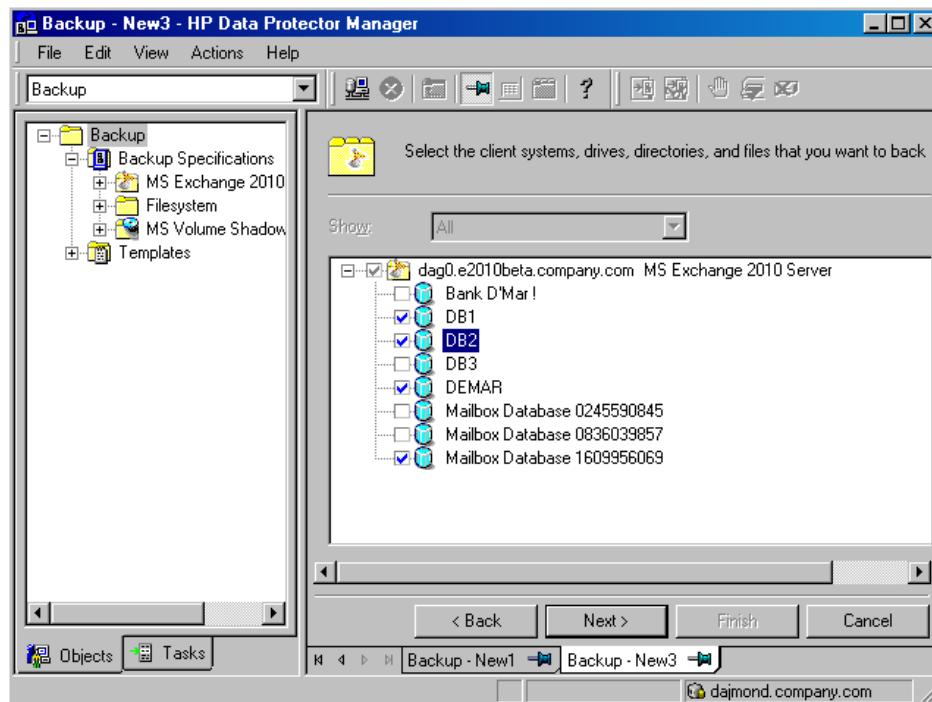
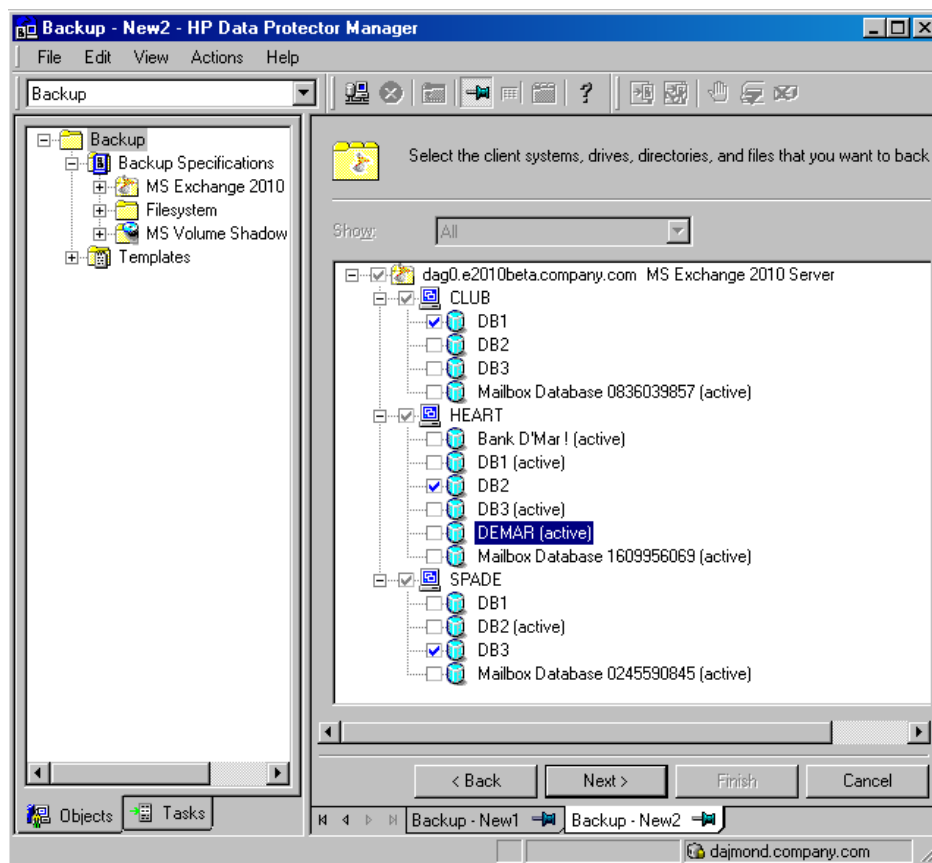



Figure 97 Selecting databases (DAG environment – by client)




**Figure 98 Selecting databases (standalone environment)**



7. This applies in DAG environments if you selected the **By Role** view type. Specify the backup policy options.

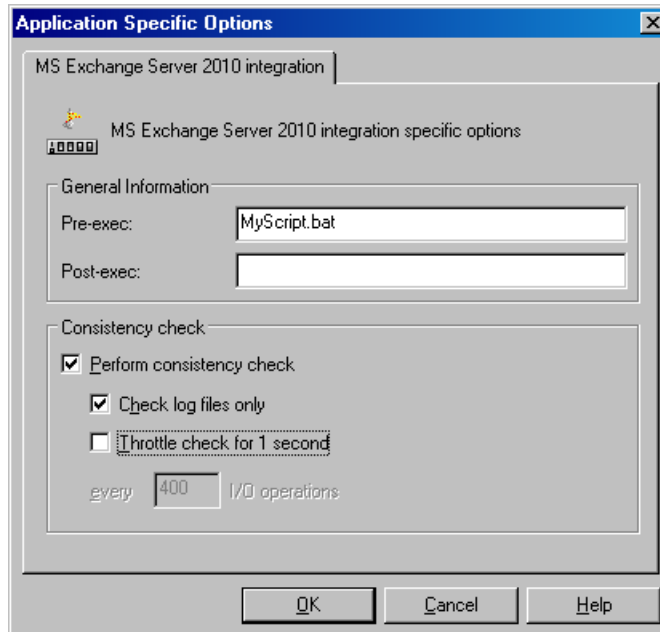
**Figure 99 Backup policy options**



For details, see “Backup policy options” (page 213).

8. Select which devices to use for the backup.  
To specify device options, right-click the device and click **Properties**. Specify the number of parallel backup streams in the **Concurrency** tab and which media pool to use.  
Click **Next**.
9. Set backup options.

**Figure 100 Application-specific option**



For information on application-specific backup options, see [“Application-specific backup options”](#) (page 214).

Click **Next**.

10. Optionally, schedule the backup. See [“Scheduling backup specifications”](#) (page 215).  
Click **Next**.
11. Save the backup specification, specifying a name and a backup specification group.



**TIP:** Preview your backup specification before using it for real. See [“Previewing backup sessions”](#) (page 216).

**Table 24 Backup policy options**

| Options                                       | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |                                               |                                                                                                                                                                                                                                                                                                               |                                             |                                                                                                                  |
|-----------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------|------------------------------------------------------------------------------------------------------------------|
| <b>Back up active database</b>                | If this option is selected, the active copy is backed up.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |                                               |                                                                                                                                                                                                                                                                                                               |                                             |                                                                                                                  |
| <b>Back up passive copy</b>                   | <p>If this option is selected, a passive copy is backed up. If a database has multiple passive copies, specify which particular copy you want to back up, using one of the following policies:</p> <table border="1"> <tr> <td><b>minimize the number of hosts</b> (default)</td><td>If this option is selected, the minimum number of clients is involved in the backup. For example, if databases to be backed up have each a passive copy residing on the same client, they are all backed up from this client (and not one database from one client and another database from another client).</td></tr> <tr> <td><b>lowest/highest activation preference</b></td><td>If this option is selected, the database copy with the lowest/highest activation preference number is backed up.</td></tr> </table> | <b>minimize the number of hosts</b> (default) | If this option is selected, the minimum number of clients is involved in the backup. For example, if databases to be backed up have each a passive copy residing on the same client, they are all backed up from this client (and not one database from one client and another database from another client). | <b>lowest/highest activation preference</b> | If this option is selected, the database copy with the lowest/highest activation preference number is backed up. |
| <b>minimize the number of hosts</b> (default) | If this option is selected, the minimum number of clients is involved in the backup. For example, if databases to be backed up have each a passive copy residing on the same client, they are all backed up from this client (and not one database from one client and another database from another client).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |                                               |                                                                                                                                                                                                                                                                                                               |                                             |                                                                                                                  |
| <b>lowest/highest activation preference</b>   | If this option is selected, the database copy with the lowest/highest activation preference number is backed up.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |                                               |                                                                                                                                                                                                                                                                                                               |                                             |                                                                                                                  |

**Table 24 Backup policy options** (continued)

| Options                            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |                                                                                                                                                   |
|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
|                                    | <b>shortest/longest replay lag time</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | If this option is selected, the database copy with the shortest/longest replay lag time is backed up.                                             |
|                                    | <b>longest truncation lag time</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | If this option is selected, the database copy with the longest truncation lag time is backed up.                                                  |
|                                    | For a brief description of the activation preference number, replay lag time and transaction lag time parameters, see “ <a href="#">Microsoft Exchange Server parameters in DAG environments</a> ” (page 203). For details, see the Microsoft Exchange Server 2010 documentation.                                                                                                                                                                                                                                                                                                                                                                                                                                              |                                                                                                                                                   |
|                                    | <b>Use active copy for backup if no passive copies are available</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Available if <b>Back up passive copy</b> is selected. If this option is selected, the active copy is backed up when no passive copy is available. |
| <b>Back up all copies</b>          | <p>Available if only one database is selected for backup.</p> <p>If this option is selected, all copies (active and passive) are backed up. This is useful when you create ZDB-to-disk or ZDB-to-disk+tape backups (that is, backups that can be used for instant recovery). If multiple copies are backed up, during instant recovery, multiple copies can be restored, as each copy has its own replica storage volumes to be restored from. For details, see “<a href="#">Instant recovery in DAG environments</a>” (page 221).</p> <p>When you create a ZDB-to-tape backup, it is enough that a single copy is backed up; you can restore different copies of a database from the ZDB-to-tape backup of a single copy.</p> |                                                                                                                                                   |
| <b>Exclude clients from backup</b> | Creates a list of clients. The database copies that reside on these clients are not backed up.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |                                                                                                                                                   |

**Table 25 Application-specific backup options**

| Options                                                                                                                                 | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Pre-exec, Post-exec</b>                                                                                                              | <p>Specifies which command line to run on a Microsoft Exchange Server system before (pre-exec) or after (post-exec) the backup.</p> <p>The command line is run only on the Microsoft Exchange Server system on which the backup session is started (that is the system on which the Data Protector Microsoft Exchange Server 2010 integration agent e2010_bar.exe is started).</p> <p>Type only the name of the command and ensure that the command is located in the <code>Data_Protector_home\bin</code> directory on the same system. Do not use double quotes.</p> <p><i>DAG environment only:</i> If you selected the DAG virtual system (host) in the <b>Application system</b> option, ensure that the command is located on the currently active node. To find out which Microsoft Exchange Server node is currently active, see “<a href="#">Tip</a>” (page 242).</p>                                                                                                                                                              |
| <b>Perform consistency check</b><br><code>[-exch_check</code><br><code>[-exch_throttle Value]  </code><br><code>-exch_checklogs]</code> | <p>If this option is selected, Microsoft Exchange Server checks the consistency of a database's backup data. If this option is not selected, the session finishes earlier, but the backup data consistency is not guaranteed.</p> <p>The check is performed on the replica storage volumes after the backup data is created. If the data is found corrupt, the replica storage volumes are discarded and the database backup fails.</p> <p>Default: selected</p> <p>If the <b>Check log files only</b> option is selected, only the backup data of the log files is checked, which is enough for Microsoft Exchange Server to guarantee data consistency.</p> <p>Default: selected</p> <p>By default, the consistency check is I/O intensive, which can negatively affect disk performance. The <b>Throttle check for 1 second</b> option throttles down the consistency check of the database file .edb to lessen impact on the disk performance. Specify after how many input/output operations the check should stop for one second.</p> |

**Table 25 Application-specific backup options** *(continued)*

| Options | Description                                                                              |
|---------|------------------------------------------------------------------------------------------|
|         | This option is not available if only the log files are checked.<br>Default: not selected |

## Modifying backup specifications

To modify your backup specification, click its name in the Scoping Pane of the Backup context, then click the appropriate tab, and apply the changes.

**NOTE:** To see all databases in the source page, not just those you selected, select **All** in the **Show** option. In a DAG environment, this not only shows all databases, but also updates the current status of databases (active or passive).

## Scheduling backup specifications

You can schedule a backup session to start automatically at specific times or periodically. For details on scheduling, see the online Help index: “scheduled backups”.

### Scheduling example

To schedule Differential backups at 8:00, 13:00, and 18:00 during week days:

1. In the **Schedule** property page of the backup specification, select the starting date in the calendar and click **Add** to open the **Schedule Backup** dialog box.
2. Under **Recurring**, select **Weekly**. Under **Time options**, select **8:00**. Under **Recurring Options**, select **Mon, Tue, Wed, Thu, and Fri**. See “[Scheduling backups](#)” (page 215). Under **Session options**, select **Differential** from the **Backup type** drop-down list.  
Click **OK**.
3. Repeat [Step 1](#) and [Step 2](#) to schedule Differential backups at 13:00 and 18:00.
4. Click **Apply** to save the changes.

**Figure 101 Scheduling backups**

**Schedule Backup**

Specify the desired backup time, frequency, duration, and type.

**Recurring**

☐ None  
☐ Daily  
☒ Weekly  
☐ Monthly

**Time options**

Time: 8:00 AM  
☐ Use starting  
3/ 5/2010

**Recurring options**

Every 1 week(s) on

☐ Sun ☒ Mon ☒ Tue ☒ Wed ☒ Thu ☒ Fri ☐ Sat

**Session options**

Backup type: Differential  
Network load: ☒ High ☐ Medium ☐ Low  
Backup protection: Default

OK Cancel Help

## Previewing backup sessions

Preview the backup session to test it. You can use the Data Protector GUI or CLI.

### Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **MS Exchange 2010 Server**. Right-click the backup specification you want to preview and click **Preview Backup**.
3. Specify **Backup type** and **Network load**. Click **OK**.

The message `Session completed successfully` is displayed at the end of a successful preview.

### Using the Data Protector CLI

1. Log in to the Cell Manager or to any client with the Data Protector User Interface component installed, under a user account that is configured as described in [“Configuring user accounts”](#) (page 204).
2. Go to the following directory:

**Windows systems:** `Data_Protector_home\bin`

**UNIX systems:** `/opt/omni/lbin`

3. Run:

```
omnib -e2010_list BackupSpecificationName -test_bar
```

## What happens during the preview?

The following are tested:

- Communication between the Microsoft Exchange Server system on which the backup session is started and the Cell Manager
- If each selected database has at least one copy available for backup after the **Backup policy** options and **Client filtering** options have been applied (this applies to backup specifications that contain backup policy options)
- If the selected databases are ready to be backed up (that is, they should not be dismantled, suspended, or in a failed state)

## Starting backup sessions

Interactive backups are run on demand. They are useful for urgent backups or restarting failed backups.

To start a backup, use the Data Protector GUI or CLI.

### Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, and then **MS Exchange 2010 Server**. Right-click the backup specification you want to use and click **Start Backup**.
3. Specify **Backup type** and **Network load**. Click **OK**.

The message `Session completed successfully` is displayed at the end of a successful backup session.

### Using the Data Protector CLI

1. Log in to the Cell Manager or to any client that has the Data Protector User Interface component installed under a user account that is configured as described in [“Configuring user accounts”](#) (page 204).



2. Go to the following directory:

**Windows systems:** `Data_Protector_home\bin`

**UNIX systems:** `/opt/omni/lbin`

3. Run:

```
omnib -e2010_list BackupSpecificationName [-barmode E2010Mode]
[LIST_OPTIONS]
```

where *E2010Mode* is one of the following:

`full|copy|incr|diff`

The default is `full`.

For *ListOptions*, see the omnib man page or the *HP Data Protector Command Line Interface Reference*.

### Examples

To start a Full backup using the backup specification *MyDatabases*, run:

```
omnib -e2010_list MyDatabases -barmode full
```

To start a Differential backup using the same backup specification, run:

```
omnib -e2010_list MyDatabases -barmode diff
```

## Backup objects

For each database (copy), Data Protector creates the following backup objects:

- *Database file object*
  - `ClientName:/Microsoft Exchange Writer (Exchange Information Store)/Microsoft Information Store/DBID/File [MSVSSW-APP]`  
(standalone database or active copy)
  - `ClientName:/Microsoft Exchange Writer (Exchange Replication Service)/Microsoft Information Store/DBID/File [MSVSSW-APP]`  
(passive copy)
- *Log file object*
  - `ClientName:/Microsoft Exchange Writer (Exchange Information Store)/Microsoft Information Store/DBID/Logs [MSVSSW-APP]`  
(standalone database or active copy)
  - `ClientName:/Microsoft Exchange Writer (Exchange Replication Service)/Microsoft Information Store/DBID/Logs [MSVSSW-APP]`  
(passive copy)
- *Database object*  
`ClientName:/DBID/DBName [E2010]`  
The database object contains information needed to construct the restore chain. For details on restore chains, see [“Restore chain” \(page 221\)](#).
- *VSS metadata object*  
`/BackupSession/Metadata [MSVSSW-APP]`

Information on whether the objects were successfully backed up or not is saved in the Data Protector IDB. On how to retrieve the information from the IDB, see [“Finding information for restore” \(page 222\)](#).

## Restore

You can restore Microsoft Exchange Server data by performing a standard restore or instant recovery session. For details, see [“Standard restore” \(page 223\)](#) and [“Instant recovery” \(page 232\)](#).

- ❗ **IMPORTANT:** After you have restored a database, start a Full backup session for the database. Otherwise, the subsequent Incremental and Differential backup sessions will fail.
- 

### Considerations

- A Microsoft Exchange Server database that was backed up using the Data Protector Microsoft Volume Shadow Copy Service integration cannot be restored using the Data Protector Microsoft Exchange Server 2010 integration, nor the reverse.

## Restore methods

There are various reasons for restoring a Microsoft Exchange Server database. Here are some examples:

- The database has become corrupt.
- The synchronization between an active and passive database copy is broken, but you want to avoid reseeding the passive copy, or simply the resume operation does not work.
- The database needs to be restored to a different point in time.
- The database's backup data needs to be restored for investigation purposes.
- The database's backup data needs to be restored to a recovery database in order to extract individual mailboxes or mailbox files.
- The database's backup data needs to be restored to a dial tone database.

To suit your needs, Data Protector offers different restore methods. You can choose from among the following:

- **Repair all passive copies with failed status**
- **Restore to the latest state**
- **Restore to a point in time**
- **Restore to a new mailbox database**
- **Restore to a temporary location**

You can specify different restore methods for different databases in the same session.

---

**NOTE:** The first three methods restore backup data to the original database and are therefore only available if the original database still exists. The last two methods restore backup data to a new location.

---

### Repair all passive copies with failed status

This method is available only for databases that are part of a DAG. It is useful if some of a database's passive copies become corrupt, acquiring the status `Failed` or `FailedAndSuspended`. The method automatically restores all the corrupt passive copies from the backup created in the last backup session (and the corresponding restore chain). After the data is restored, the copies are synchronized with the active copy, provided that the **Resume database replication** option is selected.

### Restore to the latest state

This method is used to restore a corrupt database to the latest possible point in time. Data Protector restores the database from the backup created in the last backup session (and the corresponding restore chain). For details, see [“Restore chain” \(page 221\)](#).

Once the files are restored, all the logs (not only those restored from the backup, but also any existing logs) are replayed to the database file.

---

**NOTE:** *DAG environment only:*

When a passive copy is restored, Microsoft Exchange Server ensures that the logs are replayed to the database file in accordance with the `ReplayLagTime` parameter setting.

---

## Restore to a point in time

This method is used to restore a database to a specific point in time.

---

**NOTE:** *Standard restore:*

When you restore a standalone database or active copy, the existing `.log` and `.chk` files are renamed (a `.keep` extension is added to their names). This feature is useful when you restore files without performing database recovery. It enables you to apply additional logs to the database file; just delete the `.keep` extension of the log files that you also want to be applied and start a database recovery manually. In this way, you can fine-tune the point in time the database is restored to.

When you restore a passive copy, the existing files are deleted.

---

Once the files are restored, the logs are replayed to the database file (`.edb`) if the **Perform database recovery** option is selected.

---

**NOTE:** *DAG environment only:*

- When a passive copy is restored, Microsoft Exchange Server ensures that the logs are replayed to the database file in accordance with the `ReplayLagTime` parameter setting.
  - For passive copies that are not restored, a full reseed is required once the restore session completes.
- 

## Restore to a new mailbox database

This method is used to restore data to a different database, either because the original database no longer exists or in order to move the data elsewhere.

Using it, you can restore data also to a Microsoft Exchange Server recovery database.

---

**NOTE:** *Instant recovery:*

This option is not available for replica types whose data can only be restored to the original storage volumes.

---

## Restore files to a temporary location

Using this method, you can restore database files to a location of your choice.

- When you restore from a Differential or Incremental backup session, you can restore the complete restore chain or only the files (`.log`) backed up in the selected session.
  - When you restore data from a Full backup session, you have an option to restore only the database file (`.edb`).
- 

**NOTE:** *Instant recovery:*

This option is not available for replica types whose data can only be restored to the original storage volumes.

---

## Restore destination

Backup data can be restored:

- to an existing database (standalone database, active copy, passive copy),
- to a new database,
- to a temporary location.

### Restoring to a standalone database

Restore to the original standalone database (standalone environment) progresses as follows:

1. The database is dismounted.
2. Backup data is restored.
3. Optionally, the newly-restored logs (and pre-existing ones if you are performing the **Restore to the latest state** method) are replayed to the database file `.edb` and the database is mounted.

To restore to the original standalone database, use one of the following restore methods:

- **Restore to the latest state**
- **Restore to a point in time**

### Restoring to an active copy

Restore to the active copy (DAG environment) progresses as follows:

1. The database is dismounted.
2. All replications are suspended.
3. Backup data is restored.
4. Optionally, the newly-restored logs (and pre-existing ones if you are performing the **Restore to the latest state** method) are replayed to the database file `.edb` and the database is mounted.

To restore to the active copy, use one of the following restore methods:

- **Restore to the latest state**
- **Restore to a point in time**

### Restoring to a passive copy

Restore to a passive copy (DAG environment) progresses as follows:

1. The replication is suspended.
2. Backup data is restored.
3. Optionally, the replication with the active copy is resumed.

To restore to a passive copy, use one of the following restore methods:

- **Restore all passive copies with failed status**
- **Restore to the latest state**
- **Restore to a point in time**

### Restoring data to a new database

Restore to a new database progresses as follows:

1. A new mailbox database is created.
2. Backup data is restored to the new database.

---

**NOTE:** If you restore to a recovery database, first the backup data is restored and then a recovery database is created.

---

To restore data to a new mailbox database or recovery database, use the **Restore to a new mailbox database** restore method.

## Restoring data to a temporary location

You can restore the database file (.edb and/or .log and/or .chk) to a client and directory of your choice. Select the **Restore files to a temporary location** restore method.

## Instant recovery in DAG environments

When you back up a database in a DAG environment, you can decide whether to back up all its copies or only a single copy. If all copies are backed up, during instant recovery, all copies can be restored, as each copy has its own replica storage volumes to be restored from. If only a single copy is backed up, note the following:

- In most cases, only one database copy can be restored from the backup of a single database copy. Some replica types are directly connected with the source (*dependent* replica types) while others are *independent*, allowing the data to be restored to a different location. With the latter, you can restore either the original or a different database copy.

---

**NOTE:** The following replica types are independent:

- HP P9000 XP Disk Array Family (split mirror replica type in the VSS compliant mode)
- HP P6000 EVA Disk Array Family (snapclone replica type)

---

For dependent replica types, Data Protector automatically grays out those clients in the **Target Nodes** option whose database copies were not backed up, because these copies cannot be restored.

- The only replica type that enables you to restore multiple database copies from the backup of a single database copy is snapclone (only with the HP P6000 EVA Disk Array Family). However, you must also ensure that both the **Restore using HP StorageWorks P6000 EVA SMI-S** and **Copy replica data to the source volumes** instant recovery options are selected, in which case data from the replica storage volumes is sequentially copied to multiple locations, restoring one database copy after the other.

## Restore chain

By default, when you select a Differential or Incremental backup session for restore, Data Protector restores not only the logs (.log) backed up in the selected session but also files backed up in preceding sessions (**restore chain**):

- If a Differential backup session is selected, Data Protector restores:
  1. The .edb file and .log files backed up in the most recent Full or Copy backup session.
  2. The .log files backed up in the selected Differential backup session.
- If an Incremental backup session is selected, Data Protector restores:
  1. The .edb file and .log files backed up in the most recent Full or Copy backup session.
  2. The .log files backed up in all subsequent Incremental backup sessions, up to the selected Incremental backup session.
- If a Full or Copy backup session is selected, Data Protector restores the .edb file and .log files backed up in the selected session.

---

**NOTE:**

- If the **Restore to the latest state** method is used, .log files from the Full or Copy backup session are not restored.
  - The only method that enables you to restore only .log files backed up in the selected Incremental or Differential session is **Restore to a temporary location**.
- 

## Restore chain during instant recovery

During an instant recovery session, you first select which:

- Full (ZDB-to-disk or ZDB-to-disk+tape) session or
- Copy (ZDB-to-disk or ZDB-to-disk+tape) session

to use for instant recovery. From database specific options you then specify whether additional logs should also be restored, by selecting a subsequent Incremental or Differential session in the **Restore additional logs until** option.

In an instant recovery session Data Protector restores:

1. The .edb file and .log files backed up in the selected Full or Copy backup session.
2. The .log files backed up in the selected Differential backup session or in all subsequent Incremental backup sessions, up to the selected Incremental backup session.

## Restore parallelism

If device concurrency allows, database copies are restored in parallel, except in the following cases:

- If database copies were backed up from the same client, but are now restored to different clients.
- If backup data of the same database copy is used as restore source for multiple database copies.

## Finding information for restore

You can retrieve information about backup sessions (such as information on the backup type and media used, and the messages reported during the backup) from the Data Protector IDB.

To retrieve information, use the Data Protector GUI or CLI.

## Using the Data Protector GUI

1. In the Context List, click **Internal Database**.
2. In the Scoping pane, expand **Objects** or **Sessions**.

If you expand **Objects**, backup objects are sorted according to the Microsoft Exchange Server databases for which they were created.

---

**NOTE:** The backup object name contains the database GUID. To find out which GUID belongs to which database, see the *database object /DB\_GUID/DB\_Name*.

For example, the *database object* for the database DB1 with the GUID

08bca794-c544-4e27-87e8-533fb81fd517 is:

/08bca794-c544-4e27-87e8-533fb81fd517/DB1

---

If you expand **Sessions**, backup objects are sorted according to the sessions in which they were created. For example, backup objects created in the session 2010/02/7-7 are listed under 2010/02/7-7.

To view details on a backup object, right-click the backup object and click **Properties**.



**TIP:** To view the messages reported during the session, click the **Messages** tab.

## Using the Data Protector CLI

1. Log in to the Cell Manager or to any Microsoft Exchange Server client with the Data Protector User Interface component installed under a user account that is configured as described in [“Configuring user accounts”](#) (page 204).

2. Go to the following directory:

**Windows systems:** `Data_Protector_home\bin`

**UNIX systems:** `/opt/omni/lbin`

3. Get a list of Microsoft Exchange Server backup objects created in a backup session:

```
omnidb -session SessionID
```

4. Get details on a backup object:

```
omnidb -e2010 BackupObjectName -session SessionID -catalog
```

Here is one example of a backup object name:

```
devy.company.com:/08bca794-c544-4e27-87e8-533fb81fd517/DB1
```

For details, see the `omnidb` man page or the *HP Data Protector Command Line Interface Reference*.

## Standard restore

Standard restore is restore from backup data residing on Data Protector media (for example, a tape). Such data is created in ZDB-to-disk+tape and ZDB-to-tape sessions.

You can restore multiple Microsoft Exchange Server databases in the same standard restore session, specifying a different restore method for each database. For details, see [“Restore methods”](#) (page 218).

To perform a standard restore, use the Data Protector GUI or CLI.


## Restoring using the Data Protector GUI

1. In the Context List, click **Restore**.
2. In the Scoping Pane, expand **MS Exchange 2010 Server**, expand the DAG virtual system or standalone Microsoft Exchange Server system and click **MS Exchange 2010 Server**.
3. In the Source page, Data Protector displays all Microsoft Exchange Server 2010 databases backed up from the selected DAG or standalone environment.

Select which Microsoft Exchange Server databases to restore.

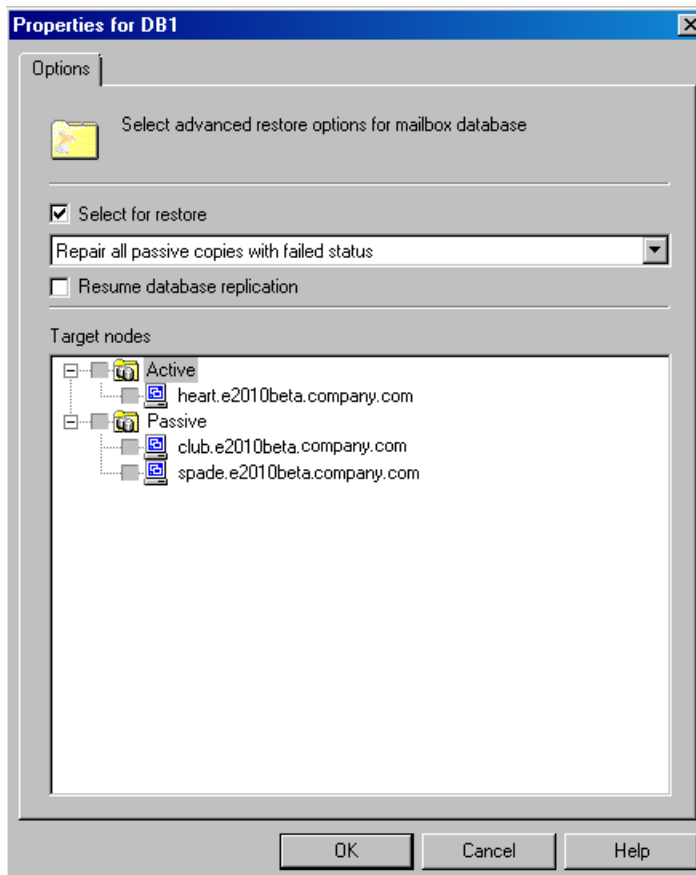
When you select a database, the Properties for Database dialog box is displayed automatically. Specify a restore method and click **OK**. For databases that are part of a DAG, the default restore method is **Repair all passive copies with failed status**. For standalone databases, the default is **Restore to the latest state**.

**Figure 102 Selecting databases for restore**



To change the restore method, right-click the database and click **Properties**.


**Figure 103 Repair all passive copies with failed status**



For details, see [“Repair all passive copies with failed status”](#) (page 239).




**Figure 104 Restore to the latest state**



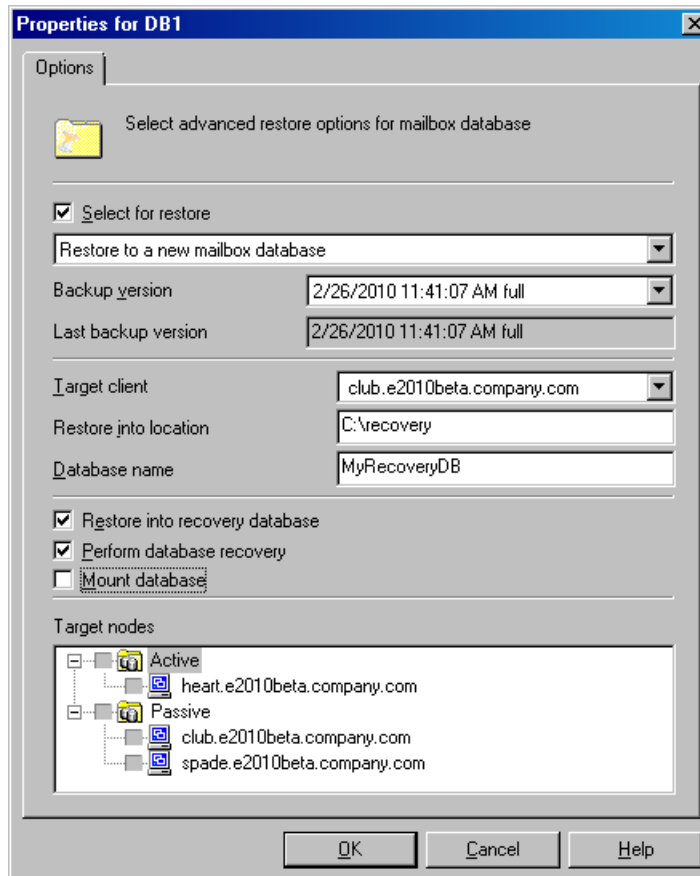
For details, see [“Restore to the latest state”](#) (page 218).

**Figure 105 Restore to a point in time**




For details, see [“Restore to a point in time”](#) (page 219).

**Figure 106 Restore to a recovery database**



For details, see [“Restore to a new mailbox database”](#) (page 219).

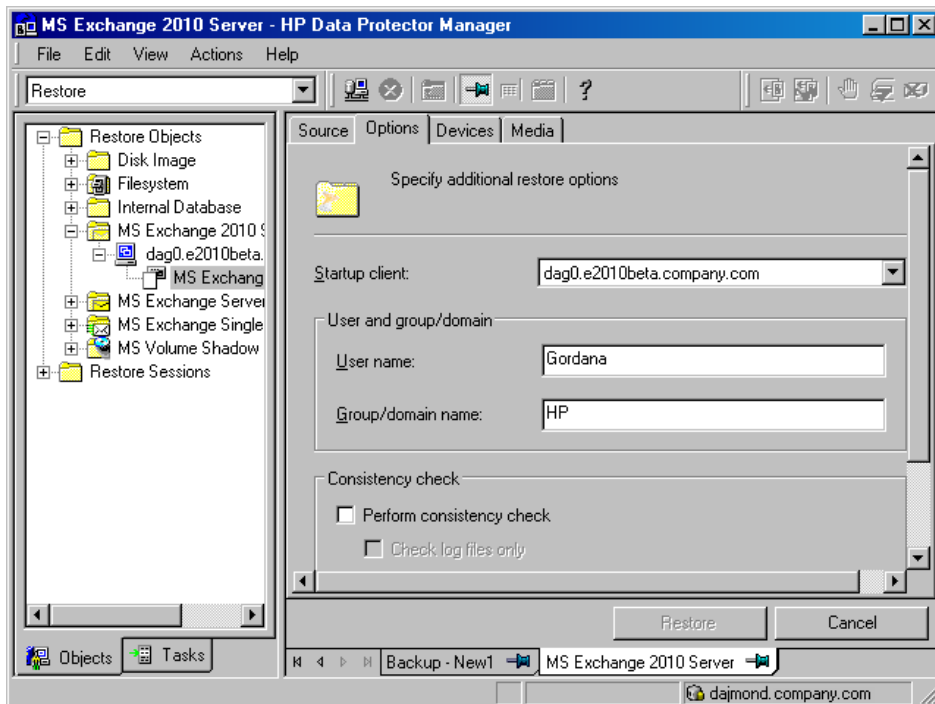
**Figure 107 Restore files to a temporary location**



For details, see “Restore files to a temporary location” (page 219).

4. In the **Options** page, specify the Data Protector Microsoft Exchange Server 2010 integration restore options. For details, see Table 31 (page 242).

**Figure 108 Restore options**



5. In the **Devices** page, select which devices to use for restore.  
For details on how to select devices to be used for restore, see the online Help index: “restore, selecting devices for”.
6. Click **Restore**.
7. In the **Start Restore Session** dialog box, click **Next**.
8. Specify **Report level** and **Network load**.  
Click **Finish** to start the restore.  
If the session succeeds, the message `Session completed successfully` is displayed at the end.

### Restoring using the Data Protector CLI

1. Log in to the Cell Manager or to any Microsoft Exchange Server client with the `User Interface` component installed under a user account that is configured as described in [“Configuring user accounts”](#) (page 204).
2. Go to the following directory:  
**Windows systems:** `Data_Protector_home\bin`  
**UNIX systems:** `opt/omni/sbin`

### 3. Run:

```
omnir -e2010
-barhost ClientName
[VSS_EXCHANGE_SPECIFIC_OPTIONS]
Database [Database ...]
[-user User:Domain]
[GENERAL_OPTIONS]

Database
{-db_name SourceDatabaseName | -db_guid SourceDatabaseGUID }
[-source SourceClientName]
{-repair | -latest | -pit | -new | -temp} E2010_METHOD_OPTIONS

E2010_REPAIR_METHOD_OPTIONS
[-no_resume_replication]

E2010_LATEST_METHOD_OPTIONS
[-node TargetNode ... | -all]
[-no_resume_replication]
[-no_recover]
[-no_mount]

E2010_PIT_METHOD_OPTIONS
-session BackupID
[-node TargetNode ... | -all]
[-no_resume_replication]
[-no_recover]
[-no_mount]

E2010_NEW_METHOD_OPTIONS
-session BackupID
-client TargetClientName
-location TargetDatabasePath
-name TargetDatabaseName
[-recoverydb]
[-no_recover]
[-no_mount]

E2010_TEMP_METHOD_OPTIONS
-session BackupID
-client TargetClientName
-location TargetDatabasePath
[-no_chain]
[-edb_only]
[-no_recover]
```

For a brief description of the options, see [“Restore options” \(page 239\)](#). For details, see the *omnir* man page or the *HP Data Protector Command Line Interface Reference*.

---

**NOTE:** A backup ID is a point in time. All objects (backup data) created in a backup session have the same backup ID, which is the same as the session ID of the backup session.

Mirrored objects and objects created in an object copy session have the same backup ID as objects created in the original backup session. Suppose the media set created in the original backup session no longer exists, but the media set created in an object copy session still exists. To restore the objects, you must specify the session ID of the original *backup* session (that is, the backup ID) and not the session ID of the *object copy* session.

The *omnir* syntax does not let you specify from which object copy to restore if several copies of the same object exist. This is only possible using the Data Protector GUI by setting the media allocation priority list.

---

### Example (Restore method – repair)

#### DAG environment

To restore all corrupt passive copies of the database DB1, which was backed up from a DAG whose virtual system name was `dag0.company.com`, and to ensure that the integration agent (`e2010_bar.exe`) is started on the client `exchange2.company.com`, run:

```
omnir -e2010 -barhost exchange2.company.com -db_name DB1 -source
dag0.company.com -repair
```

### Example (Restore method – latest)

#### Standalone environment

To restore the corrupt standalone database DB1, which resides on the client `exchange1.company.com`, to the latest possible point in time, and to ensure that the integration agent (`e2010_bar.exe`) is started on the client `exchange2.company.com`, run:

```
omnir -e2010 -barhost exchange2.company.com -db_name DB1 -source
exchange1.company.com -latest
```

#### DAG environment

Suppose you want to restore the active copy of the database DB1, which resides on the client `exchange1.company.com`, and the passive copies of the database that reside on the clients `exchange2.company.com` and `exchange3.company.com`. Suppose the database DB1 is part of a DAG whose virtual system name is `dag0.company.com`, and that you want the integration agent (`e2010_bar.exe`) to be started on the client `exchange2.company.com`. Run:

```
omnir -e2010 -barhost exchange2.company.com -db_name DB1 -source
dag0.company.com -latest -node exchange1.company.com -node
exchange2.company.com -node exchange3.company.com
```

### Example (Restore method – pit)

#### Standalone environment

Suppose you want to restore the corrupt standalone database DB1, which resides on the client `exchange1.company.com`, using the backup data created in the session `2010/5/14-1`. Suppose you want the integration agent (`e2010_bar.exe`) to be started on the client `exchange1.company.com`. Run:

```
omnir -e2010 -barhost exchange1.company.com -db_name DB1 -pit -session
2010/5/14-1
```

---

**NOTE:** Note that the `-source` option is not specified, in which case Data Protector assumes that the database was backed up from the client specified with the `-barhost` option.

---

### Example (Restore method – new)

#### DAG environment

Suppose you want to restore the backup of the database DB1 to a recovery database that should be created on the client `exchange2.company.com` and named `Recovery1`, with the files in the `C:\Recovery1Folder` directory. Suppose the database DB1 was backed up in the session `2010/5/14-1` from a DAG whose virtual system name was `dag0.company.com`. To also ensure that the integration agent (`e2010_bar.exe`) is started on the client `exchange1.company.com`, run:

```
omnir -e2010 -barhost exchange1.company.com -db_name DB1 -source
dag0.company.com -new -session 2010/5/14-1 -client exchange2.company.com
-location C:\Recovery1Folder -name Recovery1 -recoverydb
```

## Example (Restore method – temp)

### Standalone environment

Suppose you want to restore the transaction logs of the database DB1, which resides on the client exchange2.company.com. The logs were backed up in the Incremental backup session 2009/5/14-1. To restore the logs to the client exchange2.company.com to the directory C:\DB1TransactionLogFolder without performing database recovery, and to ensure that the integration agent (e2010\_bar.exe) is started on the client exchange1.company.com, run:

```
omnir -e2010 -barhost exchange1.company.com -db_name DB1 -source
exchange2.company.com -temp -session 2009/5/14-1 -client
exchange2.company.com -location C:\DB1TransactionLogFolder -no_chain
-no_recover
```

## Restoring using another device

You can restore using a device other than that used for a backup. For details, see the online Help index: “restore, selecting devices for”.

## Instant recovery

To be able to perform an instant recovery session, you need backup data that is stored on replica storage volumes. Such backup data is created in ZDB-to-disk and ZDB-to-disk+tape sessions.

You can restore multiple Microsoft Exchange Server databases in the same instant recovery session, specifying a different restore method for each database. For details, see [“Restore methods” \(page 218\)](#).

To start an instant recovery session, use the Data Protector GUI or CLI.

## Performing instant recovery using the Data Protector GUI

To perform an instant recovery session:

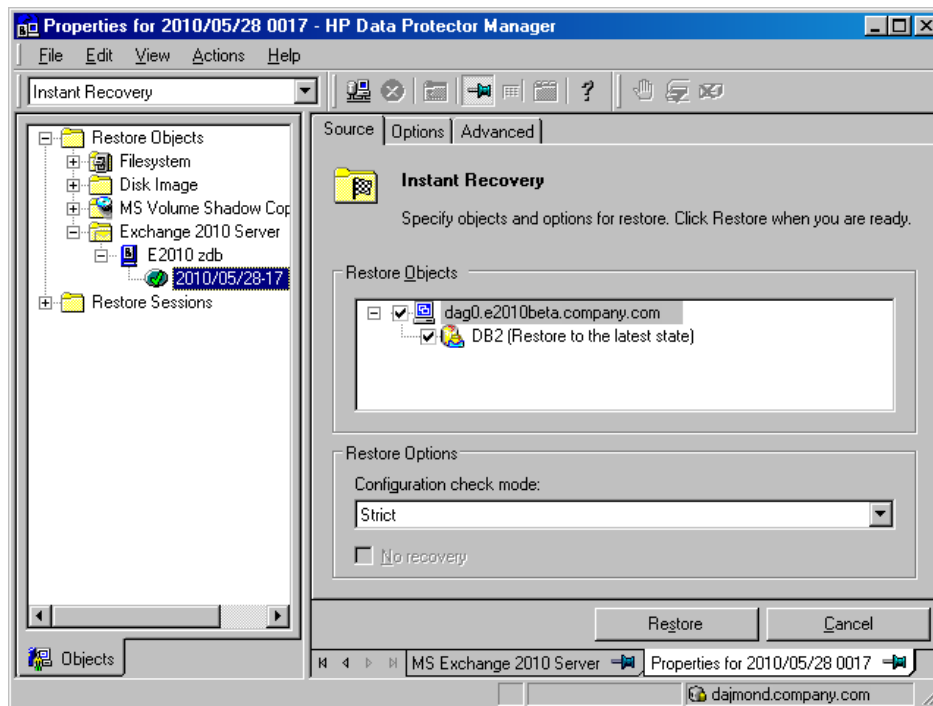
1. In the Context List, click **Instant Recovery**.
2. Expand **MS Exchange Server 2010** and select which ZDB-to-disk or ZDB-to-disk+tape session to use for instant recovery. The sessions are sorted according to the backup specifications used.
3. In the **Source** page, select which Microsoft Exchange Server databases to restore.

When you select a database, the Properties for Database dialog box is displayed automatically. Specify a restore method and click **OK**. For databases that are part of a DAG, the default restore method is **Repair all passive copies with failed status**. For standalone databases, the default is **Restore to the latest state**. To change the restore method, right-click the database and click **Properties**.

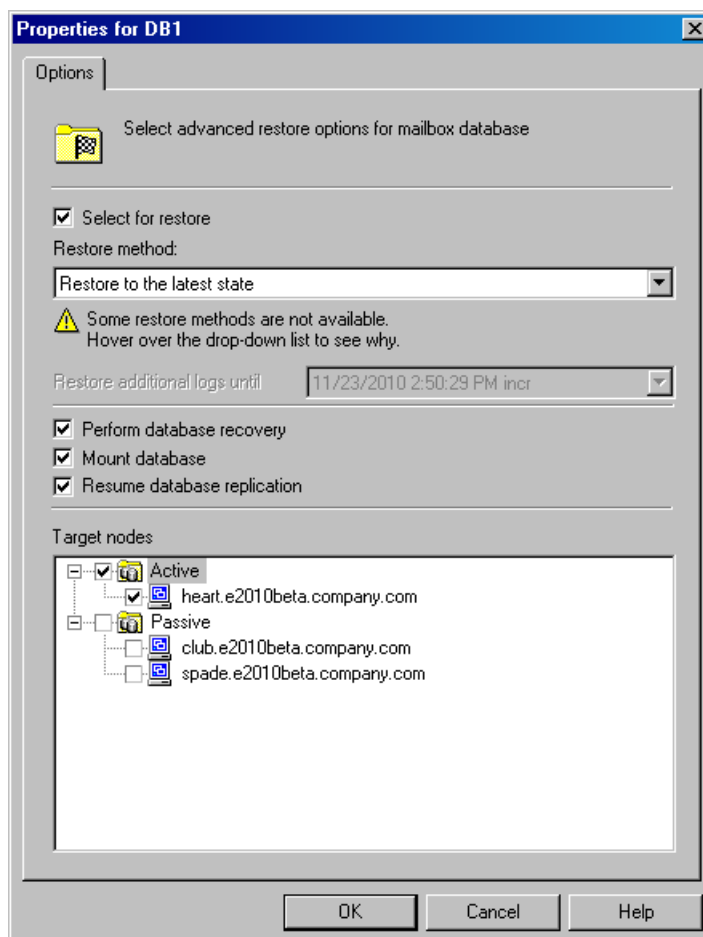
For details on **Configuration check mode**, press **F1**.



**Figure 109** Selecting databases for instant recovery

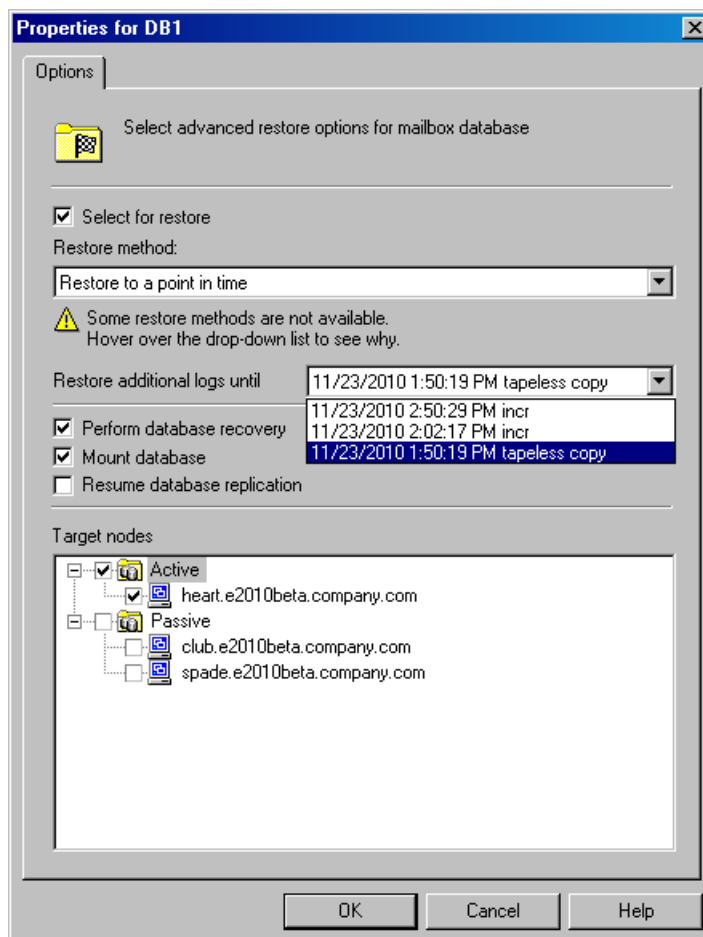


**Figure 110** Restore to the latest state



For details, see “[Restore to a point in time](#)” (page 240).

**Figure 111 Restore to a point in time**



For details, see “[Restore to a point in time](#)” (page 240).

**Figure 112 Restore to a new mailbox database**

Properties for DB2 (Restore to the latest state)

Options

Select advanced restore options for mailbox database

☒ Select for restore

Restore method:  
Restore to a new mailbox database

Restore additional logs until  
5/28/2010 4:12:07 PM tapeless

Target client  
club.e2010beta.company.com

Restore into location  
C:\RecoveryDBFiles

Database name  
Recovery1

☒ Restore into recovery database

☐ Perform database recovery

☐ Mount database

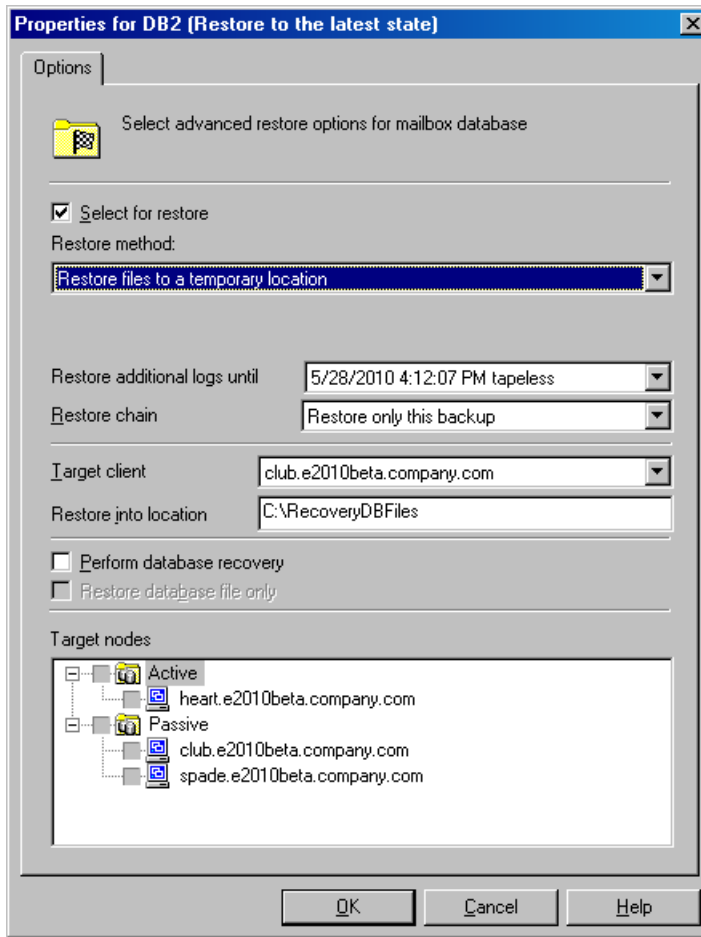
Target nodes

- Active
  - heart.e2010beta.company.com
- Passive
  - club.e2010beta.company.com
  - spade.e2010beta.company.com

OK Cancel Help

For details, see [“Restore to a new mailbox database”](#) (page 219).

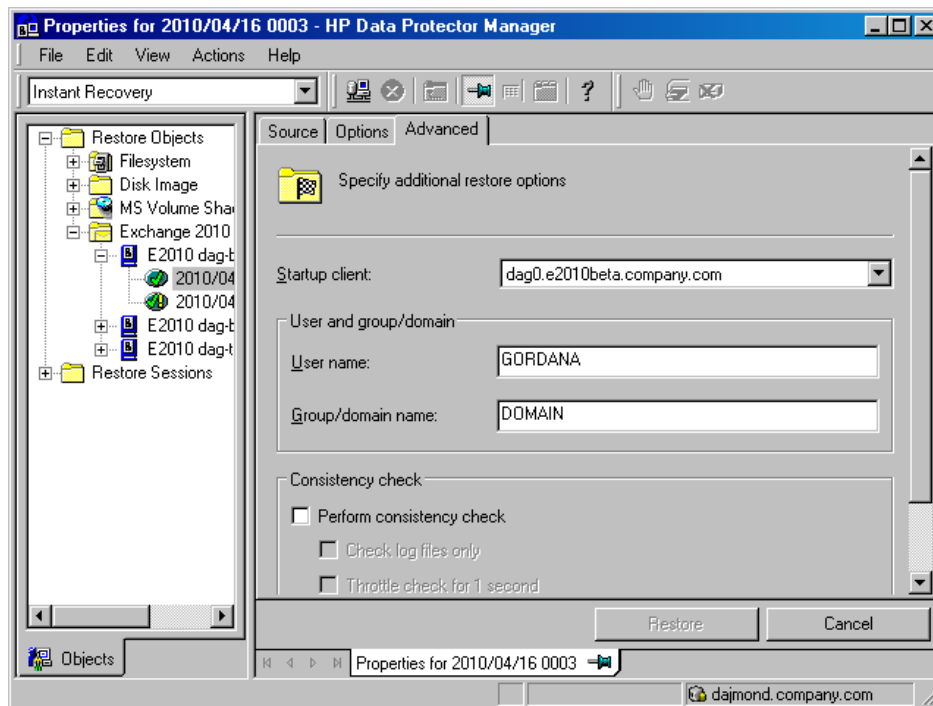
**Figure 113 Restore files to a temporary location**



For details, see [“Restore files to a temporary location”](#) (page 219).

4. In the **Options** page, specify ZDB-specific options. For details, press **F1**.
5. In the **Advanced** page, specify the Data Protector Microsoft Exchange Server 2010 integration instant recovery options. For details, see [“General restore options”](#) (page 242).

**Figure 114 Instant recovery – advanced**



6. Click **Restore**.

### Performing instant recovery using the Data Protector CLI

1. Log in to the Data Protector Cell Manager or to any Microsoft Exchange Server client under a user account as described in [“Configuring user accounts”](#) (page 204).

2. From the following directory:

**Windows:** *Data\_Protector\_home\bin*

**HP-UX:** */opt/omni/bin/*

run:

```
omnir -e2010
 -barhost ClientName
 -instant_restore
 [VSS_INSTANT_RECOVERY_OPTIONS]
 [VSS_EXCHANGE_SPECIFIC_OPTIONS]
 Database [Database ...]
 [-user User:Domain]
 [GENERAL_OPTIONS]
```

*Database*

```
{-db_name SourceDatabaseName | -db_guid SourceDatabaseGUID}
[-source SourceClientName]
{-repair | -latest | -pit | -new | -temp} E2010_METHOD_OPTIONS
```

*E2010\_REPAIR\_METHOD\_OPTIONS*  
[-no\_resume\_replication]

*E2010\_LATEST\_METHOD\_OPTIONS*  
[-node TargetNode ... | -all]  
[-no\_resume\_replication]  
[-no\_recover]  
[-no\_mount]  
[E2010\_IR\_SPECIFIC\_OPTIONS]

*E2010\_PIT\_METHOD\_OPTIONS*  
-session SessionID  
[-node TargetNode ... | -all]  
[-no\_resume\_replication]  
[-no\_recover]  
[-no\_mount]  
[E2010\_IR\_SPECIFIC\_OPTIONS]

*E2010\_NEW\_METHOD\_OPTIONS*  
-session SessionID  
-client TargetClientName  
-location TargetDatabasePath  
-name TargetDatabaseName  
[-recoverydb]  
[-no\_recover]  
[-no\_mount]  
[E2010\_IR\_SPECIFIC\_OPTIONS]

*E2010\_TEMP\_METHOD\_OPTIONS*  
-session SessionID  
-client TargetClientName  
-location TargetDatabasePath  
[-no\_chain]  
[-edb\_only]  
[-no\_recover]  
[E2010\_IR\_SPECIFIC\_OPTIONS]

*E2010\_IR\_SPECIFIC\_OPTIONS*  
[-from\_session SessionID]

For brief description of the options, see the section [“Restore options”](#) (page 239).

For details on the options, see the *HP Data Protector Command Line Interface Reference* or the omnir man page.

### Example (Restore method – latest)

#### Standalone environment

Suppose you want to restore the corrupt standalone database DB1, which resides on the client `exchange1.company.com`. To ensure the integration agent (`e2010_bar.exe`) is started on the client `exchange1.company.com`, and that the database is restored to the latest state, run:

```
omnir -e2010 -barhost exchange1.company.com -instant_restore -copy_back
-db_name DB1 -latest
```

### Example (Restore method – temp)

#### DAG environment

Suppose you want to restore the database DB1, which was part of a DAG whose DAG virtual system (host) name was `dag0.company.com`. The database was backed up in the session `2010/5/14-1`. To restore the database to a temporary location on the client `exchange1.company.com` to the directory `C:\BackupDatabase`, and to ensure that the integration agent (`e2010_bar.exe`) is started on the client `exchange1.company.com`, run:

```
omnir -e2010 -barhost exchange1.company.com -instant_restore -copy_back
-db_name DB1 -source dag0.company.com -temp -session 2010/5/14-1 -client
exchange1.company.com -location C:\BackupDatabase
```

## Restore options

**Table 26 Repair all passive copies with failed status**

| Option GUI / CLI                                                            | Description                                                                                                                                                                                                                                                                   |
|-----------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Resume database replication /</b><br><code>-no_resume_replication</code> | Available in DAG environments. Resumes the replication between the active and passive copies after the copies are restored.<br><br>Note that the CLI option <code>-no_resume_replication</code> has the opposite meaning. If it is specified, the replication is not resumed. |
| <b>Restore additional logs until</b><br><code>-session</code>               | This is an instant recovery-specific option.<br>Not available.                                                                                                                                                                                                                |
| <b>Target nodes</b>                                                         | Not available.<br><br>The clients (that is, copies) that have the status <code>Failed</code> or <code>FailedAndSuspended</code> are automatically selected.                                                                                                                   |

**Table 27 Restore to the latest state**

| Option GUI / CLI                                               | Description                                                                                                                                                                                                                                                                                                                                                                              |
|----------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Select for restore</b>                                      | Specifies whether the database should be restored.                                                                                                                                                                                                                                                                                                                                       |
| <b>Restore additional logs until</b><br><code>-session</code>  | This is an instant recovery-specific option.<br>Not available.                                                                                                                                                                                                                                                                                                                           |
| <b>Perform database recovery /</b><br><code>-no_recover</code> | Available when restoring a standalone database (standalone environment) or an active copy (DAG environment). Applies the logs to the database file ( <code>.edb</code> ) after the restore completes.<br><br>Note that the CLI option <code>-no_recover</code> has the opposite meaning. If it is specified, the database recovery is not performed.                                     |
| <b>Mount database /</b><br><code>-no_mount</code>              | Available when restoring a standalone database (standalone environment) or an active copy (DAG environment). Mounts the database after the database recovery completes. This option is available only if <b>Perform database recovery</b> is selected.<br><br>Note that the CLI option <code>-no_mount</code> has the opposite meaning. If it is specified, the database is not mounted. |

**Table 27 Restore to the latest state** *(continued)*

| Option GUI / CLI                                               | Description                                                                                                                                                                                                                                                                                  |
|----------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Resume database replication</b> /<br>-no_resume_replication | Available when restoring passive copies (DAG environment). Resumes the replication between the active and passive copies after the copies are restored.<br><br>Note that the CLI option -no_resume_replication has the opposite meaning. If it is specified, the replication is not resumed. |
| <b>Target nodes</b><br>-node   -all                            | Available only in DAG environments. Specifies which clients (that is, database copies) to restore.                                                                                                                                                                                           |

**Table 28 Restore to a point in time**

| Option GUI/CLI                                                 | Description                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Select for restore</b>                                      | See the description in “ <a href="#">Restore to the latest state</a> ” (page 239).                                                                                                                                                                                                                                                                                                                                                           |
| <b>Backup version</b> /<br>-session                            | This is a standard restore-specific option.<br>It specifies from which backup data to restore. Select a backup ID.<br>If a Differential backup session is selected, the .log files backed up in the selected Differential backup session are restored.<br>If an Incremental backup session is selected, the .log files backed up in all subsequent Incremental backup sessions, up to the selected Incremental backup session, are restored. |
| <b>Last backup version</b>                                     | This is a standard restore-specific option.<br>It shows the session in which the database was last backed up.                                                                                                                                                                                                                                                                                                                                |
| <b>Restore additional logs until</b><br>-session               | This is an instant recovery-specific option.<br>If a Differential backup session is selected, the .log files backed up in the selected Differential backup session are restored.<br>If an Incremental backup session is selected, the .log files backed up in all subsequent Incremental backup sessions, up to the selected Incremental backup session, are restored.                                                                       |
| <b>Perform database recovery</b> /<br>-no_recover              | See the description in “ <a href="#">Restore to the latest state</a> ” (page 239).                                                                                                                                                                                                                                                                                                                                                           |
| <b>Mount database</b> /<br>-no_mount                           |                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Resume database replication</b> /<br>-no_resume_replication |                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Target nodes</b> /<br>-node   -all                          | See the description in “ <a href="#">Restore to the latest state</a> ” (page 239). The node (client) hosting the active copy is automatically selected for restore.                                                                                                                                                                                                                                                                          |

**Table 29 Restore to a new mailbox database**

| Option GUI/CLI                                   | Description                                                                                                 |
|--------------------------------------------------|-------------------------------------------------------------------------------------------------------------|
| <b>Select for restore</b>                        | See the description in “ <a href="#">Restore to the latest state</a> ” (page 239).                          |
| <b>Restore additional logs until</b><br>-session | This is an instant recovery-specific option.<br>See the description in <a href="#">Table 28</a> (page 240). |
| <b>Target client</b> /<br>-client                | Specifies the client to restore to.                                                                         |



**Table 29 Restore to a new mailbox database** *(continued)*

| Option GUI/CLI                                         | Description                                                                                                                                                                                                           |
|--------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Restore into location</b> /<br>-location            | Specifies the directory to restore to (standard restore) or the directory to mount the replica storage volumes to (instant recovery).                                                                                 |
| <b>Database name</b> /<br>-name                        | Specifies the name to be used for the new database. If another database with the same name already exists, the restore fails.                                                                                         |
| <b>Restore into Recovery database</b> /<br>-recoverydb | Restores the data to a Microsoft Exchange Server recovery database. Although multiple recovery databases can exist in parallel, only one recovery database can be mounted to the Microsoft Exchange Server at a time. |
| <b>Backup version</b> /<br>-session                    | See the description in <a href="#">Table 28 (page 240)</a> .                                                                                                                                                          |
| <b>Last backup version</b>                             |                                                                                                                                                                                                                       |
| <b>Perform database recovery</b> /<br>-no_recover      | See the description in <a href="#">“Restore to the latest state” (page 239)</a> .                                                                                                                                     |
| <b>Mount database</b> /<br>-no_mount                   |                                                                                                                                                                                                                       |
| <b>Target nodes</b>                                    | Not available.                                                                                                                                                                                                        |

**Table 30 Restore files to a temporary location**

| Option GUI/CLI                                    | Description                                                                                                                                                                                                                        |
|---------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Select for restore</b>                         | See the description in <a href="#">“Restore to the latest state” (page 239)</a> .                                                                                                                                                  |
| <b>Restore additional logs until</b><br>-session  | This is an instant recovery-specific option.<br>See the description in <a href="#">Table 28 (page 240)</a> .                                                                                                                       |
| <b>Restore chain</b>                              | If this option is set to <b>Restore only this backup</b> , only files backed up in the selected session are restored.<br>If this option is set to <b>Full restore (full, incr, diff backups)</b> , the complete chain is restored. |
| <b>Target client</b> /<br>-client                 | See the description in <a href="#">“Restore to a new mailbox database” (page 240)</a> .                                                                                                                                            |
| <b>Restore into location</b> /<br>-location       |                                                                                                                                                                                                                                    |
| <b>Backup version</b> /<br>-session               | See the description in <a href="#">“Restore to a point in time” (page 240)</a> .                                                                                                                                                   |
| <b>Last backup version</b>                        |                                                                                                                                                                                                                                    |
| <b>Restore database files only</b> /<br>-edb_only | Restores only the database files (.edb). The logs (.log) and checkpoint files (.chk) are not restored.                                                                                                                             |
| <b>Perform database recovery</b> /<br>-no_recover | See the description in <a href="#">“Restore to the latest state” (page 239)</a> .                                                                                                                                                  |
| <b>Target nodes</b>                               | Not available.                                                                                                                                                                                                                     |

**Table 31 General restore options**

| Option GUI / CLI                                                                                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Startup client /</b><br>-barhost                                                                | <p>Specifies the client on which the integration agent (e2010_bar.exe) should be started. If the DAG virtual client (host) is selected, the integration agent is started on the currently active node. To find out which node is currently active, see <a href="#">“Tip” (page 242)</a>.</p> <p>Default: The same client that was specified for the backup session. If the DAG virtual client was specified, this client is now selected. However, note that the integration agent may not be started on the same physical node as during the backup session; it depends which node is currently active.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Username</b><br><b>Group/Domain name /</b><br>-user                                             | <p>Specifies which Windows domain user account to use for the restore session. Ensure that the user is configured as described in <a href="#">“Configuring user accounts” (page 204)</a>.</p> <p>If these options are not specified, the restore session is started under the user account under which the Data Protector Inet service is running.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Perform consistency check /</b><br>[-exch_check<br>[-exch_throttle Value]  <br>-exch_checklogs] | <p>If this option is selected, Microsoft Exchange Server checks the consistency of a database's backup data. If this option is not selected, the session finishes earlier, but the backup data consistency is not guaranteed.</p> <p>The check is performed on the source storage volumes after the backup data is restored. You do not need to perform the consistency check if it was already performed at the time of backup.</p> <p>Default: not selected</p> <p>If the <b>Check log files only</b> option is selected, only the log file backup data is checked, which is enough for Microsoft Exchange Server to guarantee data consistency.</p> <p>Default: not selected</p> <p>By default, the consistency check is I/O intensive, which can negatively affect disk performance. The <b>Throttle check for 1 second</b> option throttles down the consistency check of the database file .edb to lessen impact on the disk performance. Specify after how many input/output operations the check should stop for one second.</p> <p>This option is not available if only the log files are checked.</p> <p>Default: not selected</p> |



**TIP:** To find out which Microsoft Exchange Server node is currently active, connect to one of the nodes and run:

```
cluster group
```

### Example

```
C:\Users\administrator.E2010BETA>cluster group
Listing status for all available resource groups:
```

| Group             | Node  | Status  |
|-------------------|-------|---------|
| Available Storage | spade | Offline |
| Cluster Group     | club  | Online  |

The currently active node has the status Online. In the example, this is club.

## Monitoring sessions

You can monitor currently running sessions in the Data Protector GUI. When you run a backup or restore session, a monitor window shows the progress of the session. Closing the GUI does not affect the session.

You can also monitor sessions from any Data Protector client with the `User Interface` component installed, using the Monitor context.

To monitor a session, see the online Help index: “viewing currently running sessions”.

## Troubleshooting

This section lists general checks and verifications, plus problems you might encounter when using the Data Protector Microsoft Exchange Server 2010 integration.

Because the Data Protector Microsoft Exchange Server 2010 integration is based on the Data Protector Microsoft Volume Shadow Copy Service integration, also see troubleshooting information in the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

For general Data Protector troubleshooting information, see the *HP Data Protector Troubleshooting Guide*.

### Before you begin

- Ensure that the latest official Data Protector patches are installed. On how to verify this, see the online Help index: “patches”.
- See the *HP Data Protector Product Announcements, Software Notes, and References* for general Data Protector limitations, as well as recognized issues and workarounds.
- See <http://www.hp.com/support/manuals> for an up-to-date list of supported versions, platforms, and other information.

### Checks and verifications

If your browsing, backup, or restore failed:

- Examine system errors reported in the `debug.log` located in:  
`Data_Protector_home\log`.
- Check if you can do a filesystem backup and restore on the problematic client. For information, see the online Help.

### Problems

#### Problem

#### **It takes a long time to display Microsoft Exchange Server topology in the Data Protector GUI**

When you open the Data Protector GUI and try to display the source page, either in the Backup or Restore context, you must wait a long time.

This may happen if there is an unresponsive system in the same domain (for example, a system that is shut down). The problem occurs even if the unresponsive system is not part of your backup environment. This is due to Microsoft Exchange Server problems with execution of Microsoft Exchange Server Shell commands.

#### Action

Remove the system from the domain or fix the problem.

#### Problem

#### **A database backup cannot be performed**

When you start a backup session for a database, the database is not backed up, appearing to be locked by other session, though there are no other backup sessions currently running. A message similar to the following is displayed:

```
[Minor] From: OB2BAR_E2010_BAR@exch03.e2010.company.com "MS Exchange
2010 Server" Time: 6/17/2010 3:07:13 PM
```

[170:313] One or more copies of database DEMAR are already being backed up in a different session.

This may happen if the integration agent (e2010\_bar.exe) was terminated by force while a previous backup session was in progress, either because the Microsoft Exchange Server system was restarted or for some other reason, so the lock remains.

#### Action

From the *Data\_Protector\_home\bin* directory, run:

```
omnidbutil -free_cell_resources
```

---

**NOTE:** This command line removes all existing locks, so ensure that none of the existing locks is still needed.

---

#### Problem

##### Restore fails

When you try to restore a database, the session fails.

This may happen if a database has been restored before (probably unsuccessfully), and during that previous restore session, the Microsoft Exchange Server created an `.env` file in the database directory. This file now prevents the database from being restored again.

#### Action

Delete the `.env` file and start a new restore session.

#### Problem

##### Restore to the latest state fails

When you try to restore a database all of whose log files were lost, using the restore method **Restore to the latest state** with the **Perform database recovery** option selected, the database recovery fails.

This may happen if a database is restored from a Full backup (that is, the restore chain consists of only the Full backup session). Since in the **Restore to the latest state** session, only the `.edb` file is restored from the Full backup (see [“Restore chain” \(page 221\)](#)), when the database recovery is started, there are no logs to be applied to the database file, and the database recovery fails.

#### Action

Restore the database using the restore method **Restore to a point in time**. For details, see [“Restore to a point in time” \(page 219\)](#).

#### Problem

##### After instant recovery to a point in time, passive copies remain in the Failed state

When you start an instant recovery session in a DAG environment to restore active and passive copies of the same database, using the restore method **Restore to a point in time**, the data is successfully restored, but the synchronization between the active copy and passive copies fails, leaving the passive copies in the `Failed` state.

This problem occurs if, after the data is restored, the passive copy has extra log files, which are not present at the active copy side, and so synchronization cannot be established. This can happen if, during a Full backup session, multiple copies of a database are selected for backup. Data Protector first performs a Full backup of the passive copy that has the fewest logs applied to the database file, and then performs a Copy backup of all the remaining copies, with the active copy being backed up last. While the backup session is in progress, a new log may be created at the active copy side, so when the active copy is backed up, the newly created log is also backed up. If, further on in time, a failover occurs (one of the passive copies becomes the active copy) and you perform a **Restore to a point in time** instant recovery, each copy is restored from its own replica

storage volumes. This results in the active copy (which was passive at the time of backup) having fewer logs than the passive copy (which was active at the time of backup). Consequently, synchronization cannot be established.

#### Action

Perform a full reseed for all Failed passive copies.

#### Problem

##### **In a DAG, a copy-back instant recovery fails when restoring a non-original database copy**

Suppose you backed up a database copy by creating a snapclone replica (HP P6000 EVA Disk Array Family). Using the copy-back method, this replica can be used to restore the original database copy and/or the related database copies that reside on different Microsoft Exchange Server systems in the DAG. If the size of the source storage volumes on those Microsoft Exchange Server systems differs from the size of the source storage volumes that were backed up, the instant recovery session fails.

If the **Retain source for forensics** option is selected, a message similar to the following is displayed:

```
[Warning] From: SMISA@dizzy.e2008.company.com "SMISA" Time:
7/6/2010 2:51:08 PM
[236:8001] This pair of source and target storage volumes are
not the same size.
Source storage volume : 50014380025B4860\\Virtual Disks\VSS\
FizzyDizzy\DAG\dizzy\dizzy-DB1-data\ACTIVE
Source size : 4 GB
Target storage volume : 50014380025B4860\\Virtual Disks\VSS\
FizzyDizzy\DAG\fizzy\hpVSS-LUN-06Jul10 02.23.27\ACTIVE
Target size : 3 GB
```

```
[Major] From: OB2BAR_VSSBAR@dizzy.e2008.company.com "MS Exchange
2010 Server" Time: 7/6/2010 3:11:16 PM
The system failed to refresh symbolic links in kernel object namespace.
```

If the **Retain source for forensics** option is not selected, a message similar to the following is displayed:

```
[Major] From: SMISA@dizzy.e2008.company.com "SMISA" Time:
7/6/2010 2:51:08 PM
[236:8001] This pair of source and target storage volumes are not
the same size.
Source storage volume : 50014380025B4860\\Virtual Disks\VSS\
FizzyDizzy\DAG\dizzy\dizzy-DB1-data\ACTIVE
Source size : 4 GB
Target storage volume : 50014380025B4860\\Virtual Disks\VSS\
FizzyDizzy\DAG\fizzy\hpVSS-LUN-06Jul10 02.23.27\ACTIVE
Target size : 3 GB
```

#### Action

Ensure that storage volumes on different Microsoft Exchange Server systems are the same size.

---

## 6 Data Protector Microsoft SharePoint Server 2007/2010 VSS based solution

### Introduction

This chapter explains how to configure and use the Data Protector Microsoft SharePoint Server 2007/2010 VSS based solution (**VSS based solution**). In reality, the solution is based on the Data Protector Microsoft Volume Shadow Copy Service integration (**VSS integration**). For details on the VSS integration, see the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

The chapter describes concepts and methods you need to understand to back up and restore Microsoft Office SharePoint Server 2007 and Microsoft SharePoint Server 2010 data that is stored in Microsoft SQL Server databases. For example:

- The configuration database (SharePoint\_Config)
- Content databases (SharePoint\_AdminContent\_Label, WSS\_Content\_Label,...)
- Shared Services Provider databases (SSP\_DB) (Microsoft Office SharePoint Server 2007)
- SharePoint Service Applications databases (SSA\_DB) (Microsoft SharePoint Server 2010)
- Search databases (SSP\_Search\_DB)
- The Single Sign-On database (SSO)

In addition, you can also back up and restore Microsoft SharePoint Server search index files.

From now on, both Microsoft SharePoint Server versions are called **Microsoft SharePoint Server**, unless the differences are pointed out.

### Backup

Microsoft SharePoint Server data that is stored in Microsoft SQL Server databases is backed up using one of the following Microsoft SQL Server VSS writers:

- MSDE writer (for Microsoft SQL Server 2000 databases)
- SqlServerWriter (for Microsoft SQL Server 2005/2008 databases)

Microsoft Office SharePoint Server 2007 search index files are backed up using the following VSS writers:

- OSearch VSS writer
- SPSearch VSS writer

Microsoft SharePoint Server 2010 search index files are backed up using the following VSS writers:

- OSearch14 VSS writer
- SPSearch4 VSS writer

Microsoft FAST Search Server 2010 search index files are backed up:

- using the Data Protector Disk Agent (in case of the standard filesystem backup with VSS enabled)
- using the Data Protector VSS integration (in case of the ZDB filesystem backup)

You can create and run backup specifications using the Data Protector PowerShell command which is described in [“Backup” \(page 249\)](#).

## Limitations

- The only supported way to start backup sessions is using the Data Protector PowerShell command. Starting the backup sessions using the Data Protector GUI or CLI is not supported.
- With VSS based solution, the FAST Search index files can also be backed up incrementally when using the Data Protector Disk Agent. For all other Microsoft SharePoint Server data only Full backup type is supported.

## Restore

Restore can be started using the Data Protector GUI or CLI as described in [“Restore” \(page 261\)](#).

## Installation and configuration

### ZDB prerequisites

If you plan to run ZDB and instant recovery (IR) sessions, ensure that the SPSearch and OSearch index files of each SSP or SSA, and the FAST Search index files reside on a disk array.

### Microsoft Office SharePoint Server 2007

The default location for the SPSearch index files is:

```
C:\Program Files\Microsoft Office Servers\12.0\Data\Applications
```

The default location for the OSearch index files is:

```
C:\Program Files\Microsoft Office Servers\12.0\Data\Office Server\
Applications
```

To move the index files to the disk array:

1. Open the Command Prompt and change the directory to:  

```
C:\Program Files\Common Files\Microsoft Shared\Web Server
Extensions\12\BIN>
```
2. To move the SPSearch index files, run:  

```
stsadm -o spsearch -indexlocation PathToNewLocation
```
3. To move the OSearch index files, run:  

```
stsadm -o editssp -title SSPname -indexlocation PathToNewLocation
```

### Microsoft SharePoint Server 2010

The default location for the SPSearch index files is:

```
C:\Program Files\Microsoft Office Servers\14.0\Data\Applications
```

The default location for the OSearch index files is:

```
C:\Program Files\Microsoft Office Servers\14.0\Data\Office Server\
Applications
```

To move the index files to the disk array:

1. Open the Command Prompt and change the directory to:  

```
C:\Program Files\Common Files\Microsoft Shared\Web Server
Extensions\14\BIN>
```
2. To move the SPSearch index files, run:  

```
stsadm -o spsearch -indexlocation PathToNewLocation
```
3. To move the OSearch index files use the Central Administration (modify farm topology).

The FASTSearch home folder must be installed to the disk array during the FAST Search Server 2010 system installation. In case of a multiple FAST Search Server system farm, ensure that the FASTSearch home folders on all the systems have the same path (drive and path name).

## Licensing

The Data Protector VSS based solution requires one online-extension license per each Microsoft SharePoint Server client participating in the backup and restore process. This means one online-extension license for each system on which the Data Protector MS Volume Shadow Copy Integration component is installed.

## Installing the integration

For details of how to install a Data Protector cell, see the *HP Data Protector Installation and Licensing Guide*.

To be able to back up Microsoft SharePoint Server objects, install the following Data Protector components and files:

- Service Pack 2 (Windows SharePoint Services 3.0 and Microsoft Office SharePoint Server 2007)
- Windows PowerShell 2.0 and User Interface component on the Microsoft SharePoint Server system on which you plan to execute the Data Protector commands and on which you install the Microsoft Volume Shadow Copy Integration. See the next bullet.

If not already available on your Windows system, you can download Windows PowerShell from: <http://www.microsoft.com/windowsserver2003/technologies/management/powershell/default.aspx>

- MS Volume Shadow Copy Integration - on Microsoft SQL Server system and the Microsoft SharePoint Server systems that have at least one of the following services enabled:

### **Microsoft Office SharePoint Server 2007**

- Windows SharePoint Services Database
- Windows SharePoint Services Help Search
- Office SharePoint Server Search

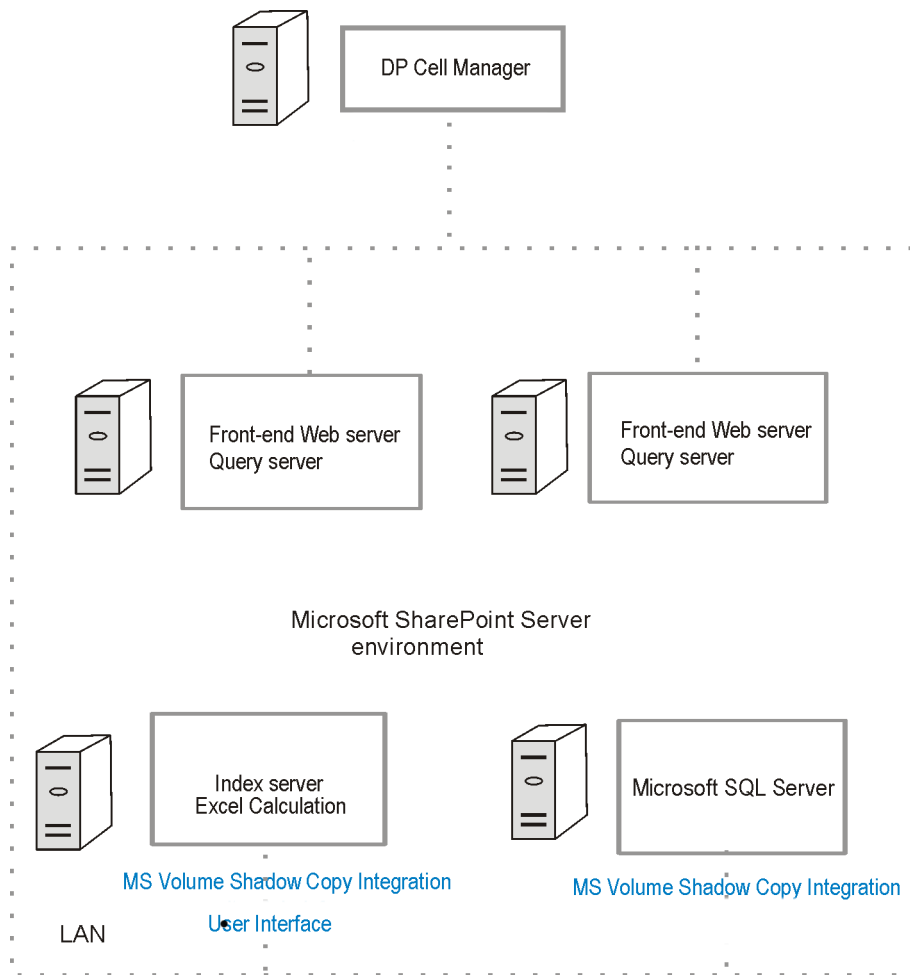
### **Microsoft SharePoint Server 2010**

- SharePoint Foundation Database
  - SharePoint Foundation Help Search
  - SharePoint Server Search
  - FAST Search Server 2010 for SharePoint
- The Data Protector Disk Agent component on each Microsoft FAST Search Server 2010 system for SharePoint (Microsoft SharePoint Server 2010, in case of the mixed filesystem + ZDB environment)

Ensure that the Volume Shadow Copy service is started on all these clients.



**Figure 115 Installing a medium farm (example)**



In Figure 115 (page 249), the Data Protector components that you need to install are colored blue.

## Configuring the integration

For details of how to configure the Data Protector VSS integration, see the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

## Configuring user accounts

Create or identify a Windows domain user account that has Windows administrative rights on the Microsoft SharePoint Server system on which you plan to execute the Data Protector commands. This user must also be granted Microsoft SharePoint Server administrative rights and must be added to the Data Protector `admin` user group.

## Backup

To back up Microsoft SharePoint Server data, create backup specifications and start backup sessions using the Data Protector PowerShell command `SharePoint_VSS_backup.ps1`.

## Prerequisites

- The Windows Remote Management service (which is used for starting and stopping Windows services remotely, and suspending and resuming FAST for Microsoft SharePoint Server 2010) must be configured on all systems.

To configure and analyze the WinRM service, run the `winrm quickconfig` command.

For more information, see the Windows Remote Management service documentation.

- In case of Microsoft SharePoint Server 2010 which uses Microsoft SQL Server 2008 for storing data, and Remote BLOB Storage (RBS) is used with the FILESTREAM provider, ensure that FILESTREAM access level is set to Full access enabled or Transact-SQL access enabled.

For details of how to configure RBS and FILESTREAM, see the Microsoft SQL Server 2008 documentation.

## Limitations

- The only supported way to start backup sessions is using the Data Protector PowerShell command. Starting the backup sessions using the Data Protector GUI or CLI is not supported.
- With VSS based solution, the FAST Search index files can also be backed up incrementally when using the Data Protector Disk Agent. For all other Microsoft SharePoint Server data only Full backup type is supported.

## Recommendations

- Use the Data Protector PowerShell command to create backup specifications and not the Data Protector GUI.
- Use the Data Protector GUI to modify backup specifications (for example, to add backup devices).
- Use the simple mode for the SQL Server databases. In case you want to use the full mode anyway, ensure that you truncate the transaction logs. Otherwise, you may run out of disk space.
- Whenever you change the farm configuration, perform a new backup.
- In case you want to back up the Single Sign-On database, do not forget to back up the encryption key as described in: <http://technet.microsoft.com/en-us/library/cc262932.aspx#Section32>.

Otherwise, you will not be able to restore the database.

## How the command works

When you execute the Data Protector PowerShell command `SharePoint_VSS_backup.ps1`, Data Protector first queries for information about the Microsoft SharePoint Server environment. Then it creates backup specifications.

The newly created backup specifications are named `SharePoint_VSS_backup_ClientName` and have the same backup device specified for use (the one that you specified at command runtime).

Once the backup specifications are created, the command starts backup sessions (one session for each backup specification).

## Microsoft Office SharePoint Server 2007

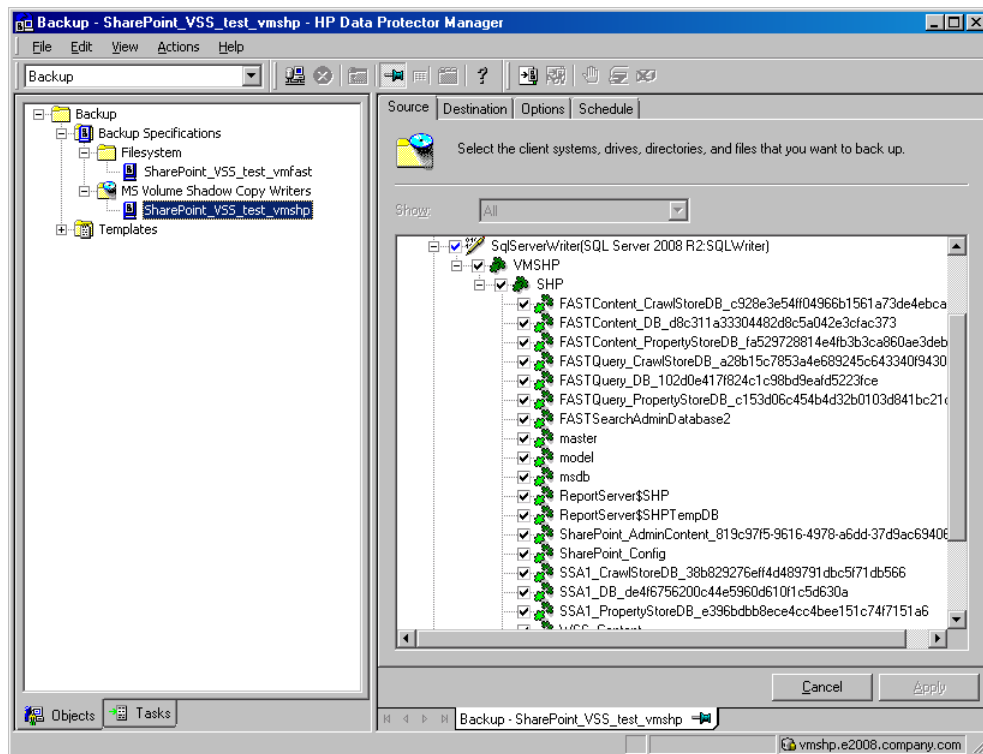
In a Microsoft Office SharePoint Server 2007 environment, the command creates a separate backup specification for each Microsoft Office SharePoint Server 2007 system that has at least one of the following services enabled:

- Windows SharePoint Services Database
- Windows SharePoint Services Help Search
- Office SharePoint Server Search

For a system with the Windows SharePoint Services Database service enabled, the command creates a backup specification that has the `SqlServerWriter` (Microsoft SQL Server

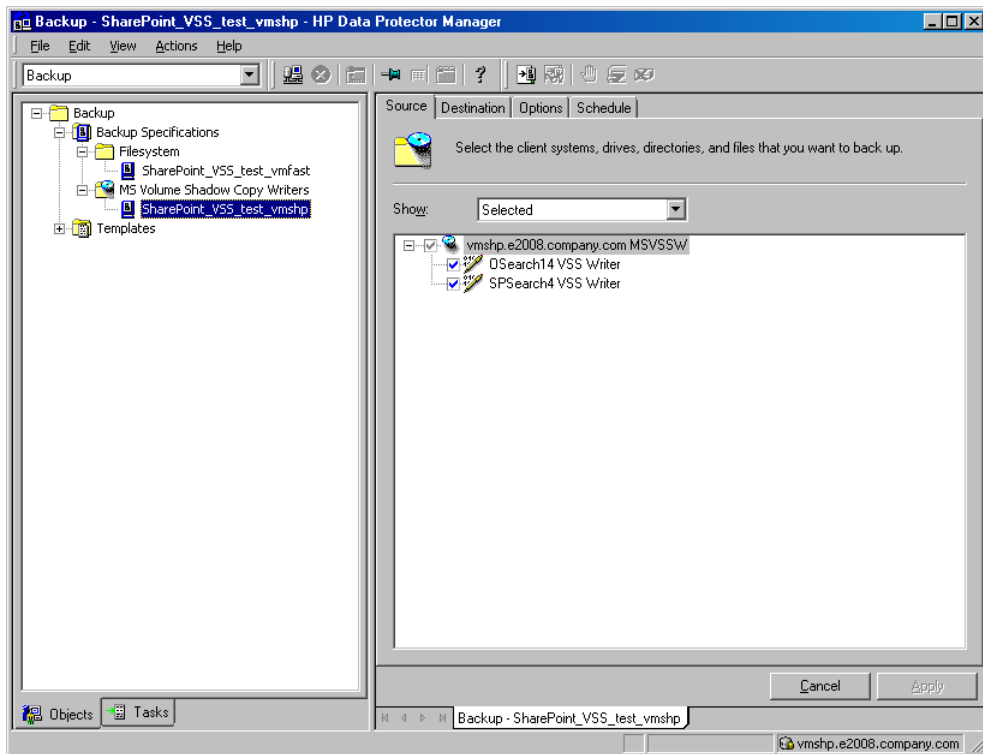
2005/2008) or MSDE writer (Microsoft SQL Server 2000) object selected (Figure 116 (page 251)).

**Figure 116 Selection of Microsoft Office SharePoint Server 2007 databases**



For a system with the Windows SharePoint Services Help Search and Office SharePoint Server Search services enabled, the command creates a backup specification that has the SPSearch VSS Writer and OSearch VSS Writer objects selected (Figure 117 (page 252)).

**Figure 117 Selecting Microsoft Office SharePoint Server 2007 search index files**



## Microsoft SharePoint Server 2010

In a Microsoft SharePoint Server 2010 environment, the command creates a separate backup specification for each Microsoft SharePoint Server 2010 system that has at least one of the following services enabled:

- SharePoint Foundation Database
- SharePoint Foundation Help Search
- SharePoint Server Search 14
- FAST Search Server 2010 for SharePoint (FAST Search)

For a system with the SharePoint Foundation Database service enabled, the command creates a backup specification that has the `SqlServerWriter` (Microsoft SQL Server 2005/2008) object selected.

For a system with the SharePoint Foundation Help Search and SharePoint Server Search services enabled, the command creates a backup specification that has the `SPSearch4 VSS Writer` and `OSearch14 VSS Writer` objects selected.

For a system with the FAST Search Server 2010 service enabled, the command with the `-hardware` option specified creates a VSS backup specification that has the complete `FASTSearch` home folder (including `bin` and `lib`) selected.

## Considerations

- Microsoft Office SharePoint Server 2007 only: If the Office SharePoint Server Search service is enabled on two separate Microsoft SharePoint Server systems so that one is assigned the Query and the other the Indexing role, the command creates a backup specification only for the system with the Indexing role. It is not created for the one with the Query role. To restore index files on the Query system, copy the files from the Indexing system to the

Query system after the restore. For details, see the section [“Restoring index files on the Query system” \(page 266\)](#)”.

- The command options enable you to split the process into two parts: first you create the backup specifications and then you start backup sessions. In this way, you can manually modify the newly-created backup specifications in the Data Protector GUI before the backup is actually started.
- If Microsoft SQL Server instances are used not only by Microsoft SharePoint Server but also by other database applications, modify the backup specifications so that only the databases that belong to Microsoft SharePoint Server are selected for backup. See the section [“Modifying backup specifications” \(page 257\)](#).
- If you have Microsoft SQL Server database mirroring enabled, a failover can occur and so a different Microsoft SQL Server system becomes active. Since the command creates backup specifications only for the currently active Microsoft SQL Server systems, it is advisable to update (recreate) the backup specifications before the backup is started.

## The command syntax

```
SharePoint_VSS_backup.ps1 -help | -version
SharePoint_VSS_backup.ps1 -createonly CreateOptions
SharePoint_VSS_backup.ps1 -backuponly BackupOptions
SharePoint_VSS_backup.ps1 -resumefarm [-preview] | -resumecert

CreateOptions
{-device DevName | -hardware {no_keep|keep|ir} [-device DevName] }
[-overwrite]
[-prefix PrefixName]
[-excludeindex]

BackupOptions
[-outfile PathToFile]
[-prefix PrefixName]
[-preview]
[-snapshot {diskonly | disktape | tapeonly}]
[-reduce]
[-mode {full | incremental | incremental1 ... | incremental9}]
[-timeout Timeout]
```



### IMPORTANT:

- The command must be executed from the *Data\_Protector\_home\bin* directory on the front-end Web Server system. Ensure that you are logged in under a user account that is configured as described in [“Configuring user accounts” \(page 249\)](#) and that you open the command prompt with administrative rights.
- Do not close the PowerShell console while the backup session is in progress. If you close the console during the backup, some actions are not performed: the backup sessions started do finish, but the farm does not resume the original state. To resume the farm, first run the command with the *-resumefarm* option and then unquiesce the farm manually using the Microsoft SharePoint Server Central Administration or *stsadm*.

## Option description

*-help*

Displays the *SharePoint\_VSS\_backup.ps1* command usage.

*-version*

Displays the *SharePoint\_VSS\_backup.ps1* version.

`-createonly`

If this option is specified, Data Protector only creates backup specifications. Backup is not started.

`-backuponly`

If this option is specified, Data Protector only starts backup sessions using the existing backup specifications. The `-device` option is not required.

`-device DevName`

Specifies which Data Protector device to use for backup. You can specify only one device.

- ❗ **IMPORTANT:** If only one device is used to back up a multi-system farm, the corresponding backup sessions cannot run in parallel. This prolongs the time during which the farm is in read-only mode. Specifically, the farm is in read-only mode from the moment when the backup sessions are started up until all VSS snapshots are created.

To enable backup sessions to run in parallel, select different or additional devices in each backup specification before the backup is started. See the section [“Modifying backup specifications” \(page 257\)](#).

`-hardware {no_keep|keep|ir}`

Specifies that the hardware provider should be used (instead of the software provider with `-device` option specified) and, consequently, ZDB options set. The default values for ZDB options are as follows:

- Keep the replica for instant recovery: selected if `ir` is specified.
- Keep the replica after the backup: selected if `ir` or `keep` is specified.
- Configuration check mode: Strict
- Replica type: Mirror/Clone (Plex)
- Numbers of replica rotated: 3

The default ZDB backup types are as follows (provided a device is also specified):

- `no_keep`: ZDB-to-tape
- `keep`: ZDB-to-disk+tape
- `ir`: ZDB-to-disk+tape

`-overwrite`

By default, Data Protector does not create backup specifications if they already exist. If this option is specified, Data Protector overwrites the existing backup specifications with the newly-created ones. Not applicable if `-backuponly` is specified.

`-prefix PrefixName`

With this option specified, the backup specifications are created under a different name: `SharePoint_VSS_backup_PrefixName_ClientName`.

In case of backup, this option specifies which backup specifications to use: those which name contains `PrefixName`.

Non-ASCII characters in `PrefixName` are not supported.

`-outfile PathToFile`

If this option is specified, backup specification names, errors, sessions outputs, and `omnir` restore commands are written to the specified file.

`-preview`

If this option is specified, Data Protector displays information about the Microsoft SharePoint Server environment and describes the related actions without actually performing them.

-snapshot {diskonly|disktape|tapeonly}

Applicable when starting ZDB backup sessions (that is, sessions that use backup specifications in which a hardware provider is specified for use). Performs a ZDB-to-disk (diskonly), ZDB-to-tape (tapeonly) or ZDB-to-disk+tape (disktape) session.

-reduce

Applicable only to Microsoft SharePoint Server 2010. If this option is specified, the command excludes mirrored query components from backup to reduce the backup size.

-excludeindex

Applicable only to a Data Protector standard filesystem backup of the FAST Search index files (Microsoft SharePoint Server 2010). If this option is specified, Data Protector excludes data\_index folder contained in the FASTSearch home folder from backup specification. This way, the backup is faster, but the restore is more time consuming. The option enables balancing between a backup size and a time to recovery.

-mode {full|incremental|incremental1... |incremental9}

Applicable only to a Data Protector standard filesystem backup of the FAST Search index files (Microsoft SharePoint Server 2010). With this option specified, either a Full or Incremental or leveled incremental backup can be started. By default, the Full backup is performed.

When the incremental option is specified and the Full backup does not exist, the option is ignored and the Full filesystem backup of the FAST Search index files is started.

-resumecert

Applicable only to Microsoft FAST Search Server 2010. If this option is specified, the FAST Search certificates for the content and the query connectors are reinstalled.

- 
- ❗ **IMPORTANT:** The `SharePoint_VSS_backup.ps1 -resumecert` command must be started on the Microsoft SharePoint Server system where the SharePoint Server Search 14 service is enabled.
- 

-resume farm

To be used after restore. This option returns the farm to a working state by resuming all background activities and crawling, unlocking sites, and starting Microsoft SharePoint Server services.

- 
- ❗ **IMPORTANT:** The command with the `-resume farm` option specified uses the WMI (Windows Management Instrumentation) to remotely start any stopped SharePoint services. To ensure its proper operation, an exception must be added to the Windows Default Firewall for Remote administration, which adds the WMI ports, or for the WMI directly. For details, see: <http://support.microsoft.com/kb/154596>
- 

-timeout *Timeout*

This option sets the timeout in minutes after which the crawl of the FAST Search index files is aborted and the farm is resumed. If not specified, the default timeout is 15 minutes.

## Starting Windows PowerShell

1. Log in to the Microsoft SharePoint Server system where Windows PowerShell and User Interface component are installed, under a user account that is configured as described in “Configuring user accounts” (page 249).
2. Open the Windows PowerShell CLI. For example:  
**Start > Programs > Accessories > Windows PowerShell > Windows PowerShell**
3. In case you have Windows User Account Control (UAC) enabled, ensure that you open the CLI with administrative rights. Otherwise, you will not be able to run the Data Protector PowerShell command.

4. Ensure that the Windows PowerShell execution policy is set to RemoteSigned or Unrestricted.

“Displaying the Data Protector PowerShell command syntax” (page 256) shows how the Windows PowerShell execution policy is set to Unrestricted and how the Data Protector PowerShell command syntax is displayed.

**Figure 118 Displaying the Data Protector PowerShell command syntax**



```
Administrator:SharePoint 2010 Management Shell
PS C:\Program Files\OmniBack\bin> .\SharePoint_VSS_backup.ps1 -help

Usage synopsis:

SharePoint_VSS_backup.ps1 -version ! -help
SharePoint_VSS_backup.ps1 -createonly CreateOptions
SharePoint_VSS_backup.ps1 -backuponly [BackupOptions]
SharePoint_VSS_backup.ps1 -resumefarm [-preview] ! -resumeecert

CreateOptions
 <-device DeviceName> ! -hardware <no_keep ! keep ! ir> [-device DeviceName]
 [-overwrite]
 [-prefix PrefixName]
 [-excludeindex]

BackupOptions
 [-outfile PathToFile]
 [-prefix PrefixName]
 [-preview]
 [-snapshot <diskonly ! disktape ! tapeonly>]
 [-reduce]
 [-mode {full ! incremental ! incremental1 ... ! incremental9}]

-version
 Shows the version of script.
-help
 Displays this help information.
-preview
 Shows all the farm information and actions to be taken. Does not actually
 perform any action and does not start the backup(s).
-createonly
 Only creates backup specifications.
-overwrite
 Overwrite the backup specifications during their creation. Not applicable
 for -backuponly.
-backuponly
 Performs backup only. Backup specification are not created, -device option
 not required with -backuponly
-device <DP device name>
 Device name to be used in created backup specifications.
 For backup specification creation either '-device' or '-hardware' option
 has to be present.
 If more than one host is backed up, the backups will not run in parallel
 with one device. In the destination page of the backup specification, you
 can select different or additional devices. For more details about
 modifying backup specification, see documentation.
-hardware <no_keep ! keep ! ir>
 With this option created backup specification uses USS hardware providers.
 Specify no_keep, keep or ir to specify whether to keep created disk copy and
 tracks it for instant recovery.
 For backup specification creation either '-device' or '-hardware' option has
 to be present.
-prefix <prefix>
 Additional prefix for backup specifications names.
-reduce
 Script will exclude mirrored query components from backup to reduce the size
 of backup. Applicable only for SharePoint 2010.
-excludeindex
 Exclude FASTSearch index data from datalist. Applicable only for FASTSearch DA datalis
-mode {full ! incremental ! incremental1 ... ! incremental9}
 This option is used for starting full or incremental or leveled incremental backup.
 If you don't use this option or if you use this option on wrong way by default
 backup mode will be full. Applicable only for FASTSearch DA datalist.
-snapshot
 Reinstall FASTSearch certificates for content and query connectors.
-snapshot <diskonly ! disktape ! tapeonly>
 This option is used for starting backup session to disk or to tape or disk+tape .
 Must be in use for backup specification that use hardware provider.
-outfile <filename>
 Writes backup specifications names/restore and/or recovery commands/session
 output to file specified.
-resumeecert
 Resumes all farm(s) activities.
```

## Creating backup specifications (examples)

1. To create backup specifications in which the backup device filelib\_writer1 is specified for use, run:  
SharePoint\_VSS\_backup.ps1 -createonly -device filelib\_writer1
2. To create backup specifications with the label weekly in their names and in which the backup device dev1 is specified for use, run:  
SharePoint\_VSS\_backup.ps1 -createonly -device dev1 -prefix weekly
3. To create ZDB backup specifications in which the backup device dev1 and the hardware provider (ZDB disk array) are specified for use, and in which the ZDB option Keep the replica for instant recovery is enabled, run:  
SharePoint\_VSS\_backup.ps1 -createonly -hardware ir -device dev1



4. Applicable only to a Data Protector standard filesystem backup of the FAST Search index files (Microsoft SharePoint Server 2010).

To create filesystem backup specifications in which the backup device `dev1` is specified for use and with the `data_index` folder, contained in the `FASTSearch` home folder, excluded from the backup of the FAST Search index files, run:

```
SharePoint_VSS_backup.ps1 -createonly -device dev1 -excludeindex
```

## Modifying backup specifications

To modify a backup specification, open the Data Protector GUI. In the Context list, select **Backup** and, under **MS Volume Shadow Copy Writers** or under **Filesystem** (if performing a standard filesystem backup of the FAST Search index files), click the name of the backup specification that you want to modify (see [Figure 116 \(page 251\)](#)).

### Source page

If you want to modify the Source page of the backup specification (for example, you want to back up individual Microsoft SharePoint Server databases), consider the following:

- The configuration database and the Central Administration content database must both be backed up during the same time period which starts with the suspend of the Microsoft SharePoint Server (SharePoint farm) and ends with the resume of the SharePoint farm to ensure the data consistency.
- **Microsoft Office SharePoint Server 2007:** The Shared Services Provider database (`SSP_DB`), Search database (`SSP_Search_DB`), and the associated search index files must all be backed up during the same time period which starts with the suspend of the Microsoft SharePoint Server (SharePoint farm) and ends with the resume of the SharePoint farm to ensure the data consistency.
- **Microsoft SharePoint Server 2010:**
  - The SharePoint Service Applications, Search database (`SSA_Search_DB`), and the associated search index files must all be backed up during the same time period which starts with the suspend of the Microsoft SharePoint Server (SharePoint farm) and ends with the resume of the SharePoint farm to ensure the data consistency.
  - The FAST search index files and the FAST Content SSA crawl components must all be backed up during the same time period which starts with the suspend of the Microsoft SharePoint Server (SharePoint farm) and ends with the resume of the SharePoint farm to ensure the data consistency.
- The Help Search database and the associated index files must all be backed up during the same time period which starts with the suspend of the Microsoft SharePoint Server (SharePoint farm) and ends with the resume of the SharePoint farm to ensure the data consistency.

Otherwise, after restore, the Microsoft SharePoint Server data may not be consistent.

### Destination page

In the Destination page of the backup specification, you can select different or additional devices and set the device and media options.

### Options page

In the Options page of the backup specification, you can modify backup options. For the standard filesystem backup of the FAST search index files leave the **Use Shadow Copy** option specified to enable the use of the VSS. To modify ZDB options, click **Advanced** in the Backup Specifications field and then, in the Backup Options dialog box, click the **Advanced backup options** tab.

## Starting backup sessions (examples)

1. To preview the actions that are performed when a backup session is started, run:  
`SharePoint_VSS_backup.ps1 -backuponly -prefix dev -preview`  
The following output is from a Microsoft Office SharePoint Server 2007 environment:

```
=====
Starting MOSS backup command
02/10/2011 03:16:30
=====

List of hosts and their services

virtual20
Application Server
 Windows SharePoint Services Help Search
 Windows SharePoint Services Database
 Information Management Policy Configuration Service
 Office SharePoint Server Search
 Shared Services Timer
 Office SharePoint Server Search Admin Web Service
 Excel Calculation Services
 Single Sign-on Service
 SSP Job Control Service
 Portal Service
 Office SharePoint Server Search
 Document Conversions Launcher Service
 Document Conversions Load Balancer Service
 Windows SharePoint Services Web Application
 Central Administration
 Windows SharePoint Services Incoming E-Mail
 Windows SharePoint Services Administration
 Windows SharePoint Services Timer

VIRTUAL21
Application Server
 Windows SharePoint Services Help Search
 Office SharePoint Server Search
 Shared Services Timer
 Office SharePoint Server Search Admin Web Service
 Single Sign-on Service
 SSP Job Control Service
 Portal Service
 Office SharePoint Server Search
 Windows SharePoint Services Web Application
 Windows SharePoint Services Administration
 Windows SharePoint Services Help Search
 Windows SharePoint Services Timer
```

```

VIRTUAL23
Application Server
 Windows SharePoint Services Help Search
 Office SharePoint Server Search
 Shared Services Timer
 Office SharePoint Server Search Admin Web Service
 Single Sign-on Service
 SSP Job Control Service
 Portal Service
 Office SharePoint Server Search
 Windows SharePoint Services Web Application
 Windows SharePoint Services Administration
 Windows SharePoint Services Help Search
 Windows SharePoint Services Timer

SQL hosts list
virtual20

Index hosts list
virtual20
VIRTUAL21
VIRTUAL23

Help search hosts list
VIRTUAL21
VIRTUAL23

SUSPENDING FARM
02/10/2011 03:16:43

Farm SharePoint_Config

Service Windows SharePoint Services Help Search on host VIRTUAL21
-> Pausing background activity ...
 ... background activity paused.

Service Windows SharePoint Services Help Search on host VIRTUAL23
-> Pausing background activity ...
 ... background activity paused.

Web applications:
 Display name: Recovery Web Application
 Alternate URL: http://virtual20:999

 Display name: SharePoint - 123
 Alternate URL: http://virtual20:123
 Web site URL: http://virtual20:123/ssp/admin
 Root title: Shared Services Administration: SSP1
 -> Setting lock state to readonly
 Crawled by: , id
 Crawl status:
 -> Pausing background activity
 ...
Quiesce status is: Quiesced

SUSPENDING END
02/10/2011 03:18:28

```

-> Starting backups...

```
Starting backup: omnib -msvssw_list SharePoint_VSS_backup_dev_virtual20 \ -barmode full
Starting backup: omnib -msvssw_list SharePoint_VSS_backup_dev_VIRTUAL21 \ -barmode full
Starting backup: omnib -msvssw_list SharePoint_VSS_backup_dev_VIRTUAL23 \ -barmode full
```

```
Waiting while VSS creates Volume Shadow Copies ...
Please wait. DO NOT close PowerShell console!
After shadow copies are created, the command will resume farm
and display Data Protector backup session(s) output(s).
SUCCESS: Volume Shadow Copy successfully created.
Host : virtual20
SUCCESS: Volume Shadow Copy successfully created.
Host : VIRTUAL21
SUCCESS: Volume Shadow Copy successfully created.
Host : VIRTUAL23
```

```

RESUMING FARM
02/10/2011 03:18:28

```

```
Service Windows SharePoint Services Help Search on host VIRTUAL21
-> Resuming background activity ...
... background activity resumed
```

```
Service Windows SharePoint Services Help Search on host VIRTUAL23
-> Resuming background activity ...
... background activity resumed
```

```
Web site URL: http://virtual20:123/ssp/admin
Root title: Shared Services Administration: SSP1
-> Reverting lock for site http://virtual20:123/ssp/admin to none
-> Resuming background activity
...
```

```

RESUMING END
02/10/2010 03:19:18

```

```
=====
MOSS backup command finished
02/10/2011 03:19:18
Running time 00:02:48.3336122
=====
```

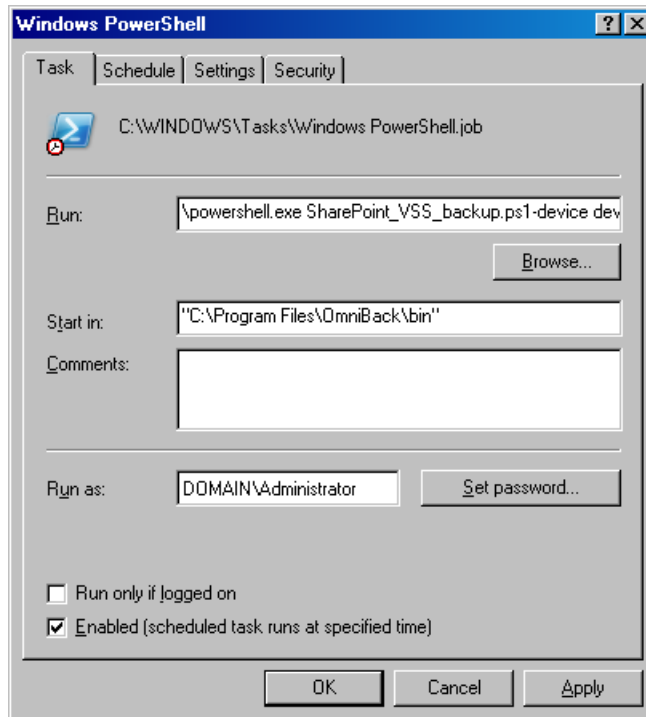
2. To start backup sessions using the existing backup specifications that have no prefix in their names, run:  
SharePoint\_VSS\_backup.ps1 -backuponly
3. To start backup sessions using the existing backup specifications that have the prefix weekly in their names, run:  
SharePoint\_VSS\_backup.ps1 -backuponly -prefix weekly
4. To start backup sessions using the existing backup specifications that have no prefix in their names and to save the output of the sessions and the associated restore commands to the file c:\logs\shp.log, run:  
SharePoint\_VSS\_backup.ps1 -backuponly -outfile C:\logs\shp.log
5. To start ZDB-to-disk backup sessions using the existing ZDB backup specifications that have no prefix in their names, run:  
SharePoint\_VSS\_backup.ps1 -backuponly -snapshot diskonly
6. To start incremental filesystem backup sessions of the FAST Search index files (Microsoft SharePoint Server 2010), run:  
SharePoint\_VSS\_backup.ps1 -backuponly -mode incremental

## Scheduling backup sessions

You can schedule backup sessions using the Windows system scheduler.

1. On the front-end Web server system, create a Windows PowerShell scheduled task. Go to:  
**Start > Settings > Control Panel > Scheduled Tasks > Add Scheduled Task**
2. Open advanced properties for the task.

**Figure 119 Scheduling a backup session using the Windows scheduler**



In **Run**, type:

`Windows_PowerShell_home\powershell.exe SharePoint_VSS_backup.ps1 [Options]`

For details on Options, see [“The command syntax”](#) (page 253).

In **Start in**, type:

`Data_Protector_home\bin`

In **Run as**, type a Windows domain user account `DOMAIN\UserName` that is configured as described in [“Configuring user accounts”](#) (page 249).

## Restore

To restore Microsoft SharePoint Server data:

- Stop Microsoft SharePoint Server services
- Restore the data.
- Return the farm to a working state.

For details, see the following sections.

## Before you begin

- Stop and disable the following services:
  - IIS Admin Service (Only for Internet Information Services 6.0 on Windows Server 2003, when the whole farm is restored)
  - Office SharePoint Server Search (Microsoft Office SharePoint Server 2007)
  - SharePoint Server Search 14 (Microsoft SharePoint Server 2010)

In addition, stop the following services:

- Microsoft Office SharePoint Server 2007
  - Windows SharePoint Services Administration
  - Windows SharePoint Services Search
  - Windows SharePoint Services Timer
- Microsoft SharePoint Server 2010
  - SharePoint 2010 Administration
  - SharePoint Foundation Search V4
  - SharePoint 2010 Timer
  - SharePoint 2010 Tracing
  - FAST Search for SharePoint
  - FAST Search for SharePoint Monitoring
- Put the Microsoft SQL Server instance offline if you plan to restore one of the following Microsoft SQL Server databases:
  - master
  - model
  - msdb
  - a database for which Microsoft SQL Server mirroring is enabled

---

### NOTE:

- If you use `SqlServerWriter`, you can restore the `model` and `msdb` databases also when the Microsoft SQL Server instance is online. This is one advantage over `MSDE writer`.
  - *Microsoft SQL Server mirroring*: If the original and mirror database reside in separate Microsoft SQL Server instances, put offline both Microsoft SQL Server instances.
- 

## Restoring data

You can restore Microsoft SharePoint Server data using the Data Protector GUI or CLI.

### Considerations

- The configuration database and the Central Administration content database must both be restored using backups from the same point in time (the backups performed in the same period in which the Microsoft SharePoint Server (SharePoint farm) was in suspended mode) to ensure the data consistency. Since the configuration database and the Central Administration content database contain system-specific information, you can restore them only to the original environment or to an environment that has precisely the same configuration, software updates, server names, and number of servers.
- **Microsoft Office SharePoint Server 2007**: The Shared Services Provider database (`SSP_DB`), Search database (`SSP_Search_DB`), and the associated search index files must all be restored

using backups from the same point in time (the backups performed in the same period in which the Microsoft SharePoint Server (SharePoint farm) was in suspended mode) to ensure the data consistency.

- **Microsoft SharePoint Server 2010:**

- The SharePoint Service Applications, Search database (SSA\_Search\_DB), and the associated search index files must all be restored using backups from the same point in time (the backups performed in the same period in which the Microsoft SharePoint Server (SharePoint farm) was in suspended mode) to ensure the data consistency.
  - Since the FAST configuration database and the FAST Search home folder contain system-specific information, you can restore them only to the original environment or to an environment that has precisely the same configuration, software updates, server names, and number of servers.
  - The FAST Search index files and the FAST Content SSA crawl components must all be restored using backups from the same point in time (the backups performed in the same period in which the Microsoft SharePoint Server (SharePoint farm) was in suspended mode) to ensure the data consistency.
- The Help Search database and the associated index files must all be restored using backups from the same point in time (the backups performed in the same period in which the Microsoft SharePoint Server (SharePoint farm) was in suspended mode) to ensure the data consistency.
  - The following table shows which VSS restore modes are supported for which writers:

**Table 32 VSS supported restore modes and writers**

| Writers                                     | VSS restore modes         |                                     |
|---------------------------------------------|---------------------------|-------------------------------------|
|                                             | Restore to another client | Restore files to temporary location |
| MSDE writer<br>SqlServerWriter              | No                        | Yes (manual attach needed)          |
| OSearch VSS writer<br>OSearch14 VSS writer  | Yes                       | No                                  |
| SPSearch VSS writer<br>SPSearch4 VSS writer | Yes                       | No                                  |

## Prerequisites

- Applicable only to a Data Protector filesystem restore of the FAST Search index files (Microsoft SharePoint Server 2010). Before restoring the FAST Search index files, the **Overwrite** option must remain selected to ensure the data consistency. It is selected by default.

## Restoring using the Data Protector GUI

1. In the Context List, click **Restore**.
2. In the Scoping Pane, expand **MS Volume Shadow Copy Writers**, expand the client which data you want to restore, and then click **MS Volume Shadow Copy Writers**.  
If performing a filesystem restore of the FAST Search index files (Microsoft SharePoint Server 2010), expand **Filesystem**, expand the client which data you want to restore, and then click the filesystem object.
3. In the Source page, select the data that you want to restore.

Figure 120 Selecting Microsoft Office SharePoint Server 2007 databases for restore

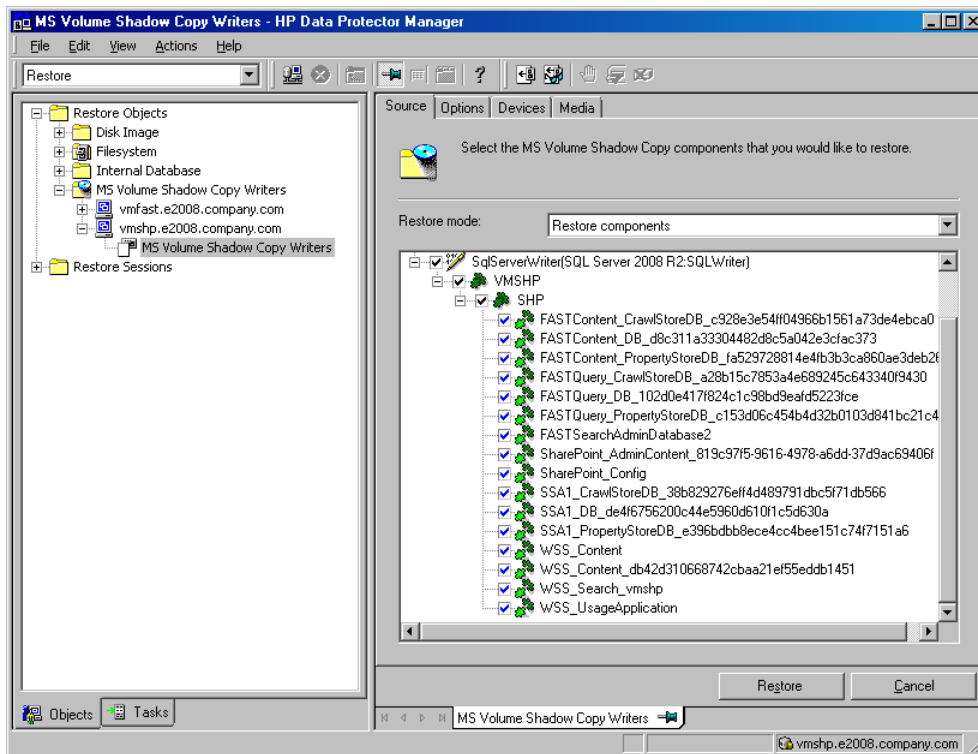
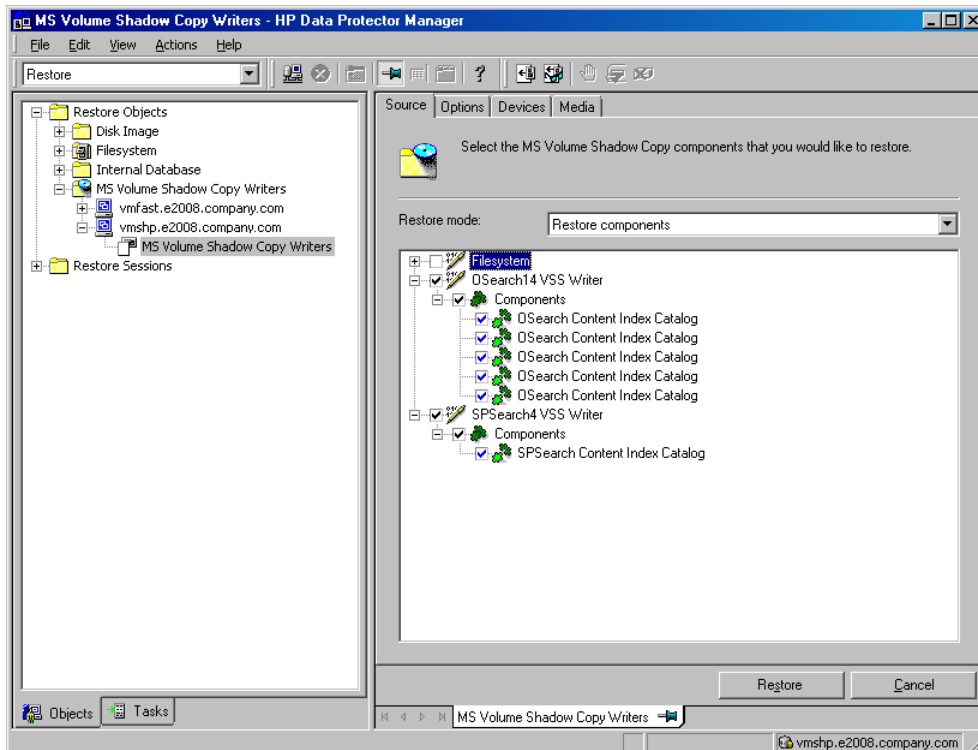


Figure 121 Selecting Microsoft Office SharePoint Server 2007 Search index files for restore



4. In the Options page, specify the restore options.
5. In the Devices page, select devices to use for restore.
6. Click **Restore**, review your selection, and click **Finish**.



## Restoring using the Data Protector CLI

You can restore Microsoft SharePoint Server data using the Data Protector `omnir` command. For details, see the `omnir` man page or the *HP Data Protector Command Line Interface Reference*.

If you specified the `-outfile` option when you ran backup sessions, you can find the necessary `omnir` commands in the specified file. The following is an example of the `omnir` command from such a file.

```
omnir -vss -barhost SHP-APP
-session 2011/09/25-13
-tree "/SqlServerWriter(SQL Server 2005:SQLWriter)/SHP-APP/master"
-session 2011/09/25-13
-tree "/SqlServerWriter(SQL Server2005:SQLWriter)/SHP-APP/model"
-session 2011/09/25-13
-tree "/SqlServerWriter(SQL Server 2005:SQLWriter)/SHP-APP/msdb"
-session 2011/09/25-13
-tree "/SqlServerWriter(SQL Server 2005:SQLWriter)/SHP-APP/
WSS_Content_SSPOAdminAccounting"
-session 2011/09/25-13
-tree "/SqlServerWriter(SQL Server 2005:SQLWriter)/SHP-APP/SSP_Accounting"
-session 2011/09/25-13
-tree "/SqlServerWriter(SQL Server 2005:SQLWriter)/
SHP-APP/SSP_Accounting_Search"
```

### Limitations

The `omnir` command syntax should not contain more than 8191 characters. If you have so many `-tree` objects that the syntax exceeds 8191 characters, split the objects and run two separate sessions.

## After the restore

After the restore:

1. Enable and start the service IIS Admin Service (Only for IIS 6 on Windows 2003, when the whole farm was restored)
2. Enable the service Office SharePoint Server Search or SharePoint Server Search 14.
3. Bring the Microsoft SQL Server instances online (if offline).
4. Return the farm to a working state (that is, resume background activities and crawling, unlock sites, and start the Microsoft SharePoint Server services) by running:

```
SharePoint_VSS_backup.ps1 -resumefarm
```

---

### NOTE:

- The command uses the WMI (Windows Management Instrumentation) to remotely start any stopped SharePoint services. Ensure its proper operation by adding an exception to the Windows Default Firewall for Remote administration, which adds the WMI ports, or for the WMI directly. For details, see: <http://support.microsoft.com/kb/154596>
- If the FAST Search certificates for the content and query connectors are out of sync, you can reinstall them by running:

```
SharePoint_VSS_backup.ps1 -resumecert
```

Start the command on the Microsoft SharePoint Server system where the SharePoint Server Search 14 service is enabled.

---

## Restoring index files on the Query system

This section is applicable for Microsoft Office SharePoint Server 2007 only. The Office SharePoint Server Search service is enabled on two separate Microsoft Office SharePoint Server 2007 systems, so that one is assigned the Indexing and the other the Query role.

To copy the newly restored index files from the Indexing system to the Query system, perform the following steps (depending on which Microsoft Office SharePoint Server 2007 and Windows Shared Services service pack you have):

- **Service Pack 1**

1. On the Query system, stop and disable the service Office SharePoint Server Search.
2. Copy the index files from the Indexing to the Query system.

By default, index files are located in the C:\Program Files\Microsoft Office Servers\12.0\Data\Office Server\Applications directory.

3. On the Query system, enable and start the service Office SharePoint Server Search.

- **Service Pack 2**

On the Query system, run:

```
stsadm -o search -reprovisionindex -ssp SSPName
```

for each Shared Services Provider separately.

## Troubleshooting

This section lists general checks and verifications, plus problems you might encounter when using the Data Protector Microsoft SharePoint Server VSS based solution.

For Microsoft Volume Shadow Copy troubleshooting information, see the troubleshooting chapter in the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

For general Data Protector troubleshooting information, see the *HP Data Protector Troubleshooting Guide*.

## Before you begin

- Ensure that the latest official Data Protector patches are installed. On how to verify this, see the online Help index: "patches".
- See the *HP Data Protector Product Announcements, Software Notes, and References* for general Data Protector limitations, as well as recognized issues and workarounds.
- See <http://www.hp.com/support/manuals> for an up-to-date list of supported versions, platforms, and other information.

## Checks and verifications

If your browsing, backup, or restore failed:

- Examine system errors reported in the `debug.log` located in:  
`Data_Protector_home\log`.
- Check if you can do a filesystem backup and restore on the problematic client. For information, see the online Help.

## After restore, you cannot connect to the Central Administration web page

### Problem

After restore, when you try to connect to the Microsoft SharePoint Central Administration web page, an error similar to the following is displayed in your web browser:

Windows Internet Explorer:

```
Retrieving the COM class factory for component with CLSID
(BDEADDEE2-C265-11D0-BCED-00A0C90AB50F) failed due to the following error
800703fa.
```

Mozilla Firefox:

An unexpected error has occurred.

### Action

1. Restart Microsoft SharePoint Server services on all clients in the farm.
2. Open the Internet Information Services (IIS) Manager and restart all application pools.
3. In case an application pool fails to be restarted with the following error:  
Cannot Restore Application Pool. There was an error while performing this operation.  
wait for a few seconds and then restart the operation.
4. Delete browsing history in your web browser.
5. Log in to the Central Administration web page.

## Backup fails with the error Failed to resume Service Windows SharePoint Services Help Search

### Problem

When you start backup sessions, an error similar to the following is displayed:

```
Service Windows SharePoint Services Help Search on host
```

```
MOSS07-INDEX
```

```
-> Resuming background activity ...
```

```
ERROR: Failed to resume Service Windows SharePoint Services Help Search
on host MOSS07-INDEX
```

```
Web site URL: http://moss07-web:2001
```

```
Root title: as
```

```
-> Resuming background activity
```

### Action

Run:

```
SharePoint_VSS_backup.ps1-resumefarm
```

## After restore, a quiesce operation fails

### Problem

After you have restored the configuration database and executed `SharePoint_VSS_backup.ps1-resumefarm`, the data in the Microsoft SharePoint Server file system caches on front-end Web Server systems is not consistent with the data in the newly-restored configuration database. When you try to quiesce the farm, the operation fails with the following error:

An unhandled exception occurred in the user interface. Exception Information: An update conflict has occurred, and you must re-try this action. The object SessionStateService Parent=SPFarm Name=< farm\_config\_database\_name > is being updated by < domain\username >, in the w3wp process, on machine < servername >. View the tracing log for more information about the conflict.

#### Action

Clear the Microsoft Office SharePoint Server file system cache on all server systems in the farm. For details, see:

<http://support.microsoft.com/kb/939308>

## After restore, you cannot connect to the FAST Search Server

#### Problem

After restore, when you try to connect to the Microsoft FAST Search Server 2010 system for SharePoint, the operation fails.

FAST Query SSA search operations display an error similar to the following:

The search request was unable to connect to the Search Service.

#### Action

Run:

```
SharePoint_VSS_backup.ps1 -resumecert
```

---

**NOTE:** The VSS based solution copies the FAST Search certificate `FASTSearchCert.pfx` from the FAST Admin Server system to the SharePoint Server system and installs it. Also, the SharePoint certificate is copied and installed to each FAST Search Server system. For details, see: <http://technet.microsoft.com/en-us/library/ff381244.aspx>.

---

## The SharePoint\_VSS\_backup.ps1 script stops responding and the farm stays in read only mode

#### Problem

When starting a back up, the `SharePoint_VSS_backup.ps1` script stops responding when a crawl of the Microsoft SharePoint Server is being performed. The issue can appear due to external conditions such as a corrupted SSA index, the need to reissue the certificate manually and so on.

As a result, the farm stays in read-only mode.

#### Action

Normally, the crawl should be aborted automatically after 15 minutes. If this does not happen:

1. Abort the script by pressing **Ctrl-C**.
2. Manually resume the farm.

You can specify a different timeout after which the crawl is aborted and the farm is resumed by using the `-timeout` option.

---

# A Appendix

## In this appendix

This appendix gives information on the following topics:

- “Reconfiguring an Oracle instance for instant recovery” (page 269)
- “ZDB integrations omnirc variables” (page 271)

## Reconfiguring an Oracle instance for instant recovery

If the control files or redo logs are located on the same volume group (if LVM is used) or source volume as the database files, the control files and online redo logs are overwritten during instant recovery. In such case, you may want to reconfigure the Oracle instance. For details on the required configuration, see “Oracle backup set ZDB concepts” (page 24) and to “Oracle proxy-copy ZDB concepts” (page 28). For additional examples on how to move the redo logs and control files, see “Examples for moving the control files and redo logs to different locations ” (page 270).

### Moving online redo logs

To move the *online redo log files* from the source volumes to be replicated, to other locations:

1. List the online redo log files:  

```
$ sqlplus
SQL> select member from v$logfile;
```
2. Shut down the database:  

```
SQL> connect user/password@service as sysdba;
SQL> shutdown
SQL> exit
```
3. Move the log files to a different location using operating system tools.
4. Start the database in mount mode:  

```
$ sqlplus
SQL> connect user/password@service as sysdba;
SQL> startup mount;
```
5. Register the new locations for each moved file:  

```
SQL> alter database rename file 'OldPathName' to 'NewPathName';
where OldPathName and NewPathName are full paths to the log file.
```
6. Open the database in normal mode:  

```
SQL> alter database open;
```

### Moving control files

To move the *control files* from the source volumes to be replicated, to other locations:

1. Determine if the database uses the SPFILE parameter:  

```
SQL> show parameter SPFILE
```
2. If the database does not use SPFILE:
  - a. Shut down the database.  

```
SQL> shutdown
```
  - b. Move the control files to a different location using operating system tools.

- c. Edit the `CONTROL_FILES` parameter in the database's initialization parameter file (usually located in the `$ORACLE_HOME/dbs/initSID.ora` directory) to change the existing control file names:

```
control_files = ("NewPathName", ...)
```

- d. Restart the database:

```
SQL> startup
```

If the database uses SPFILE:

- a. Specify the new location for control files by running the following command:

```
SQL> alter system set control_files='NewPathName1',
'NewPathName2',..., scope=spfile
```

- b. Shut down the database.

```
SQL> shutdown
```

- c. Move the control files to a different location.

- d. Restart the database:

```
SQL> startup
```

## Examples for moving the control files and redo logs to different locations

### Example - moving online redo logs

In the following example for Oracle10g on HP-UX, the data files are on the same source volume as the control files and redo logs, which is `/opt/oracle/product/10.2.0`.

To move the *online redo log files* from `/opt/oracle/product/10.2.0` to `/oracle/logs` (which is not replicated):

1. List the online redo log files:

```
$ sqlplus
```

```
SQL> select member from v$logfile;
```

```
/opt/oracle/product/10.2.0/oradata/redo01.log
/opt/oracle/product/10.2.0/oradata/redo02.log
/opt/oracle/product/10.2.0/oradata/redo03.log
```

List the filenames and tablespaces to check whether they are on the same source volumes as the control files:

```
SQL> select FILE_NAME, TABLESPACE_NAME, BYTES from dba_data_files;
```

```
FILE_NAME
```

```

TABLESPACE_NAME BYTES

/opt/oracle/product/10.2.0/oradata/system01.dbf
SYSTEM 419430400
/opt/oracle/product/10.2.0/oradata/undotbs01.dbf
UNDOTBS1 377487360
/opt/oracle/product/10.2.0/oradata/cwmlite01.dbf
CWMLITE 20971520
```

2. Shut down the database:

```
SQL> connect user/password@service as sysdba;
```

```
SQL> shutdown
```

```
SQL> exit
```

3. Move the log files to a different location.

```
$ mv /opt/oracle/product/10.2.0/oradata/redo* /oracle/logs
```

4. Start the database in mount mode:
 

```
$ sqlplus
SQL> connect user/ password@service as sysdba;
SQL> startup mount;
```
5. Rename the new locations for each moved file:
 

```
alter database rename file
'/opt/oracle/product/10.2.0/oradata/redo01.log' to
'/oracle/logs/redo01.log';

Database altered.

alter database rename file
'/opt/oracle/product/10.2.0/oradata/redo02.log' to
'/oracle/logs/redo01.log';

Database altered.

alter database rename file
'/opt/oracle/product/10.2.0/oradata/redo03.log' to
'/oracle/logs/redo01.log';

Database altered.
```
6. Open the database in normal mode:
 

```
SQL> alter database open;
```

#### Example - moving control files for Oracle10g

In the following example, the Oracle10g database uses SPFILE. To move the control files from /opt/oracle/product/10.2.0/ to /oracle/oractl:

1. Determine if the database uses the SPFILE parameter:
 

```
SQL> show parameter spfile;
```

| NAME   | TYPE   | VALUE             |
|--------|--------|-------------------|
| spfile | string | ?/dbs/spfile@.ora |
2. Specify the new location for the control files by running the following command (in a single line and without the "\" characters):
 

```
SQL> alter system setcontrol_files='/oracle/logs/RCVCAT \
/control01.ctl','/oracle/logs/RCVCAT/control02','/oracle \
/logs/RCVCAT/control03.ctl' scope=spfile;
```
3. Shut down the database:
 

```
SQL> shutdown
```
4. Move the control files to the new location:
 

```
mv /opt/oracle/product/10.2.0/oradata/control* /oracle/oractl
```
5. Restart the database:
 

```
SQL> startup
```

## ZDB integrations omnirc variables

The Data Protector ZDB integrations use environment variables, which can be set in the /opt/omni/.omnirc (on UNIX systems) or *Data\_Protector\_home\omnirc* file (on Windows systems), on both the application and backup systems. These variables are used for Data Protector ZDB integrations customizing. See the online Help index: "omnirc options" for information on how to use the omnirc file.

For information on Data Protector ZDB agents omnirc file variables, see the *HP Data Protector Zero Downtime Backup Administrator's Guide*.

This section explains the `omnirc` file variables that can be set for the Data Protector ZDB integrations.

**ZDB\_ORA\_INCLUDE\_CF\_OLF:** Data Protector Oracle Server integration and Data Protector SAP R/3 integration-related variable.

---

**NOTE:** This variable is not supported on EMC.

---

The default value is 0. Possible values are 0 and 1.

If an offline backup is performed using the Data Protector SAP R/3 integration, the variable is ignored and the integration behaves as if the variable was set to 1.

#### Instant recovery

The instant recovery process depends on whether the control file and redo logs reside on the same disk array source volume as datafiles or not:

- If this variable is set to 0 (default), during a ZDB session, Data Protector creates target volumes only for the source volumes containing Oracle datafiles. Target volumes for source volumes containing Oracle control file and Oracle online redo logs are not created.  
For Oracle proxy-copy or backup set ZDB and restore concepts when this variable is set to 0, see [“Oracle backup set ZDB concepts”](#) (page 24) and [“Oracle proxy-copy ZDB concepts”](#) (page 28). For SAP R/3 backup and restore concept when this variable is set to 0, see [“Integration concepts”](#) (page 105).
- If this variable is set to 1, Data Protector creates target volumes for all source volumes containing Oracle datafiles, Oracle control file, and if Oracle integration is used, Oracle online redo logs.

---

❗ **IMPORTANT:** If the `ZDB_ORA_INCLUDE_CF_OLF` variable is set to 1 the control files and redo logs are overwritten during instant recovery.

---

#### Opening the database on the backup system

To successfully open the database on the backup system for *other* purposes than Data Protector, note:

- With Oracle proxy-copy ZDB method, set this variable to 1.
- With Oracle backup set ZDB method, used with the Oracle integration, you can always open the database on the backup system.

#### Prerequisites


The prerequisites for this variable to be set to 1 are:

- Data Protector Oracle Server integration: Oracle datafiles, Oracle control file, and Oracle online redo logs must be installed on a disk array.
- Data Protector SAP R/3 integration: Oracle datafiles and Oracle control file must be installed on a disk array.

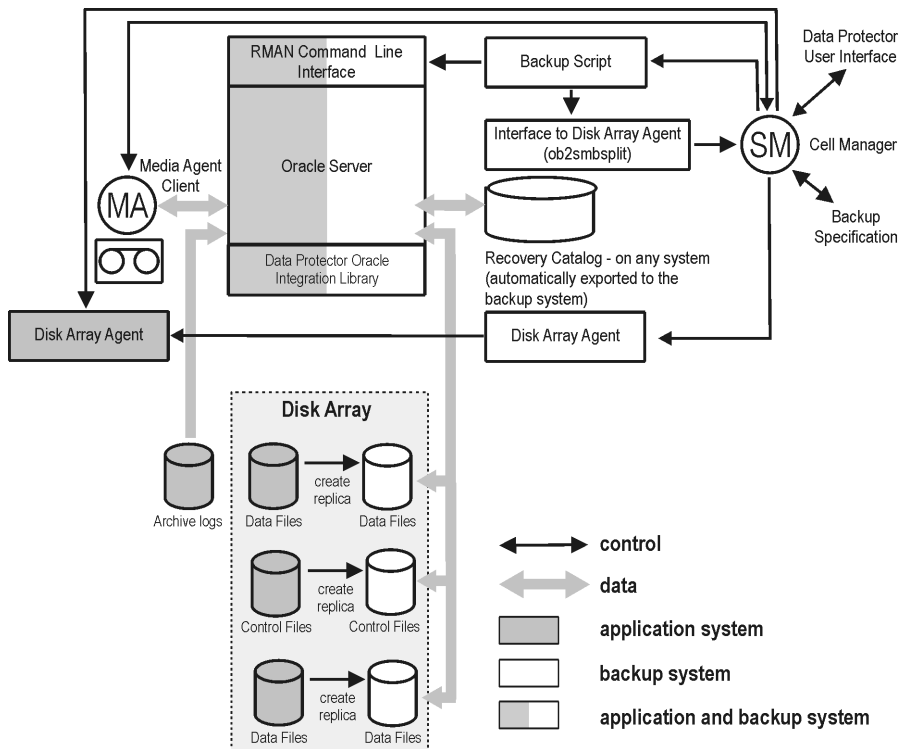
See [“Oracle proxy-copy ZDB and restore concepts when the `ZDB\_ORA\_INCLUDE\_CF\_OLF` variable is set to 1”](#) (page 273) and [“Oracle backup set ZDB and restore concept when the `ZDB\_ORA\_INCLUDE\_CF\_OLF` variable is set to 1”](#) (page 273) for Oracle backup and restore concepts when this variable is set to 1. See [“SAP R/3 backup and restore concept when the `ZDB\_ORA\_INCLUDE\_CF\_OLF` variable is set to 1 with online backup, or in case of offline backup”](#) (page 274) for SAP R/3 backup and restore concepts when this variable is set to 1.



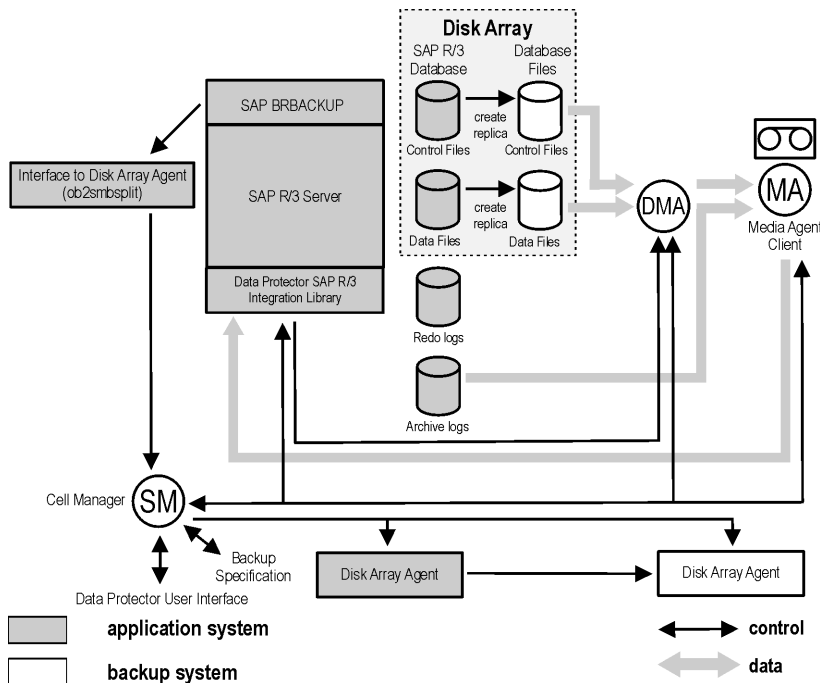
**Figure 122 Oracle proxy-copy ZDB and restore concepts when the ZDB\_ORA\_INCLUDE\_CF\_OLF variable is set to 1**



**Figure 123 Oracle backup set ZDB and restore concept when the ZDB\_ORA\_INCLUDE\_CF\_OLF variable is set to 1**



**Figure 124 SAP R/3 backup and restore concept when the ZDB\_ORA\_INCLUDE\_CF\_OLF variable is set to 1 with online backup, or in case of offline backup**



**ZDB\_ORA\_INCLUDE\_SPF:** Data Protector Oracle Server integration-related variable.

The default value is 0. Possible values are 0 and 1.

The variable is ignored and the integration behaves as if the variable was set to 1 if offline backup is performed using the Data Protector SAP R/3 integration.

If this variable is set to 0, during ZDB sessions, Data Protector checks if Oracle Server datafiles and the Oracle Server SPFILE are located on the same source volume. If the datafiles and the SPFILE are located on the same volume and instant recovery is enabled in the backup specification, the ZDB session fails. If this variable is set to 1, Data Protector skips the check. To enable instant recovery, leave this variable set to the default value.

**CAUTION:** If this variable is set to 1 and the datafiles are located on the same volume as the SPFILE, the SPFILE is overwritten during instant recovery, potentially resulting in a data loss.

**ZDB\_ORA\_NO\_CHECKCONF\_IR:** Data Protector Oracle Server integration-related variable.

The default value is 0. Possible values are 0 and 1.

By default, the Oracle Server configuration is checked whether it is instant recovery-enabled or not (whether the Oracle control file, the Oracle Server SPFILE, and the Oracle Server online redo logs are located on volumes of a different volume group than Oracle Server datafiles or not). For Oracle Server configuration check, the Data Protector command `omniresolve` is used internally. On UNIX systems, the `omniresolve` file must have the `setuid` bit set. When this variable is set to 1, the configuration check is omitted.

**CAUTION:** Checking the Oracle Server configuration for instant recovery suitability is an essential step to ensure the instant recovery session does not result in a data loss. It is therefore not recommended to set this variable to 1 unless you ensure that the Oracle Server is and remains configured appropriately for instant recovery.

**OB2MARAWREAD\_KB:** This variable sets the read block size for Oracle and SAP R/3 ZDB integrations on UNIX systems with Oracle tablespaces or datafiles installed on disk images and when using the proxy-copy method (when using DMA).

The default value is 64 kB. The specified value must be in the range between 1 kB and 1 MB.

The specified size is automatically adjusted to a size which is a multiple of the block size. The values above 256 kB could cause the DMA to fail.

**ZDB\_TAKE\_CLUSRES\_ONLINE:** This variable specifies how many times Data Protector tries to connect to Microsoft SQL Server in case the first connection fails. A reconnection is triggered every 30 seconds. This means that Data Protector waits up to  $ZDB\_TAKE\_CLUSRES\_ONLINE \times 30$  seconds for the Microsoft SQL Server resources to start up.

**SSEA\_ATOMIC\_SPLIT:** Determines if the Data Protector HP StorageWorks P9000 XP Agent should use the atomic split configuration of a disk array of the HP P9000 XP Disk Array Family to ensure data consistency of replicas of the Oracle Server data in configurations where Automatic Storage Management (ASM) is used.

*Default:* 0 (disabled). Possible: 0 | 1.

---

# Glossary

## A

|                                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>access rights</b>                      | See user rights.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>ACSLs</b>                              | (StorageTek specific term) The Automated Cartridge System Library Server (ACSLs) software that manages the Automated Cartridge System (ACS).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Active Directory</b>                   | (Windows specific term) The directory service in a Windows network. It contains information about resources on the network and makes them accessible to users and applications. The directory services provide a consistent way to name, describe, locate, access and manage resources regardless of the physical system they reside on.                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>AES 256-bit encryption</b>             | Data Protector software encryption, based on the AES-CTR (Advanced Encryption Standard in Counter Mode) encryption algorithm that uses random keys of 256-bit length. The same key is used for both encryption and decryption. With AES 256-bit encryption, data is encrypted before it is transferred over a network and before it is written to media.                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>AML</b>                                | (ADIC/GRAU specific term) Automated Mixed-Media library.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>AMU</b>                                | (ADIC/GRAU specific term) Archive Management Unit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>application agent</b>                  | A component needed on a client to back up or restore online database integrations.<br>See also Disk Agent.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>application system</b>                 | (ZDB specific term) A system the application or database runs on. The application or database data is located on source volumes.<br>See also backup system and source volume.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>archive logging</b>                    | (Lotus Domino Server specific term) Lotus Domino Server database mode where transaction log files are overwritten only after they have been backed up.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>archived redo log</b>                  | (Oracle specific term) Also called offline redo log. If the Oracle database operates in the ARCHIVELOG mode, as each online redo log is filled, it is copied to an archived log destination. This copy is the archived redo log. The presence or absence of an archived redo log is determined by the mode the database is using: <ul style="list-style-type: none"><li>• ARCHIVELOG - The filled online redo log files are archived before they are reused. The database can be recovered if an instance or a disk fails. A "hot" backup can be performed only when the database is running in this mode.</li><li>• NOARCHIVELOG - The filled online redo log files are not archived.</li></ul> See also online redo log.                                |
| <b>ASR set</b>                            | A collection of files stored on several diskettes, required for proper reconfiguration of the replacement disk (disk partitioning and logical volume configuration) and automatic recovery of the original system configuration and user data that was backed up during the full client backup. These files are stored as an ASR archive file on the Cell Manager in the directory <code>Data_Protector_program_data\Config\Server\dr\asr</code> (Windows Server 2008), <code>Data_Protector_home\Config\Server\dr\asr</code> (other Windows systems), or <code>/etc/opt/omni/server/dr/asr</code> (UNIX systems) as well as on the backup medium. After a disaster occurs, the ASR archive file is extracted to diskettes which you need to perform ASR. |
| <b>audit logs</b>                         | Data files to which auditing information is stored.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>audit report</b>                       | User-readable output of auditing information created from data stored in audit log files.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>auditing information</b>               | Data about every backup session that was performed over an extended, user-defined period for the whole Data Protector cell.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>autochanger</b>                        | See library.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>autoloader</b>                         | See library.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Automatic Storage Management (ASM)</b> | (Oracle specific term) A filesystem and volume manager integrated into Oracle which manages Oracle database files. It eliminates complexity associated with data and disk management and optimizes performance by providing striping and mirroring capabilities.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

|                       |                                                                                                                                                                                                                                                                                                           |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>automigration</b>  | (VLS specific term) The functionality that allows data backups to be first made to the VLS' virtual tapes and then migrated to physical tapes (one virtual tape emulating one physical tape) without using an intermediate backup application.<br>See also Virtual Library System (VLS) and virtual tape. |
| <b>auxiliary disk</b> | A bootable disk that has a minimal operating system with networking and Data Protector Disk Agent installed. It can be carried around and used to boot the target system in Phase 1 of Disk Delivery Disaster Recovery of UNIX clients.                                                                   |

## B

|                             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>BACKINT</b>              | (SAP R/3 specific term) SAP R/3 backup programs can call the Data Protector backint interface program via an open interface, which enables them to communicate with Data Protector software. For backup and restore, SAP R/3 programs issue orders for the Data Protector backint interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>backup API</b>           | The Oracle interface between the Oracle backup/restore utility and the backup/restore media management layer. The interface defines a set of routines to allow the reading and writing of data to the backup media, the creation, searching and removing the backup files.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>backup chain</b>         | See restore chain.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>backup device</b>        | A device configured for use with Data Protector, which can write data to and read data from storage media. This can be, for example, a standalone DDS/DAT drive or a library.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>backup generation</b>    | One backup generation includes one full backup and all incremental backups until the next full backup.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>backup ID</b>            | An identifier of an integration object that equals the session ID of the backup of this object. The backup ID is preserved when an object is copied, exported, or imported.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>backup object</b>        | A backup unit that contains all items backed up from one disk volume (logical disk or mount point). The backed up items can be any number of files, directories, or the entire disk or mount point. Additionally, a backup object can be a database/application entity or a disk image (rawdisk).<br><br>A backup object is defined by: <ul style="list-style-type: none"> <li>• Client name: Hostname of the Data Protector client where the backup object resides.</li> <li>• Mount point: For filesystem objects — the access point in a directory structure on the client where the backup object is located (drive on Windows and mount point on UNIX). For integration objects — backup stream identification, indicating the backed up database/application items.</li> <li>• Description: For filesystem objects — uniquely defines objects with identical client name and mount point. For integration objects — displays the integration type (for example, SAP or Lotus).</li> <li>• Type: Backup object type. For filesystem objects — filesystem type (for example, WinFS). For integration objects — "Bar".</li> </ul> |
| <b>backup owner</b>         | Each backup object in the IDB has an owner. The default owner of a backup is the user who starts the backup session.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>backup session</b>       | A process that creates a copy of data on storage media. The activities are specified in a backup specification or an interactive session. All clients configured in one backup specification are backed up together in one backup session using the same backup type. The result of a backup session is a set of media, which was written to, also called the backup or media set.<br>See also backup specification, full backup, and incremental backup.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>backup set</b>           | A complete set of integration objects associated with a backup.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>backup set</b>           | (Oracle specific term) A logical grouping of backed up files that are created using the RMAN backup command. A backup set is a complete set of files associated with a backup. The files can be multiplexed to increase performance. A backup set contains either datafiles or archived logs, but not both together.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>backup specification</b> | A list of objects to be backed up, together with a set of devices or drives to be used, backup options for all objects in the specification, and days and time that you want backups to be performed. The objects are entire disks/volumes or parts of them such as files, directories, or                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

|                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                   | even the Windows Registry for example. File selection lists such as include-lists and exclude-lists can be specified.                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>backup system</b>              | (ZDB specific term) A system connected to a disk array together with one or multiple application systems. The backup system is typically connected to a disk array to create target volumes (a replica) and is used for mounting the target volumes (the replica).<br>See also application system, target volume, and replica.                                                                                                                                                                                |
| <b>backup types</b>               | See incremental backup, differential backup, transaction backup, full backup, and delta backup.                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>backup view</b>                | Data Protector provides different views for backup specifications:<br>By Type - according to the type of data available for backups/templates. Default view.<br>By Group - according to the group to which backup specifications/templates belong.<br>By Name - according to the name of backup specifications/templates.<br>By Manager - if you are running MoM, you can also set the Backup view according to the Cell Manager to which backup specifications/templates belong.                             |
| <b>BC</b>                         | (EMC Symmetrix specific term) Business Continuance are processes that allow customers to access and manage instant copies of EMC Symmetrix standard devices.<br>See also BCV.                                                                                                                                                                                                                                                                                                                                 |
| <b>BC Process</b>                 | (EMC Symmetrix specific term) A protected storage environment solution that has defined specially configured EMC Symmetrix devices as mirrors or Business Continuance Volumes to protect data on EMC Symmetrix standard devices.<br>See also BCV.                                                                                                                                                                                                                                                             |
| <b>BCV</b>                        | (EMC Symmetrix specific term) Business Continuance Volumes, or BCV devices, are dedicated SLDs that are pre-configured in the ICDA on which the business continuation operation runs. BCV devices are assigned separate SCSI addresses, differing from the addresses used by the SLDs they mirror. The BCV devices are used as splittable mirrors of the primary EMC Symmetrix SLDs that need to be protected.<br>See also BC and BC Process.                                                                 |
| <b>Boolean operators</b>          | The Boolean operators for the full text search functionality of the online Help system are AND, OR, NOT, and NEAR. Used when searching, they enable you to define your query precisely by creating a relationship between search terms. If no operator is specified in a multi-word search, AND is used by default. For example, the query manual disaster recovery is equivalent to manual AND disaster AND recovery.                                                                                        |
| <b>boot volume/disk/partition</b> | A volume/disk/partition with files required for the initial step of the boot process. Microsoft terminology defines the boot volume/disk/partition as a volume/disk/partition containing the operating system files.                                                                                                                                                                                                                                                                                          |
| <b>BRARCHIVE</b>                  | (SAP R/3 specific term) An SAP R/3 backup tool that allows you to archive redo log files. BRARCHIVE also saves all the logs and profiles of the archiving process.<br>See also BRBACKUP and BRRESTORE.                                                                                                                                                                                                                                                                                                        |
| <b>BRBACKUP</b>                   | (SAP R/3 specific term) An SAP R/3 backup tool that allows an online or offline backup of the control file, of individual data files, or of all tablespaces and, if necessary, of the online redo log files.<br>See also BRARCHIVE and BRRESTORE.                                                                                                                                                                                                                                                             |
| <b>BRRESTORE</b>                  | (SAP R/3 specific term) An SAP R/3 tool that can be used to restore files of the following type: <ul style="list-style-type: none"> <li>• Database data files, control files, and online redo log files saved with BRBACKUP</li> <li>• Redo log files archived with BRARCHIVE</li> <li>• Non-database files saved with BRBACKUP</li> </ul> You can specify files, tablespaces, complete backups, log sequence numbers of redo log files, or the session ID of the backup.<br>See also BRBACKUP and BRARCHIVE. |
| <b>BSM</b>                        | The Data Protector Backup Session Manager controls the backup session. This process always runs on the Cell Manager system.                                                                                                                                                                                                                                                                                                                                                                                   |

## C

|                                                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>CAP</b>                                           | <i>(StorageTek specific term)</i> Cartridge Access Port is a port built into the door panel of a library. The purpose is to enter or eject media.                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>catalog protection</b>                            | Defines how long information about backed up data (such as file names and file versions) is kept in the IDB.<br><i>See also</i> data protection.                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>CDB</b>                                           | The Catalog Database is a part of the IDB that contains information about backup, restore, object copy, object consolidation, object verification, and media management sessions. Depending on the selected logging level, it also contains file names and file versions. This part of the database is always local to the cell.<br><i>See also</i> MMDB.                                                                                                                                                                                     |
| <b>CDF file</b>                                      | <i>(UNIX specific term)</i> A Context Dependent File is a file consisting of several files grouped under the same pathname. The system ordinarily selects one of the files using the context of the process. This mechanism allows machine dependent executables, system data, and device files to work correctly from all hosts in a cluster while using the same pathname.                                                                                                                                                                  |
| <b>cell</b>                                          | A set of systems that are under the control of a Cell Manager. The cell typically represents the systems on a site or an organizational entity, which are connected to the same LAN or SAN. Central control is available to administer the backup and restore policies and tasks.                                                                                                                                                                                                                                                             |
| <b>Cell Manager</b>                                  | The main system in the cell where the essential Data Protector software is installed and from which all backup and restore activities are managed. The GUI used for management tasks can be located on a different system. Each cell has one Cell Manager system.                                                                                                                                                                                                                                                                             |
| <b>centralized licensing</b>                         | Data Protector allows you to configure centralized licensing for the whole enterprise environment consisting of several cells. All Data Protector licenses are installed and kept on the Enterprise Cell Manager system. You can then allocate licenses to specific cells to suit your needs.<br><i>See also</i> MoM.                                                                                                                                                                                                                         |
| <b>Centralized Media Management Database (CMMDB)</b> | <i>See</i> CMMDB.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Certificate Server</b>                            | A Windows Certificate Server can be installed and configured to provide certificates for clients. It provides customizable services for issuing and managing certificates for the enterprise. These services issue, revoke, and manage certificates employed in public key-based cryptography technologies.                                                                                                                                                                                                                                   |
| <b>Change Journal</b>                                | <i>(Windows specific term)</i> A Windows filesystem feature that logs a record of each change as it occurs to the files and directories on a local NTFS volume.                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Change Log Provider</b>                           | <i>(Windows specific term)</i> A module that can be queried to determine which objects on a filesystem have been created, modified, or deleted.                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>channel</b>                                       | <i>(Oracle specific term)</i> An Oracle Recovery Manager resource allocation. Every allocated channel starts a new Oracle process, which performs backup, restore, and recovery actions. The type of channel allocated determines the type of media used: <ul style="list-style-type: none"> <li>• type 'disk'</li> <li>• type 'sbt_tape'</li> </ul> If the specified channel is of type 'sbt_tape' and Oracle is integrated with Data Protector, the server process will attempt to read backups from or write data files to Data Protector. |
| <b>circular logging</b>                              | <i>(Microsoft Exchange Server and Lotus Domino Server specific term)</i> Circular logging is a Microsoft Exchange Server database and Lotus Domino Server database mode, in which the transaction log file contents are periodically overwritten after the corresponding data is committed to the database. Circular logging reduces disk storage space requirements.                                                                                                                                                                         |
| <b>client backup</b>                                 | A backup of all volumes (filesystems) mounted on a Data Protector client. What is actually backed up depends on how you select objects in a backup specification: <ul style="list-style-type: none"> <li>• If you select the check box next to the client system name, a single backup object of the Client System type is created. As a result, at the time of the backup, Data Protector first</li> </ul>                                                                                                                                   |

detects all volumes that are mounted on the selected client and then backs them up. On Windows clients, CONFIGURATION is also backed up.

- If you individually select all volumes that are mounted on the client system, a separate backup object of the Filesystem type is created for each volume. As a result, at the time of the backup, only the selected volumes are backed up. Volumes that have been potentially mounted on the client after the backup specification was created are not backed up.

**client or client system**

Any system configured with any Data Protector functionality and configured in a cell.

**cluster continuous replication**

(Microsoft Exchange Server specific term) Cluster continuous replication (CCR) is a high availability solution that uses cluster management and failover options to create and maintain an exact copy (CCR copy) of a storage group. A storage group is replicated to a separate server. CCR removes any single point of failure in your Exchange back-end servers. You can perform backups using VSS on your passive Exchange Server node where a CCR copy is located and thus reducing the load on the active node.

A CCR copy is used for disaster recovery since you can switch to the CCR copy in a few seconds. A replicated storage group is represented as a new instance of Exchange writer called Exchange Replication Service and can be backed up (using VSS) like an ordinary storage group.

See also Exchange Replication Service and local continuous replication.

**cluster-aware application**

It is an application that supports the cluster Application Programming Interface. Each cluster-aware application declares its own critical resources (disk volumes (on Microsoft Cluster Server), volume groups (on MC/ServiceGuard), application services, IP names and addresses, and so on).

**CMD script for Informix Server**

(Informix Server specific term) A Windows CMD script that is created in INFORMIXDIR when an Informix Server database is configured. The CMD script is a set of system commands that export environment variables for Informix Server.

**CMMDB**

The Data Protector Centralized Media Management Database (CMMDB) is the result of merging MMDBs from several cells in the MoM environment. It allows you to share high-end devices and media across multiple cells in a MoM environment. One cell can control the robotics, serving the devices that are connected to systems in other cells. The CMMDB must reside on the Manager-of-Managers. A reliable network connection between the MoM cell and the other Data Protector cells is highly recommended

See also MoM.

**COM+ Class Registration Database**

(Windows specific term) The COM+ Class Registration Database and the Windows Registry store application attributes, class attributes, and computer-level attributes. This guarantees consistency among these attributes and provides common operation on top of these attributes.

**command device**

(HP P9000 XP Disk Array Family specific term) A dedicated volume in the disk array which acts as the interface between a management application and the disk array's storage system. It cannot be used for data storage and only accepts requests for operations that are then executed by the disk array.

**Command View VLS**

(VLS specific term) A web browser-based GUI that is used to configure, manage, and monitor the VLS through a LAN.

See also Virtual Library System (VLS).

**command-line interface (CLI)**

A set of DOS and UNIX like commands that you can use in shell scripts to perform Data Protector configuration, backup, restore, and management tasks.

**concurrency**

See Disk Agent concurrency.

**container**

(HP P6000 EVA Disk Array Family specific term) Space on a disk array, which is pre-allocated for later use as a standard snapshot, vsnap, or snapclone.

**control file**

(Oracle and SAP R/3 specific term) An Oracle data file that contains entries specifying the physical structure of the database. It provides database consistency information used for recovery.

**copy set**

(HP P6000 EVA Disk Array Family specific term) A pair that consists of the source volumes on a local P6000 EVA and their replica on a remote P6000 EVA.

See also source volume, replica, and HP Continuous Access + Business Copy (CA+BC) P6000 EVA.

**CRS**

The Cell Request Server process (service), which runs on the Data Protector Cell Manager, and starts and controls the backup and restore sessions. The service is started as soon as Data Protector



is installed on the Cell Manager. On Windows systems, the CRS runs under the account of the user specified at installation time. On UNIX systems, it runs under the account `root`.

**CSM** The Data Protector Copy and Consolidation Session Manager process controls the object copy and object consolidation sessions and runs on the Cell Manager system.

## D

**data file** (*Oracle and SAP R/3 specific term*) A physical file created by Oracle that contains data structures such as tables and indexes. A data file can only belong to one Oracle database.

**data protection** Defines how long the backed up data on media remains protected, that is, Data Protector will not overwrite it. When the protection expires, Data Protector will be able to reuse the media in one of the next backup sessions.  
See also catalog protection.

**data replication (DR) group** (*HP P6000 EVA Disk Array Family specific term*) A logical grouping of HP P6000 EVA Disk Array Family virtual disks. It can contain up to eight copy sets provided they have common characteristics and share a common HP CA P6000 EVA log.  
See also copy set.

**data stream** Sequence of data transferred over the communication channel.

**Data\_Protector\_home** A reference to the directory containing Data Protector program files (on Windows Vista, Windows 7, and Windows Server 2008) or the directory containing Data Protector program files and data files (on other Windows operating systems). Its default path is `%ProgramFiles%\OmniBack`, but the path can be changed in the Data Protector Setup Wizard at installation time.  
See also `Data_Protector_program_data`.

**Data\_Protector\_program\_data** A reference to the directory containing Data Protector data files on Windows Vista, Windows 7, and Windows Server 2008. Its default path is `%ProgramData%\OmniBack`, but the path can be changed in the Data Protector Setup Wizard at installation time.  
See also `Data_Protector_home`.

**database library** A Data Protector set of routines that enables data transfer between Data Protector and a server of an online database integration, for example, Oracle Server.

**database parallelism** More than one database is backed up at a time if the number of available devices allows you to perform backups in parallel.

**database server** A computer with a large database stored on it, such as the SAP R/3 or Microsoft SQL database. A server has a database that can be accessed by clients.

**Dbobject** (*Informix Server specific term*) An Informix Server physical database object. It can be a blob space, db space, or logical log file.

**DC directory** The Detail Catalog (DC) directory contains DC binary files, which store information about file versions. It represents the DCBF part of the IDB, which occupies approximately 80% of the IDB. The default DC directory is called the `dcbf` directory and is located on the Cell Manager in the directory `Data_Protector_program_data\db40` (Windows Server 2008), `Data_Protector_home\db40` (other Windows systems), or `/var/opt/omni/server/db40` (UNIX systems). You can create more DC directories and use a custom location. Up to 50 DC directories are supported per cell. The default maximum size of a DC directory is 16 GB.

**DCBF** The Detail Catalog Binary Files (DCBF) part of the IDB stores information about file versions and attributes. It occupies approximately 80% of the IDB. One DC binary file is created for each Data Protector medium used for backup. Its maximum size is limited by the filesystem settings.

**delta backup** A delta backup is a backup containing all the changes made to the database from the last backup of any type.  
See also backup types.

**device** A physical unit which contains either just a drive or a more complex unit such as a library.

**device chain** A device chain consists of several standalone devices configured for sequential use. When a medium in one device gets full, the backup automatically continues on a medium in the next device in the device chain.

**device group** (*EMC Symmetrix specific term*) A logical unit representing several EMC Symmetrix devices. A device cannot belong to more than a single device group. All devices in a device group must be

|                                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                           | on the same EMC Symmetrix unit. You can use a device group to identify and work with a subset of the available EMC Symmetrix devices.                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>device streaming</b>                   | A device is streaming if it can feed enough data to the medium to keep it moving forward continuously. Otherwise, the tape has to be stopped, the device waits for more data, reverses the tape a little and resumes to write to the tape, and so on. In other words, if the data rate written to the tape is less or equal the data rate which can be delivered to the device by the computer system, then the device is streaming. Streaming significantly improves the performance of the device and use of space.            |
| <b>DHCP server</b>                        | A system running the Dynamic Host Configuration Protocol (DHCP) providing dynamic IP address assignment and network configuration for DHCP clients.                                                                                                                                                                                                                                                                                                                                                                              |
| <b>differential backup</b>                | An incremental backup that backs up changes made since the last full backup. To perform this type of backup, specify the Incr1 backup type.<br><i>See also</i> incremental backup.                                                                                                                                                                                                                                                                                                                                               |
| <b>differential backup</b>                | <i>(Microsoft SQL Server specific term)</i> A database backup that records only the data changes made to the database after the last full database backup.<br><i>See also</i> backup types.                                                                                                                                                                                                                                                                                                                                      |
| <b>differential database backup</b>       | A differential database backup records only those data changes made to the database after the last full database backup.                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>directory junction</b>                 | <i>(Windows specific term)</i> Directory junctions use the reparse point concept of Windows. An NTFS 5 directory junction allows you to redirect a directory/file request to another location.                                                                                                                                                                                                                                                                                                                                   |
| <b>disaster recovery</b>                  | A process to restore a client's main system disk to a state close to the time when a (full) backup was performed.                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>disaster recovery operating system</b> | <i>See</i> DR OS.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Disk Agent</b>                         | A component needed on a client to back it up and restore it. The Disk Agent controls reading from and writing to a disk. During a backup session, the Disk Agent reads data from a disk and sends it to the Media Agent, which then moves it to the device. During a restore session the Disk Agent receives data from the Media Agent and writes it to the disk. During an object verification session the Disk Agent receives data from the Media Agent and performs the verification process, but no data is written to disk. |
| <b>Disk Agent concurrency</b>             | The number of Disk Agents that are allowed to send data to one Media Agent concurrently.                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>disk group</b>                         | <i>(Veritas Volume Manager specific term)</i> The basic unit of data storage in VxVM system. A disk group can consist of one or more physical volumes. There can be more than one disk group on the system.                                                                                                                                                                                                                                                                                                                      |
| <b>disk image (rawdisk) backup</b>        | A high-speed backup where Data Protector backs up files as bitmap images. A disk image (rawdisk) backup does not track the files and directory structure stored on the disk, but stores a disk image structure on byte level. You can perform a disk image backup of either specific disk sections or a complete disk.                                                                                                                                                                                                           |
| <b>disk quota</b>                         | A concept to manage disk space consumption for all or a subset of users on a computer system. This concept is used by several operating system platforms.                                                                                                                                                                                                                                                                                                                                                                        |
| <b>disk staging</b>                       | The process of backing up data in several phases to improve the performance of backups and restores, reduce costs of storing the backed up data, and increase the data availability and accessibility for restore. The backup stages consist of backing up data to one media type first (for example disk) and later copying it to a different media type (for example tape).                                                                                                                                                    |
| <b>distributed file media format</b>      | A media format, available with the file library, which supports a space efficient type of synthetic backup called virtual full backup. Using this format is a prerequisite for virtual full backup.<br><i>See also</i> virtual full backup.                                                                                                                                                                                                                                                                                      |
| <b>Distributed File System (DFS)</b>      | A service that connects file shares into a single namespace. The file shares can reside on the same or on different computers. DFS provides client access to the resources in a location-transparent manner.                                                                                                                                                                                                                                                                                                                     |
| <b>DMZ</b>                                | The Demilitarized Zone (DMZ) is a network inserted as a "neutral zone" between a company's private network (intranet) and the outside public network (Internet). It prevents outside users from getting direct access to company servers in the intranet.                                                                                                                                                                                                                                                                        |

|                               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>DNS server</b>             | In the DNS client-server model, this is the server containing information about a portion of the DNS database that makes computer names available to client resolvers querying for name resolution across the Internet.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>domain controller</b>      | A server in a network that is responsible for user security and verifying passwords within a group of other servers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>DR image</b>               | Data required for temporary disaster recovery operating system (DR OS) installation and configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>DR OS</b>                  | An operating system environment in which disaster recovery runs. It provides Data Protector with a basic runtime environment (disk, network, tape, and filesystem access). It has to be installed on disk or loaded into memory and configured before the Data Protector disaster recovery can be performed. DR OS can be either temporary or active. A temporary DR OS is used exclusively as a host environment for some other operating system restore along with the target operating system configuration data. It is deleted after the target system is restored to the original system configuration. An active DR OS not only hosts the Data Protector disaster recovery process but can also be a part of the restored system because it replaces its own configuration data with the original configuration data. |
| <b>drive</b>                  | A physical unit that receives data from a computer system and can write it onto a magnetic medium (typically a tape drive). It can also read the data from the medium and send it to the computer system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>drive index</b>            | A number that identifies the mechanical position of a drive inside a library device. This number is used by the robotic control to access a drive.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>drive-based encryption</b> | Data Protector drive-based encryption uses the encryption functionality of the drive. While performing the backup, the drive encrypts both the data and the meta-data that is written to the medium.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## E

|                                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|----------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>EMC Symmetrix Agent</b>             | A Data Protector software module that prepares the EMC Symmetrix environment for backup and restore operations.                                                                                                                                                                                                                                                                                                                                                   |
| <b>emergency boot file</b>             | <i>(Informix Server specific term)</i> The Informix Server configuration file <code>ixbar.server_id</code> that resides in the directory <code>INFORMIXDIR/etc</code> (on Windows) or <code>INFORMIXDIR\etc</code> (on UNIX). <code>INFORMIXDIR</code> is the Informix Server home directory and <code>server_id</code> is the value of the <code>SERVENUM</code> configuration parameter. Each line of the emergency boot file corresponds to one backup object. |
| <b>encrypted control communication</b> | Data Protector secure communication between the clients in the Data Protector cell is based on Secure Socket Layer (SSL) that uses SSLv3 algorithms to encrypt control communication. Control communication in a Data Protector cell is all communication between Data Protector processes, except the data transfer from Disk Agent (and Integrations) to Media Agent, and the other way round.                                                                  |
| <b>encryption key</b>                  | A 256-bit randomly generated number used by the Data Protector encryption algorithm to encode information during backups for which AES 256-bit software encryption or drive-based encryption has been specified. The same key is used for subsequent decryption of the information. Encryption keys for a Data Protector cell are stored in a central keystore on the Cell Manager.                                                                               |
| <b>encryption KeyID-StoreID</b>        | Combined identifier used by the Data Protector Key Management Server to identify and administer encryption keys used by Data Protector. <code>KeyID</code> identifies the key within the keystore. <code>StoreID</code> identifies the keystore on the Cell Manager. If Data Protector has been upgraded from an earlier version with encryption functionality, there may several <code>StoreIDs</code> used on the same Cell Manager.                            |
| <b>enhanced incremental backup</b>     | Conventional incremental backup backs up files that have changed since a previous backup, but has certain limitations in detection of changes. Unlike conventional incremental backup, enhanced incremental backup reliably detects and backs up also renamed and moved files, as well as files with changes in attributes.                                                                                                                                       |
| <b>enterprise backup environment</b>   | Several cells can be grouped together and managed from a central cell. The enterprise backup environment includes all clients located in several Data Protector cells which are managed and administered from a central cell using the Manager-of-Managers concept.<br><i>See also MoM.</i>                                                                                                                                                                       |

|                                             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Event Log (Data Protector Event Log)</b> | A central repository of all Data Protector-related notifications. By default, all notifications are sent to the Event Log. The events are logged on the Cell Manager into the file <code>Data_Protector_program_data\log\server\Ob2EventLog.txt</code> (Windows Server 2008), <code>Data_Protector_home\log\server\Ob2EventLog.txt</code> (other Windows systems), or <code>/var/opt/omni/server/log/Ob2EventLog.txt</code> (UNIX systems). The Event Log is accessible only to users of the Data Protector Admin user group and to users who are granted the Data Protector Reporting and notifications user rights. You can view or delete all events in the Event Log. |
| <b>Event Logs</b>                           | ( <i>Windows specific term</i> ) Files in which Windows logs all events, such as the starting or stopping of services and the logging on and off of users. Data Protector can back up Windows Event Logs as part of the Windows configuration backup.                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Exchange Replication Service</b>         | ( <i>Microsoft Exchange Server specific term</i> ) The Microsoft Exchange Server service that represents storage groups that were replicated using either local continuous replication (LCR) or cluster continuous replication (CCR) technology.<br>See also cluster continuous replication and local continuous replication.                                                                                                                                                                                                                                                                                                                                             |
| <b>exchanger</b>                            | Also referred to as SCSI Exchanger.<br>See also library.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>exporting media</b>                      | A process that removes all data about backup sessions, such as systems, objects, and file names, which reside on the media from the IDB. Information about the media and their relation to a pool is also removed from the IDB. The data on the media remains unchanged.<br>See also importing media.                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Extensible Storage Engine (ESE)</b>      | ( <i>Microsoft Exchange Server specific term</i> ) A database technology used as a storage system for information exchange in Microsoft Exchange Server.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>F</b>                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>failover</b>                             | Transferring of the most important cluster data, called group (on Windows) or package (on UNIX) from one cluster node to another. A failover can occur mostly because of software or hardware failures or maintenance on the primary node.                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>failover</b>                             | ( <i>HP P6000 EVA Disk Array Family specific term</i> ) An operation that reverses the roles of source and destination in HP Continuous Access + Business Copy (CA+BC) P6000 EVA configurations.<br>See also HP Continuous Access + Business Copy (CA+BC) P6000 EVA.                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>FC bridge</b>                            | See Fibre Channel bridge.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Fibre Channel</b>                        | An ANSI standard for high-speed computer interconnection. Using either optical or copper cables, it allows the high speed bi-directional transmission of large data files and can be deployed between sites kilometers apart. Fibre Channel connects nodes using three physical topologies: point-to-point, loop, and switched.                                                                                                                                                                                                                                                                                                                                           |
| <b>Fibre Channel bridge</b>                 | A Fibre Channel bridge or multiplexer provides the ability to migrate existing parallel SCSI devices, like RAID arrays, solid state disks (SSD), and tape libraries to a Fibre Channel environment. On one side of the bridge or multiplexer there is a Fibre Channel interface while on the other side there are parallel SCSI ports. The bridge or multiplexer enables SCSI packets to be moved between the Fibre Channel and parallel SCSI devices.                                                                                                                                                                                                                    |
| <b>file depot</b>                           | A file containing the data from a backup to a file library device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>file jukebox device</b>                  | A device residing on disk consisting of multiple slots used to store file media.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>file library device</b>                  | A device which resides on a disk emulating a library with several media, hence containing multiple files, referred to as file depots.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>File Replication Service (FRS)</b>       | A Windows service that replicates the domain controller store logon scripts and group policies. FRS also enables replication of Distributed File System (DFS) shares between systems and allows any server to perform replication activity.                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>file tree walk</b>                       | ( <i>Windows specific term</i> ) The process of traversing a filesystem to determine which objects have been created, modified, or deleted.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>file version</b>                         | The same file can be backed up multiple times in case of full backups and incremental backups (if the file changed). If the log level ALL is selected for backup, Data Protector retains one entry in the IDB for the filename itself and one for each version (date/time) of the file.                                                                                                                                                                                                                                                                                                                                                                                   |

|                             |                                                                                                                                                                                                                                                                                                                                                                                          |
|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>filesystem</b>           | The organization of files on a hard disk. A filesystem is backed up so that the file attributes and the file contents are stored on the backup media.                                                                                                                                                                                                                                    |
| <b>first-level mirror</b>   | <i>(HP P9000 XP Disk Array Family specific term)</i> A mirror of an internal disk (LDEV) of a disk array of the HP P9000 XP Disk Array Family which can be further mirrored itself, producing second-level mirrors. For Data Protector zero downtime backup and instant recovery purposes, only first-level mirrors can be used.<br>See also primary volume and mirror unit (MU) number. |
| <b>flash recovery area</b>  | <i>(Oracle specific term)</i> A directory, filesystem, or Automatic Storage Management (ASM) disk group managed by Oracle that serves as a centralized storage area for files related to backup, restore, and database recovery (recovery files).<br>See also recovery files.                                                                                                            |
| <b>fnames.dat</b>           | The <code>fnames.dat</code> files of the IDB contain information on the names of the backed up files. Typically, these files occupy about 20% of the IDB, if filenames are stored.                                                                                                                                                                                                       |
| <b>formatting</b>           | A process that erases any data contained on a medium and prepares it for use with Data Protector. Information about media (medium ID, description, and location) is saved in the IDB as well as on the respective media (media header). Data Protector media with protected data are not formatted until the protection expires or the media are unprotected/recycled.                   |
| <b>free pool</b>            | An auxiliary source of media for use by media pools when they run out of media. The media pools must be configured to use free pools.                                                                                                                                                                                                                                                    |
| <b>full backup</b>          | A backup in which all selected objects are backed up, whether or not they have been recently modified.<br>See also backup types.                                                                                                                                                                                                                                                         |
| <b>full database backup</b> | A backup of all data in a database, not only the data that has been changed after the last (full or incremental) backup of the database. A full database backup does not depend on any other backup.                                                                                                                                                                                     |
| <b>full mailbox backup</b>  | A full mailbox backup is a backup of the entire mailbox content.                                                                                                                                                                                                                                                                                                                         |
| <b>full ZDB</b>             | A ZDB-to-tape or ZDB-to-disk+tape session in which all selected objects are streamed to tape, even if there are no changes from the previous backup.<br>See also incremental ZDB.                                                                                                                                                                                                        |

## G

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>global options file</b> | A file that allows you to customize Data Protector. It explains the global options, which cover various aspects of Data Protector, typically time-outs and limits, and affect the entire Data Protector cell. The file is located on the Cell Manager in the directory<br><code>Data_Protector_program_data\Config\Server\Options</code> (Windows Server 2008),<br><code>Data_Protector_home\Config\Server\Options</code> (other Windows systems), or<br><code>/etc/opt/omni/server/options</code> (HP-UX, Solaris, and Linux systems). |
| <b>group</b>               | <i>(Microsoft Cluster Server specific term)</i> A collection of resources (for example disk volumes, application services, IP names, and addresses) that are needed to run a specific cluster-aware applications.                                                                                                                                                                                                                                                                                                                       |
| <b>GUI</b>                 | A graphical user interface provided by Data Protector for easy access to all configuration, administration, and operation tasks. Besides the original Data Protector GUI that runs on Windows, Data Protector also provides a Java-based graphical user interface with the same look and feel, which runs on numerous platforms.                                                                                                                                                                                                        |

## H

|                      |                                                                                                                                                                                       |
|----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>hard recovery</b> | <i>(Microsoft Exchange Server specific term)</i> A Microsoft Exchange Server database recovery that is performed after a restore by the database engine, using transaction log files. |
| <b>heartbeat</b>     | A cluster data set with a time stamp carrying information about the operational status of a particular cluster node. This data set or packet is distributed among all cluster nodes.  |

|                                                               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Hierarchical Storage Management (HSM)</b>                  | A method for optimizing the use of expensive hard disk storage by migrating less frequently used data to less expensive optical platters. When needed, the data is migrated back to hard disk storage. This balances the need for fast retrieval from hard disk with the lower cost of optical platters.                                                                                                                                                                                                                                                                                                                                                                |
| <b>Holidays file</b>                                          | A file that contains information about holidays. You can set different holidays by editing the Holidays file on the Cell Manager in the directory<br><i>Data_Protector_program_data\Config\Server\holidays</i> (Windows Server 2008),<br><i>Data_Protector_home\Config\Server\holidays</i> (other Windows systems), or<br><i>/etc/opt/omni/server/Holidays</i> (UNIX systems).                                                                                                                                                                                                                                                                                          |
| <b>hosting system</b>                                         | A working Data Protector client used for Disk Delivery Disaster Recovery with a Data Protector Disk Agent installed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>HP Business Copy (BC) P6000 EVA</b>                        | <i>(HP P6000 EVA Disk Array Family specific term)</i> A local replication software solution that enables creation of point-in-time copies (replicas) of the source volumes using the snapshot and clone capabilities of the P6000 EVA firmware.<br>See also replica, source volume, snapshot, and HP Continuous Access + Business Copy (CA+BC) P6000 EVA.                                                                                                                                                                                                                                                                                                               |
| <b>HP Business Copy (BC) P9000 XP</b>                         | <i>(HP P9000 XP Disk Array Family specific term)</i> An HP P9000 XP Disk Array Family configuration that enables creation and maintenance of internal copies of LDEVs for various purposes, such as data duplication and backup. The copies (secondary volumes or S-VOLs) can be separated from the primary volumes (P-VOLs) and connected to a different system. For Data Protector zero downtime backup purposes, P-VOLs should be available to the application system, and one of the S-VOL sets should be available to the backup system.<br>See also LDEV, HP Continuous Access (CA) P9000 XP, Main Control Unit, application system, and backup system.           |
| <b>HP Command View (CV) EVA</b>                               | <i>(HP P6000 EVA Disk Array Family specific term)</i> The user interface that enables you to configure, manage, and monitor your P6000 EVA storage system. It is used to perform various storage management tasks, for example, creating virtual disk families, managing storage system hardware, and creating snapshots, snapclones, and mirrorclones of virtual disks. The HP Command View EVA software runs on the HP Storage Management Appliance, and is accessed by a Web browser.<br>See also HP StorageWorks P6000 EVA SMI-S Agent and HP StorageWorks SMI-S P6000 EVA Array provider.                                                                          |
| <b>HP Continuous Access (CA) P9000 XP</b>                     | <i>(HP P9000 XP Disk Array Family specific term)</i> An HP P9000 XP Disk Array Family configuration that enables creation and maintenance of remote copies of LDEVs for purposes such as data duplication, backup, and disaster recovery. HP CA P9000 XP operations involve main (primary) disk array units and remote (secondary) disk array units. The main disk array units are connected to the application system and contain primary volumes (P-VOLs), which store original data. The remote disk array units are connected to the backup system and contain secondary volumes (S-VOLs).<br>See also HP Business Copy (BC) P9000 XP, Main Control Unit, and LDEV. |
| <b>HP Continuous Access + Business Copy (CA+BC) P6000 EVA</b> | <i>(HP P6000 EVA Disk Array Family specific term)</i> An HP P6000 EVA Disk Array Family configuration that enables creation and maintenance of copies (replicas) of the source volumes on a remote P6000 EVA, and later use of these copies as the source for local replication on this remote array.<br>See also HP Business Copy (BC) P6000 EVA, replica, and source volume.                                                                                                                                                                                                                                                                                          |
| <b>HP SMI-S P6000 EVA Array provider</b>                      | An interface used for controlling HP P6000 EVA Disk Array Family. SMI-S P6000 EVA Array provider runs as a separate service on the HP Storage Management Appliance system and acts as a gateway between incoming requests and HP Command View EVA. With the Data Protector HP P6000 EVA Disk Array Family integration, SMI-S P6000 EVA Array provider accepts standardized requests from the P6000 EVA SMI-S Agent, communicates with HP Command View EVA for information or method invocation, and returns standardized responses.<br>See also HP StorageWorks P6000 EVA SMI-S Agent and HP Command View (CV) EVA.                                                     |
| <b>HP StorageWorks P6000 EVA SMI-S Agent</b>                  | A Data Protector software module that executes all tasks required for the HP P6000 EVA Disk Array Family integration. With the P6000 EVA SMI-S Agent, the control over the array is established through HP SMI-S P6000 EVA Array provider, which directs communication between incoming requests and HP CV EVA.                                                                                                                                                                                                                                                                                                                                                         |

|                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                                  | See also HP Command View (CV) EVA and HP SMI-S P6000 EVA Array provider.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>HP StorageWorks P9000 XP Agent</b>            | A Data Protector component that executes all tasks needed by the Data Protector HP P9000 XP Disk Array Family integration. It uses RAID Manager Library for communication with a P9000 XP Array storage system.<br>See also RAID Manager Library.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>HP Operations Manager</b>                     | HP Operations Manager provides powerful capabilities for operations management of a large number of systems and applications in a network. Data Protector provides an integration into this management product. This integration is implemented as a SMART Plug-In for HP Operations Manager management servers on Windows, HP-UX, Solaris, and Linux. Earlier versions of HP Operations Manager were called IT/Operations, Operations Center, Vantage Point Operations, and OpenView Operations.                                                                                                                                                                                                                                                                                             |
| <b>HP Operations Manager SMART Plug-In (SPI)</b> | A fully integrated, out-of-the-box solution which "plugs into" HP Operations Manager, extending the managed domain. Through the Data Protector integration, which is implemented as an HP Operations Manager SMART Plug-In, a user can have an arbitrary number of Data Protector Cell Managers monitored as an extension to HP Operations Manager.                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>ICDA</b>                                      | (EMC Symmetrix specific term) EMC's Symmetrix Integrated Cached Disk Arrays (ICDA) is a disk array device that combines a set of physical disks, a number of FWD SCSI channels, an internal cache memory, and control and diagnostic software commonly referred to as the microcode.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>IDB</b>                                       | The Data Protector Internal Database. IDB is an embedded database located on the Cell Manager and keeps information regarding which data was backed up, to which media it was backed up, how backup, restore, and other sessions were run, which devices, libraries, and disk arrays are configured, and so on.                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>IDB recovery file</b>                         | An IDB file (obrindex.dat) with information about IDB backups, media, and devices used for the backup. This information can significantly simplify IDB recovery. It is recommended to relocate the file, together with IDB transaction logs, to a separate physical disk from other IDB directories, and, additionally, to make an additional copy of the file.                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>importing media</b>                           | A process that re-reads all data about backup sessions which are on the medium back into the IDB. This then allows for fast and convenient access to the data on the media.<br>See also exporting media.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>incremental (re)-establish</b>                | (EMC Symmetrix specific term) A BCV or SRDF control operation. In BCV control operations, an incremental establish causes the BCV device to be synchronized incrementally and to function as an EMC Symmetrix mirrored medium. The EMC Symmetrix devices must have been previously paired. In SRDF control operations, an incremental establish causes the target (R2) device to be synchronized incrementally and to function as an EMC Symmetrix mirrored medium. The EMC Symmetrix devices must have been previously paired.                                                                                                                                                                                                                                                               |
| <b>incremental backup</b>                        | A backup that selects only files that have changed since a previous backup. Several levels of incremental backup are available, which enables detailed control of restore chain length.<br>See also backup types.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>incremental backup</b>                        | (Microsoft Exchange Server specific term) A backup of the Microsoft Exchange Server data that has changed since the last full or incremental backup. With the incremental backup, only the transaction log files are backed up.<br>See also backup types.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>incremental mailbox backup</b>                | An incremental mailbox backup backs up all the changes made to the mailbox after the last backup of any type.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>incremental restore</b>                       | (EMC Symmetrix specific term) A BCV or SRDF control operation. In BCV control operations, an incremental restore reassigns a BCV device as the next available mirror of the standard device in the pair. However, the standard devices are updated with only the data that was written to the BCV device during the time of the original pair split, and the data that was written to the standard device during the split is overwritten with data from the BCV mirror. In SRDF control operations, an incremental restore reassigns a target (R2) device as the next available mirror of the source (R1) device in the pair. However, the source (R1) devices are updated with only the data that was written to the target (R2) device during the time of the original pair split, and the |



|                                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                            | data that was written to the source (R1) device during the split is overwritten with data from the target (R2) mirror.                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>incremental ZDB</b>                     | A filesystem ZDB-to-tape or ZDB-to-disk+tape session in which only changes from the last protected full or incremental backup are streamed to tape.<br>See also full ZDB.                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>incremental1 mailbox backup</b>         | An incremental1 mailbox backup backs up all the changes made to the mailbox after the last full backup.                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Inet</b>                                | A process that runs on each UNIX system or service that runs on each Windows system in the Data Protector cell. It is responsible for communication between systems in the cell and for starting other processes needed for backup and restore. The Inet service is started as soon as Data Protector is installed on a system. The Inet process is started by the inetd daemon.                                                                                                                                                                                       |
| <b>Information Store</b>                   | <i>(Microsoft Exchange Server specific term)</i> The Microsoft Exchange Server service that is responsible for storage management. Information Store in Microsoft Exchange Server manages two kinds of stores: mailbox stores and public folder stores. A mailbox store consists of mailboxes that belong to individual users. A public folder store contains public folders and messages that are shared among several users.<br>See also Key Management Service and Site Replication Service.                                                                        |
| <b>Informix Server</b>                     | <i>(Informix Server specific term)</i> Refers to Informix Dynamic Server.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>initializing</b>                        | See formatting.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Installation Server</b>                 | A computer system that holds a repository of the Data Protector installation packages for a specific architecture. The Installation Server is used for remote installation of Data Protector clients. In mixed environments at least two Installation Servers are needed: one for UNIX systems and one for Windows systems.                                                                                                                                                                                                                                            |
| <b>instant recovery</b>                    | <i>(ZDB specific term)</i> A process in which a replica, produced by a ZDB-to-disk or a ZDB-to-disk+tape session, is used to restore the contents of the source volumes to their states at the time at which the replica was created, avoiding the need to perform a restore from tape. Depending on the application or database concerned, this may be all that is required, or other steps, such as the application of transaction log files, may be required for full recovery.<br>See also replica, zero downtime backup (ZDB), ZDB to disk, and ZDB to disk+tape. |
| <b>integration object</b>                  | A backup object of a Data Protector integration, such as Oracle or SAP DB.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Internet Information Services (IIS)</b> | <i>(Windows specific term)</i> Microsoft Internet Information Services is a network file and application server that supports multiple protocols. Primarily, IIS transmits information in Hypertext Markup Language (HTML) pages by using the Hypertext Transport Protocol (HTTP).                                                                                                                                                                                                                                                                                     |
| <b>ISQL</b>                                | <i>(Sybase specific term)</i> A Sybase utility used to perform system administration tasks on Sybase SQL Server.                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>J</b>                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Java GUI Client</b>                     | The Java GUI Client is a component of the Java GUI that contains only user interface related functionalities (the Cell Manager graphical user interface and the Manager-of-Managers (MoM) graphical user interface) and requires connection to the Java GUI Server to function.                                                                                                                                                                                                                                                                                        |
| <b>Java GUI Server</b>                     | The Java GUI Server is a component of the Java GUI that is installed on the Data Protector Cell Manager system. The Java GUI Server receives requests from the Java GUI Client, processes them and then sends the responses back to the Java GUI Client. The communication is done through Hypertext Transfer Protocol (HTTP) on port 5556.                                                                                                                                                                                                                            |
| <b>jukebox</b>                             | See library.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>jukebox device</b>                      | A device consisting of multiple slots used to store either optical or file media. When being used to store file media, the jukebox device is known as the "file jukebox device".                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>K</b>                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Key Management Service</b>              | <i>(Microsoft Exchange Server specific term)</i> The Microsoft Exchange Server service that provides encryption functionality for enhanced security.<br>See also Information Store and Site Replication Service.                                                                                                                                                                                                                                                                                                                                                       |



|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>keychain</b>                                     | A tool that eliminates the supply of a passphrase manually when decrypting the private key. It needs to be installed and configured on the Installation Server if you perform remote installation using secure shell.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>keystore</b>                                     | All encryption keys are centrally stored in the keystore on the Cell Manager and administered by the Key Management Server (KMS).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>KMS</b>                                          | Key Management Server (KMS) is a centralized service that runs on the Cell Manager and provides key management for the Data Protector encryption functionality. The service is started as soon as Data Protector is installed on the Cell Manager.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>L</b>                                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>LBO</b>                                          | <i>(EMC Symmetrix specific term)</i> A Logical Backup Object (LBO) is an object of data storage/retrieval in the EMC Symmetrix environment. It is stored/retrieved by EMC Symmetrix as one entity and can only be restored as a whole.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>LDEV</b>                                         | <i>(HP P9000 XP Disk Array Family specific term)</i> A logical partition of a physical disk of a disk array of the HP P9000 XP Disk Array Family. An LDEV is the entity that can be replicated using the split-mirror or snapshot functionality of such disk array.<br><i>See also</i> HP Business Copy (BC) P9000 XP, HP Continuous Access (CA) P9000 XP, and replica.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>library</b>                                      | Also called autochanger, jukebox, autoloader, or exchanger. A library contains media in repository slots. Each slot holds one medium (for example, DDS/DAT). Media are moved between slots and drives by a robotic mechanism, allowing random access to media. The library can contain multiple drives.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>lights-out operation or unattended operation</b> | A backup or restore operation that takes place outside of normal business hours without an operator. This implies that no operator personnel is present to work with the backup application or service mount requests, for example.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>LISTENER.ORA</b>                                 | <i>(Oracle specific term)</i> An Oracle configuration file that describes one or more Transparent Network Substrate (TNS) listeners on a server.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>load balancing</b>                               | By default, Data Protector automatically balances the usage of devices selected for backup, so that they are used evenly. Load balancing optimizes the device usage by balancing the number of objects written to each device. Since load balancing is done automatically during backup time, you do not need to manage how the data is actually backed up. You just specify the devices to be used. If you do not want to use load balancing, you can select which device will be used for each object in the backup specification. Data Protector will access the devices in the specified order.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>local and remote recovery</b>                    | Remote recovery is performed if all Media Agent hosts specified in the SRD file are accessible. If any of them fails, the disaster recovery process fails over to the local mode. This means that the target system is searched for locally attached devices. If only one device is found, it is automatically used. Otherwise, Data Protector prompts you to select the device, which will be used for restore.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>local continuous replication</b>                 | <i>(Microsoft Exchange Server specific term)</i> Local continuous replication (LCR) is a single-server solution that creates and maintains an exact copy (LCR copy) of a storage group. An LCR copy is located on the same server as the original storage group. When an LCR copy is created, it is kept up to date through change propagation (log replay) technology. The replication feature in LCR guarantees that logs that have not been replicated are not deleted. The implication of this behavior is that running backups in a mode that deletes logs may not actually free space if replication is sufficiently far behind in its log copying.<br><br>An LCR copy is used for disaster recovery because you can switch to the LCR copy in a few seconds. If an LCR copy is used for backup and if it is located on a different disk than the original data, then the I/O load on a production database is minimal.<br><br>A replicated storage group is represented as a new instance of Exchange writer called Exchange Replication Service and can be backed up (using VSS) as a normal storage group.<br><i>See also</i> cluster continuous replication and Exchange Replication Service. |
| <b>lock name</b>                                    | You can configure the same physical device several times with different characteristics, by using different device names. The lock name is a user specified string that is used for locking all such                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

device configurations to prevent collision if several such devices (device names) are used concurrently. Use an identical lock name for all device definitions which use the same physical device.

|                                                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-----------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>log_full shell script</b>                              | <i>(Informix Server UNIX specific term)</i> A script provided by ON-Bar that you can use to start backing up logical log files when Informix Server issues a logfull event alarm. The Informix Server ALARMPROGRAM configuration parameter defaults to the <code>INFORMIXDIR/etc/log_full.sh</code> , where <code>INFORMIXDIR</code> is the Informix Server home directory. If you do not want logical logs to be backed up continuously, set the ALARMPROGRAM configuration parameter to <code>INFORMIXDIR/etc/no_log.sh</code> .                                                                                                                                                                                                                                                      |
| <b>logging level</b>                                      | The logging level determines the amount of details on files and directories written to the IDB during backup, object copying, or object consolidation. You can always restore your data, regardless of the logging level used during backup. Data Protector provides four logging levels: Log All, Log Directories, Log Files, and No Log. The different logging level settings influence the IDB growth, backup speed, and the convenience of browsing data for restore.                                                                                                                                                                                                                                                                                                               |
| <b>logical-log files</b>                                  | This applies to online database backup. Logical-log files are files in which modified data is first stored before being flushed to disk. In the event of a failure, these logical-log files are used to roll forward all transactions that have been committed as well as roll back any transactions that have not been committed.                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>login ID</b>                                           | <i>(Microsoft SQL Server specific term)</i> The name a user uses to log on to Microsoft SQL Server. A login ID is valid if Microsoft SQL Server has an entry for that user in the system table syslogin.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>login information to the Oracle Target Database</b>    | <i>(Oracle and SAP R/3 specific term)</i> The format of the login information is <code>user_name/password@service</code> , where: <ul style="list-style-type: none"><li>• <code>user_name</code> is the name by which a user is known to Oracle Server and to other users. Every user name is associated with a password and both have to be entered to connect to an Oracle Target Database. This user must have Oracle SYSDBA or SYSOPER rights.</li><li>• <code>password</code> must be the same as the password specified in the Oracle password file (<code>orapwd</code>), which is used for authentication of users performing database administration.</li><li>• <code>service</code> is the name used to identify an SQL*Net server process for the target database.</li></ul> |
| <b>login information to the Recovery Catalog Database</b> | <i>(Oracle specific term)</i> The format of the login information to the Recovery (Oracle) Catalog Database is <code>user_name/password@service</code> , where the description of the user name, password, and service name is the same as in the Oracle SQL*Net V2 login information to the Oracle target database. In this case, <code>service</code> is the name of the service to the Recovery Catalog Database, not the Oracle target database.<br><br>Note that the Oracle user specified here must be the owner of the Oracle Recovery Catalog.                                                                                                                                                                                                                                  |
| <b>Lotus C API</b>                                        | <i>(Lotus Domino Server specific term)</i> An interface for the exchange of backup and recovery information between Lotus Domino Server and a backup solution, like Data Protector.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>LVM</b>                                                | A Logical Volume Manager is a subsystem for structuring and mapping physical disk space to logical volumes on UNIX systems. An LVM system consists of several volume groups, where each volume group has several volumes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## M

|                                |                                                                                                                                                                                                                                                                                                                    |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Magic Packet</b>            | See Wake ONLAN.                                                                                                                                                                                                                                                                                                    |
| <b>mailbox</b>                 | <i>(Microsoft Exchange Server specific term)</i> The location to which e-mail is delivered, which is set up by the administrator for each user. If a set of personal folders is designated as the e-mail delivery location, e-mail is routed from the mailbox to this location.                                    |
| <b>mailbox store</b>           | <i>(Microsoft Exchange Server specific term)</i> A part of the Information Store that maintains information in user mailboxes. A mailbox store consists of a binary rich-text <code>.edb</code> file and a streaming native internet content <code>.stm</code> file.                                               |
| <b>Main Control Unit (MCU)</b> | <i>(HP P9000 XP Disk Array Family specific term)</i> An HP P9000 XP Disk Array Family unit that contains primary volumes (P-VOLs) for the HP CA P9000 XP or HP CA+BC P9000 XP configuration and acts as a master device.<br>See also HP Business Copy (BC) P9000 XP, HP Continuous Access (CA) P9000 XP, and LDEV. |

|                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>make_net_recovery</b>         | <code>make_net_recovery</code> is an Ignite-UX command, which allows you to create a recovery archive over the network onto the Ignite-UX server or any other specified system. The target system can be recovered across subnets after booting either from a bootable tape created by the Ignite-UX <code>make_boot_tape</code> command or the system boots directly from the Ignite-UX server. Booting directly from the Ignite-UX server can be automated with the Ignite-UX <code>bootsys</code> command or interactively specified on the boot console. |
| <b>make_tape_recovery</b>        | <code>make_tape_recovery</code> is a command on Ignite-UX which creates a bootable recovery (installation) tape, customized for your system and enables you unattended disaster recovery by connecting the backup device directly to the target system and booting the target system from the bootable recovery tape. The backup device has to be locally connected to the client during the creation of the archive and recovery of the client.                                                                                                             |
| <b>Manager-of-Managers (MoM)</b> | See MoM.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>MAPI</b>                      | ( <i>Microsoft Exchange Server specific term</i> ) The MAPI (Messaging Application Programming Interface) is the programming interface that lets applications and messaging clients interact with messaging and information systems.                                                                                                                                                                                                                                                                                                                         |
| <b>MCU</b>                       | See Main Control Unit (MCU).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Media Agent</b>               | A process that controls reading from and writing to a device, which reads from or writes to a medium (typically a tape). During a backup session, a Media Agent receives data from the Disk Agent and sends it to the device for writing it to the medium. During a restore or object verification session, a Media Agent locates data on the backup medium and sends it to the Disk Agent for processing. For a restore session, the Disk Agent then writes the data to the disk. A Media Agent also manages the robotics control of a library.             |
| <b>media allocation policy</b>   | Determines in which sequence media are used for backup. The Strict allocation policy directs Data Protector to prompt for a specific medium. The Loose policy directs Data Protector to prompt for any suitable medium. The Formatted First policy directs Data Protector to give preference to unknown media, even if unprotected media are available in the library.                                                                                                                                                                                       |
| <b>media condition</b>           | The quality of a medium as derived from the media condition factors. Heavy usage and age result in an increased number of read and write errors with tape media. Media need to be replaced when they are marked as POOR.                                                                                                                                                                                                                                                                                                                                     |
| <b>media condition factors</b>   | The user-assigned age threshold and overwrite threshold used to determine the state of a medium.                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>media label</b>               | A user-defined identifier used to describe a medium.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>media location</b>            | A user-defined physical location of a medium, such as "building 4" or "off-site storage".                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>media management session</b>  | A session performing some action on a medium, such as initializing, scanning the content, verifying data on a medium, or copying a medium.                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>media pool</b>                | A set of media of the same type (such as DDS) used and tracked as a group. Media are formatted and assigned to a media pool.                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>media set</b>                 | The result of a backup session is data backed up on a group of media called media set. Depending on the media usage policy, several sessions can share the same media.                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>media type</b>                | The physical type of media, such as DDS or DLT.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>media usage policy</b>        | The media usage policy controls how new backups are added to the already used media. It can be Appendable, Non-Appendable, or Appendable for incrementals only.                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>medium ID</b>                 | A unique identifier assigned to a medium by Data Protector.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>merging</b>                   | This defines one mode to resolve file conflicts during restore. If the file to be restored already exists at the destination, the one with the more recent modification date is kept. Files not present on the disk are always restored.<br>See also <code>overwrite</code> .                                                                                                                                                                                                                                                                                |
| <b>Microsoft Exchange Server</b> | A "client-server" messaging and a workgroup system that offers a transparent connection to many different communication systems. It provides users with an electronic mail system, individual and group scheduling, online forms, and workflow automation tools. It provides a developer with a platform on which to build custom information-sharing and messaging-service applications.                                                                                                                                                                    |

|                                                                               |                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Microsoft Management Console (MMC)</b>                                     | ( <i>Windows specific term</i> ) An administration model for Windows-based environments. It provides a simple, consistent, and integrated administration user interface allowing management of many applications through the same GUI, provided that the applications adhere to the MMC model.                                                                                                                                           |
| <b>Microsoft SQL Server</b>                                                   | A database management system designed to meet the requirements of distributed "client-server" computing.                                                                                                                                                                                                                                                                                                                                 |
| <b>Microsoft Volume Shadow Copy Service (VSS)</b>                             | A software service that provides a unified communication interface to coordinate backup and restore of a VSS-aware application regardless of its specific features. This service collaborates with the backup application, writers, shadow copy providers, and the operating system kernel to implement the management of volume shadow copies and shadow copy sets.<br>See also shadow copy, shadow copy provider, replica, and writer. |
| <b>mirror (EMC Symmetrix and HP P9000 XP Disk Array Family specific term)</b> | See target volume.                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>mirror rotation (HP P9000 XP Disk Array Family specific term)</b>          | See replica set rotation.                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>mirror unit (MU) number</b>                                                | ( <i>HP P9000 XP Disk Array Family specific term</i> ) A non-negative integer number that determines a secondary volume (S-VOL) of an internal disk (LDEV) located on a disk array of the HP P9000 XP Disk Array Family.<br>See also first-level mirror.                                                                                                                                                                                 |
| <b>mirrorclone</b>                                                            | ( <i>HP P6000 EVA Disk Array Family specific term</i> ) A dynamic replica of a storage volume, which is kept updated with changes made to the original storage volume via a local replication link. Replication between the original storage volume and its mirrorclone can be suspended. For each storage volume, a single mirrorclone can be created on the disk array.                                                                |
| <b>MMD</b>                                                                    | The Media Management Daemon process (service) runs on the Data Protector Cell Manager and controls media management and device operations. The process is started when Data Protector is installed on the Cell Manager.                                                                                                                                                                                                                  |
| <b>MMDB</b>                                                                   | The Media Management Database (MMDB) is a part of the IDB that contains information about media, media pools, devices, libraries, library drives, and slots configured in the cell, as well as the Data Protector media used for backup. In an enterprise backup environment, this part of the database can be common to all cells.<br>See also CMMDB and CDB.                                                                           |
| <b>MoM</b>                                                                    | Several cells can be grouped together and managed from a central cell. The management system of the central cell is the Manager-of-Managers (MoM). The cells are called MoM clients. The MoM enables you to configure and manage multiple cells from a central point.                                                                                                                                                                    |
| <b>mount point</b>                                                            | The access point in a directory structure for a disk or logical volume, for example /opt or d:. On UNIX, the mount points are displayed using the bdf or df command.                                                                                                                                                                                                                                                                     |
| <b>mount request</b>                                                          | A screen prompt that tells you to insert a specific medium into a device. Once you respond to the mount request by providing the required medium and confirm the mount request, the session continues.                                                                                                                                                                                                                                   |
| <b>MSM</b>                                                                    | The Data Protector Media Session Manager, which runs on the Cell Manager and controls media sessions, such as copying media.                                                                                                                                                                                                                                                                                                             |
| <b>multisnapping</b>                                                          | ( <i>HP P6000 EVA Disk Array Family specific term</i> ) Simultaneous creation of target volumes so that the backup data is consistent not only on each individual target volume, but also across all the volumes that constitute a snapshot.<br>See also snapshot.                                                                                                                                                                       |
| ○                                                                             |                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>OBDR capable device</b>                                                    | A device that can emulate a CD-ROM drive loaded with a bootable disk and can thus be used as a backup or boot device for disaster recovery purposes.                                                                                                                                                                                                                                                                                     |
| <b>obdrindex.dat</b>                                                          | See IDB recovery file.                                                                                                                                                                                                                                                                                                                                                                                                                   |

|                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>object</b>                       | See backup object.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>object consolidation</b>         | The process of merging a restore chain of a backup object, consisting of a full backup and at least one incremental backup, into a new, consolidated version of this object. The process is a part of the synthetic backup procedure. The result is a synthetic full backup of the specified backup object.                                                                                                                                                                                                                                          |
| <b>object consolidation session</b> | A process that merges a restore chain of a backup object, consisting of a full backup and at least one incremental backup, into a new, consolidated version of this object.                                                                                                                                                                                                                                                                                                                                                                          |
| <b>object copy</b>                  | A copy of a specific object version that is created during an object copy session or a backup session with object mirroring.                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>object copy session</b>          | A process that creates an additional copy of the backed up data on a different media set. During an object copy session, the selected backed up objects are copied from the source to the target media.                                                                                                                                                                                                                                                                                                                                              |
| <b>object copying</b>               | The process of copying selected object versions to a specific media set. You can select object versions from one or several backup sessions to be copied.                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>object ID</b>                    | ( <i>Windows specific term</i> ) The object IDs (OIDs) enable access to NTFS 5 files no matter where in the system the files reside. Data Protector treats the OIDs as alternate streams of the files.                                                                                                                                                                                                                                                                                                                                               |
| <b>object mirror</b>                | A copy of a backup object created using object mirroring. Object mirrors are often referred to as object copies.                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>object mirroring</b>             | The process of writing the same data to several media sets during a backup session. Data Protector enables you to mirror all or some backup objects to one or more media sets.                                                                                                                                                                                                                                                                                                                                                                       |
| <b>object verification</b>          | The process of verifying the data integrity of backup objects, from the Data Protector point of view, and the ability of Data Protector to deliver them to the required destination. The process can be used to provide a level of confidence in the ability to restore object versions created by backup, object copy, or object consolidation sessions.                                                                                                                                                                                            |
| <b>object verification session</b>  | A process that verifies the data integrity of specified backup objects or object versions and the ability of selected Data Protector network components to deliver them to a specified host. Object verification sessions can be run interactively, or as specified in automated post-backup, or scheduled specifications.                                                                                                                                                                                                                           |
| <b>offline backup</b>               | A backup during which an application database cannot be used by the application. In an offline backup session, the database is generally put into a quiescent state that allows use by the backup system, but not the application, for the time period of the data replication process. For instance, for backup to tape, until streaming of data to the tape is finished. Normal database operation is resumed before potential post-backup operations are started.<br>See also zero downtime backup (ZDB) and online backup.                       |
| <b>offline recovery</b>             | Offline recovery is performed if the Cell Manager is not accessible, for example, due to network problems. Only standalone and SCSI library devices can be used for offline recovery. Recovery of the Cell Manager is always offline.                                                                                                                                                                                                                                                                                                                |
| <b>offline redo log</b>             | See archived redo log.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>ON-Bar</b>                       | ( <i>Informix Server specific term</i> ) A backup and restore system for Informix Server. ON-Bar enables you to create a copy of your Informix Server data and later restore the data. The ON-Bar backup and restore system involves the following components: <ul style="list-style-type: none"> <li>• the onbar command</li> <li>• Data Protector as the backup solution</li> <li>• the XBSA interface</li> <li>• ON-Bar catalog tables, which are used to back up dbobjects and track instances of dbobjects through multiple backups.</li> </ul> |
| <b>ONCONFIG</b>                     | ( <i>Informix Server specific term</i> ) An environment variable that specifies the name of the active ONCONFIG configuration file. If the ONCONFIG environment variable is not present, Informix Server uses the configuration values from the onconfig file in the directory <i>INFORMIXDIR/etc</i> (on Windows) or <i>INFORMIXDIR/etc/</i> (on UNIX).                                                                                                                                                                                             |

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>online backup</b>     | <p>A backup performed while a database application remains available for use. The database is placed into a special backup mode of operation for the time period of the data replication process. For instance, for backup to tape, until streaming of data to tape is finished. During this period, the database is fully operational, but there may be a small performance impact and log files may grow very quickly. Normal database operation is resumed before potential post-backup operations are started.</p> <p>In some cases, transaction logs may also have to be backed up to allow a consistent database to be restored.</p> <p>See also zero downtime backup (ZDB) and offline backup.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>online recovery</b>   | <p>Online recovery is performed when Cell Manager is accessible. In this case, most of the Data Protector] functionalities are available (Cell Manager runs the session, restore sessions are logged in the IDB, you can monitor the restore progress using the GUI, and so on).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>online redo log</b>   | <p>(Oracle specific term) Redo logs that have not been archived, but are either available to the instance for recording database activity or are filled and waiting to be archived or reused.</p> <p>See also archived redo log.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Oracle Data Guard</b> | <p>(Oracle specific term) Oracle Data Guard is Oracle's primary disaster recovery solution. Oracle Data Guard is able to maintain up to nine standby databases, each of which is a real-time copy of the production (primary) database, to protect against corruptions, data failures, human errors, and disasters. If a failure occurs on the production (primary) database, then a failover to one of the standby databases which becomes the new primary database is possible. In addition, planned downtime for maintenance can be reduced because the production processing can be moved from the current primary database to a standby database and back quickly.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Oracle instance</b>   | <p>(Oracle specific term) Each installation of an Oracle database on one or more systems. One computer system can have several instances of a database running.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>ORACLE_SID</b>        | <p>(Oracle specific term) A unique name for an Oracle Server instance. To switch among Oracle Servers, specify the desired <code>ORACLE_SID</code>. The <code>ORACLE_SID</code> is included in the <code>CONNECT DATA</code> parts of the connect descriptor in a <code>TNSNAMES.ORA</code> file and in the definition of the TNS listener in the <code>LISTENER.ORA</code> file.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>original system</b>   | <p>The system configuration backed up by Data Protector before a computer disaster hits the system.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>overwrite</b>         | <p>An option that defines one mode to resolve file conflicts during restore. All files are restored from a backup even if they are older than existing files.</p> <p>See also merging.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>ownership</b>         | <p>Backup ownership affects the ability of users to see and restore data. Each backup session and all the data backed up within it is assigned an owner. The owner can be the user that starts an interactive backup, the account under which the CRS process is running, or the user specified as the owner in the backup specification options.</p> <p>If a user starts an existing backup specification without modifying it, the backup session is not considered as interactive.</p> <p>If a modified backup specification is started by a user, the user is the owner unless the following is true:</p> <ul style="list-style-type: none"> <li>• The user has the Switch Session Ownership user right.</li> <li>• The backup session owner is explicitly defined in the backup specification, where the username, group or domain name, and the system name are specified.</li> </ul> <p>If a backup is scheduled on a UNIX Cell Manager, the session owner is root:sys unless the above conditions are true.</p> <p>If a backup is scheduled on a Windows Cell Manager, the session owner is the user specified during the installation, unless the above conditions are true.</p> <p>When copying or consolidating objects, by default the owner is the user who starts the operation, unless a different owner is specified in the copy or consolidation specification.</p> |

## P

|                 |                                                                                                                                                                                                            |
|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>P1S file</b> | <p>P1S file contains information on how to format and partition all disks installed in the system during Enhanced Automated Disaster Recovery (EADR). It is created during full backup and is saved on</p> |
|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

|                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                     | <p>backup medium and on Cell Manager into the directory<br/> <i>Data_Protector_program_data\Config\Server\dr\pls</i> (Windows Server 2008),<br/> <i>Data_Protector_home\Config\Server\dr\pls</i> (other Windows systems), or<br/> <i>/etc/opt/omni/server/dr/pls</i> (UNIX systems) with the filename <i>recovery.pls</i>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>package</b>                      | (MC/ServiceGuard and Veritas Cluster specific term) A collection of resources (for example volume groups, application services, IP names, and addresses) that are needed to run a specific cluster-aware application.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>pair status</b>                  | <p>(HP P9000 XP Disk Array Family specific term) The status of a disk pair (secondary volume and its corresponding primary volume) of a disk array of the HP P9000 XP Disk Array Family. Depending on the circumstances, the paired disks can be in various states. The following states are particularly important for the operation of the Data Protector HP StorageWorks P9000 XP Agent:</p> <ul style="list-style-type: none"> <li>• PAIR – The secondary volume is prepared for zero downtime backup. If it is a mirror, it is completely synchronized, and if it is a volume to be used for snapshot storage, it is empty.</li> <li>• SUSPENDED – The link between the disks is suspended. However, the pair relationship is still maintained, and the secondary disk can be prepared for zero downtime backup again at a later time.</li> <li>• COPY – The disk pair is currently busy and making a transition into the PAIR state. If the secondary volume is a mirror, it is re-synchronizing with the primary volume, and if it is a volume to be used for snapshot storage, its contents are getting cleared.</li> </ul> |
| <b>parallel restore</b>             | Restoring backed up data to multiple disks at the same time (that is, in parallel) by running multiple Disk Agents, that receive data from one Media Agent. For the parallel restore to work, select data that is located on different disks or logical volumes and during backup, the data from the different objects must have been sent to the same device using a concurrency of 2 or more. During a parallel restore, the data for multiple objects selected for restore is read from media at the same time, thereby improving performance.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>parallelism</b>                  | The concept of reading multiple data streams from an online database.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>phase 0 of disaster recovery</b> | Preparation for disaster recovery - the prerequisite condition for a successful disaster recovery.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>phase 1 of disaster recovery</b> | Installation and configuration of DR OS, establishing previous storage structure.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>phase 2 of disaster recovery</b> | Restoration of operating system (with all the configuration information that defines the environment) and Data Protector.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>phase 3 of disaster recovery</b> | Restoration of user and application data.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>physical device</b>              | A physical unit that contains either a drive or a more complex unit such as a library.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>post-exec</b>                    | <p>A backup option that executes a command or script after the backup of an object or after the entire session completes. Post-exec commands are not supplied by Data Protector. You need to create your own. They can be written as executables or batch files on Windows and as shell scripts on UNIX.</p> <p>See also pre-exec.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>pre- and post-exec commands</b>  | Pre- and post-exec commands are used to perform additional action before and after a backup or restore session. They are not supplied by Data Protector. You need to create your own commands. They can be written as executables or batch files on Windows and as shell scripts on UNIX.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>pre-exec</b>                     | <p>A backup option that executes a command or script before the backup of an object or before the entire session is started. Pre-exec commands are not supplied by Data Protector. You need to create your own. They can be written as executables or batch files on Windows and as shell scripts on UNIX.</p> <p>See also post-exec.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>prealloc list</b>                | A subset of media in a media pool that specifies the order in which media are used for backup.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>primary volume (P-VOL)</b>       | (HP P9000 XP Disk Array Family specific term) An internal disk (LDEV) of a disk array of the HP P9000 XP Disk Array Family for which a secondary volume (S-VOL), either its mirror or a volume                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |



|                                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                      | to be used for its snapshot storage, exists. In the HP CA P9000 XP and HP CA+BC P9000 XP configurations, primary volumes are located in the Main Control Unit (MCU).<br>See also secondary volume (S-VOL) and Main Control Unit (MCU).                                                                                                                                                                                                                                                                                                                     |
| <b>protection</b>                    | See data protection and also catalog protection.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>public folder store</b>           | ( <i>Microsoft Exchange Server specific term</i> ) The part of the Information Store that maintains information in public folders. A public folder store consists of a binary rich-text .edb file and a streaming native internet content .stm file.                                                                                                                                                                                                                                                                                                       |
| <b>public/private backed up data</b> | When configuring a backup, you can select whether the backed up data will be: <ul style="list-style-type: none"> <li>• public, that is visible (and accessible for restore) to all Data Protector users</li> <li>• private, that is, visible (and accessible for restore) only to the owner of the backup and administrators</li> </ul>                                                                                                                                                                                                                    |
| <b>R</b>                             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>RAID</b>                          | Redundant Array of Independent Disks.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>RAID Manager Library</b>          | ( <i>HP P9000 XP Disk Array Family specific term</i> ) A software library that is used for accessing the configuration, status, and performance measurement data of a P9000 XP Array storage system, and for invoking operations on the disk array. It translates function calls into sequences of low-level SCSI commands.<br>See also HP StorageWorks P9000 XP Agent.                                                                                                                                                                                    |
| <b>RAID Manager P9000 XP</b>         | ( <i>HP P9000 XP Disk Array Family specific term</i> ) A software application that provides a command-line interface to disk arrays of the HP P9000 XP Disk Array Family. It offers an extensive set of commands for reporting and controlling the status of a P9000 XP Array storage system, and for performing various operations on the disk array.                                                                                                                                                                                                     |
| <b>rawdisk backup</b>                | See disk image backup.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>RCU</b>                           | See Remote Control Unit (RCU).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>RDBMS</b>                         | Relational Database Management System.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>RDF1/RDF2</b>                     | ( <i>EMC Symmetrix specific term</i> ) A type of SRDF device group. Only RDF devices can be assigned to an RDF group. An RDF1 group type contains source (R1) devices and an RDF2 group type contains target (R2) devices.                                                                                                                                                                                                                                                                                                                                 |
| <b>RDS</b>                           | The Raima Database Server process (service) runs on the Data Protector Cell Manager and manages the IDB. The process is started when Data Protector is installed on the Cell Manager.                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Recovery Catalog</b>              | ( <i>Oracle specific term</i> ) A set of Oracle tables and views that are used by Recovery Manager to store information about Oracle databases. This information is used by Recovery Manager to manage the backup, restore, and recovery of Oracle databases. The recovery catalog contains information about: <ul style="list-style-type: none"> <li>• The physical schema of the Oracle target database</li> <li>• Data file and archived log backup sets</li> <li>• Data file copies</li> <li>• Archived Redo Logs</li> <li>• Stored scripts</li> </ul> |
| <b>Recovery Catalog Database</b>     | ( <i>Oracle specific term</i> ) An Oracle database that contains a recovery catalog schema. You should not store the recovery catalog in your target database.                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>recovery files</b>                | ( <i>Oracle specific term</i> ) Recovery files are Oracle specific files that reside in the flash recovery area: the current control file, online redo logs, archived redo logs, flashback logs, control file autobackups, datafile copies, and backup pieces.<br>See also flash recovery area.                                                                                                                                                                                                                                                            |
| <b>Recovery Manager (RMAN)</b>       | ( <i>Oracle specific term</i> ) An Oracle command-line interface that directs an Oracle Server process to back up, restore, or recover the database it is connected to. RMAN uses either the recovery catalog or the control file to store information about backups. This information can be used later in restore sessions.                                                                                                                                                                                                                              |



|                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>RecoveryInfo</b>                          | When backing up Windows configuration files, Data Protector collects the information about the current system configuration (information on disk layout, volume, and network configuration). This information is needed for disaster recovery.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>recycle or unprotect</b>                  | A process that removes the data protection from all backed up data on a medium, allowing Data Protector to overwrite it during one of the next backups. Data that belongs to the same session(s) but resides on other media is also unprotected. Recycling does not actually alter the data on the medium.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>redo log</b>                              | <i>(Oracle specific term)</i> Every Oracle database has a set of two or more redo log files. The set of redo log files for the database is known as the database's redo log. Oracle uses the redo log to record all changes made to data.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Remote Control Unit (RCU)</b>             | <i>(HP P9000 XP Disk Array Family specific term)</i> An HP P9000 XP Disk Array Family unit that acts as a slave device to the Main Control Unit (MCU) in the HP CA P9000 XP or HP CA+BC P9000 XP configuration. In bidirectional configurations, the RCU can also act as an MCU.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Removable Storage Management Database</b> | <i>(Windows specific term)</i> A Windows service used for managing removable media (such as tapes and disks) and storage devices (libraries). Removable Storage allows applications to access and share the same media resources.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>reparse point</b>                         | <i>(Windows specific term)</i> A system-controlled attribute that can be associated with any directory or file. The value of a reparse attribute can have user-controlled data. The format of the data is understood by the application that stored the data and a filesystem filter that was installed to interpret the data and process such files. Whenever the filesystem encounters a file with a reparse point, it attempts to find the filesystem filter associated with the data format.                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>replica</b>                               | <i>(ZDB specific term)</i> An image, at a particular point in time, of the data in source volumes that contain user-specified backup objects. Depending on the hardware or software with which it is created, the image may be an independent exact duplicate (clone) of the storage blocks at the physical disk level (for example, a split mirror or snapclone), or a virtual copy (for example, a snapshot). From perspective of a basic operating system, the complete physical disk containing backup objects is replicated. However, if a volume manager is used on UNIX, the whole volume or disk group containing a backup object (logical volume) is replicated. If partitions are used on Windows, the whole physical volume containing the selected partition is replicated.<br><i>See also</i> snapshot, snapshot creation, split mirror, and split mirror creation. |
| <b>replica set</b>                           | <i>(ZDB specific term)</i> A group of replicas, all created using the same backup specification.<br><i>See also</i> replica and replica set rotation.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>replica set rotation</b>                  | <i>(ZDB specific term)</i> The use of a replica set for regular backup production: Each time the same backup specification requiring the use of a replica set is run, a new replica is created and added to the set, until the maximum number for the set is reached. After that, the oldest replica in the set is replaced, maintaining the maximum number of replicas in the set.<br><i>See also</i> replica and replica set.                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>restore chain</b>                         | All backups that are necessary for a restore of a backup object to a certain point in time. A restore chain consists of a full backup of the object and any number of related incremental backups.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>restore session</b>                       | A process that copies data from backup media to a client.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>resync mode</b>                           | <i>(HP P9000 XP Disk Array Family VSS provider specific term)</i> One of two P9000 XP Array VSS hardware provider operation modes. When the P9000 XP Array provider is in the resync mode, the source volume (P-VOL) and its replica (S-VOL) are in the suspended mirror relationship after a backup. The maximum number of replicas (S-VOLs per a P-VOL) rotated is three provided that MU range is 0-2 or 0, 1, 2. Restore from a backup in such a configuration is possible only by re-synchronization of an S-VOL with its P-VOL.<br><i>See also</i> VSS compliant mode, source volume, primary volume (P-VOL), replica, secondary volume (S-VOL), mirror unit (MU) number, and replica set rotation.                                                                                                                                                                        |
| <b>RMAN (Oracle specific term)</b>           | <i>See</i> Recovery Manager.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>RSM</b>                                   | The Data Protector Restore Session Manager controls restore and object verification sessions. This process always runs on the Cell Manager system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>RSM</b>                      | ( <i>Windows specific term</i> ) Removable Storage Manager (RSM) includes a media management service that facilitates communication among applications, robotic changers, and media libraries. It enables multiple applications to share local robotic media libraries and tape or disk drives and to manage removable media.                                                                                                                                                                                                                                                     |
| <b>S</b>                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>SAPDBA</b>                   | ( <i>SAP R/3 specific term</i> ) An SAP R/3 user interface that integrates the BRBACKUP, BRARCHIVE, and BRRESTORE tools.                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>scanning</b>                 | A function which identifies the media in a device. This synchronizes the MMDB with the media that are actually present at the selected locations (for example, slots in a library). It is useful to perform a scan and check the actual media in the device if someone has manually manipulated media without using Data Protector to eject or enter, for example.                                                                                                                                                                                                                |
| <b>Scheduler</b>                | A function that controls when and how often automatic backups occur. By setting up a schedule, you automate the start of backups.                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>secondary volume (S-VOL)</b> | ( <i>HP P9000 XP Disk Array Family specific term</i> ) An internal disk (LDEV) of a disk array of the HP P9000 XP Disk Array Family which is paired with another LDEV: a primary volume (P-VOL). It can act as a mirror of the P-VOL or as a volume to be used for the P-VOL's snapshot storage. An S-VOL is assigned a SCSI address different from the one used for the P-VOL. In an HP CA P9000 XP configuration, the S-VOLs acting as mirrors can be used as failover devices in a MetroCluster configuration.<br>See also primary volume (P-VOL) and Main Control Unit (MCU). |
| <b>session</b>                  | See backup session, media management session, and restore session.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>session ID</b>               | An identifier of a backup, restore, object copy, object consolidation, object verification, or media management session, consisting of the date when the session ran and a unique number.                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>session key</b>              | This environment variable for the pre-exec and post-exec script is a Data Protector unique identification of any session, including preview sessions. The session key is not recorded in the database, and it is used for specifying options for the omnimnt, omnistat, and omniabort commands.                                                                                                                                                                                                                                                                                   |
| <b>shadow copy</b>              | ( <i>Microsoft VSS specific term</i> ) A volume that represents a duplicate of the original volume at a certain point in time. The data is then backed up from the shadow copy and not from the original volume. The original volume continues to change as the backup process continues, but the shadow copy of the volume remains constant.<br>See also Microsoft Volume Shadow Copy Service and replica.                                                                                                                                                                       |
| <b>shadow copy provider</b>     | ( <i>Microsoft VSS specific term</i> ) An entity that performs the work on creating and representing the volume shadow copies. Providers own the shadow copy data and expose the shadow copies. Providers can be software (for example, system providers) or hardware (local disks, disk arrays).<br>See also shadow copy.                                                                                                                                                                                                                                                        |
| <b>shadow copy set</b>          | ( <i>Microsoft VSS specific term</i> ) A collection of shadow copies created at the same point in time.<br>See also shadow copy and replica set.                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>shared disks</b>             | A Windows disk on another system that has been made available to other users on the network. Systems with shared disks can be backed up without a Data Protector Disk Agent installed.                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>SIBF</b>                     | The Serverless Integrations Binary Files (SIBF) is a part of the IDB that stores raw NDMP meta data. This data is necessary to perform restore of NDMP objects.                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Site Replication Service</b> | ( <i>Microsoft Exchange Server specific term</i> ) The Microsoft Exchange Server 2003 service that permits compatibility with Microsoft Exchange Server 5.5 by emulating the Exchange Server 5.5 directory service.<br>See also Information Store and Key Management Service.                                                                                                                                                                                                                                                                                                     |
| <b>slot</b>                     | A mechanical position in a library. Each slot can hold a medium, such as a DLT tape. Data Protector references each slot with a number. To read a medium, a robotic mechanism moves the medium from a slot into the drive.                                                                                                                                                                                                                                                                                                                                                        |
| <b>smart copy</b>               | ( <i>VLS specific term</i> ) A copy of the backed up data created from the virtual tape to the physical tape library. The smart copy process allows Data Protector to distinguish between the source and the target medium thus enabling media management.<br>See also Virtual Library System (VLS).                                                                                                                                                                                                                                                                              |

|                                                                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>smart copy pool</b>                                                   | <i>(VLS specific term)</i> A pool that defines which destination library slots are available as smart copy targets for a specified source virtual library.<br>See also Virtual Library System (VLS) and smart copy.                                                                                                                                                                                                                                                                                                                                         |
| <b>SMB</b>                                                               | See split mirror backup.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>SMBF</b>                                                              | The Session Messages Binary Files (SMBF) part of the LDB stores session messages generated during backup, restore, object copy, object consolidation, object verification, and media management sessions. One binary file is created per session. The files are grouped by year and month.                                                                                                                                                                                                                                                                  |
| <b>SMI-S Agent (SMISA)</b>                                               | See HP StorageWorks P6000 EVA SMI-S Agent.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>snapshot</b>                                                          | <i>(HP P6000 EVA Disk Array Family, HP P9000 XP Disk Array Family, and HP P4000 SAN Solutions specific term)</i> A type of target volumes created using a specific replication technology. Depending on the disk array model and the chosen replication technique, a range of snapshot types with different characteristics is available. Basically, each snapshot may be either a virtual copy, still reliant upon the contents of the source volume, or an independent duplicate (clone) of the source volume.<br>See also replica and snapshot creation. |
| <b>snapshot backup</b>                                                   | See ZDB to tape, ZDB to disk, and ZDB to disk+tape.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>snapshot creation</b>                                                 | <i>(HP P6000 EVA Disk Array Family, HP P9000 XP Disk Array Family, and HP P4000 SAN Solutions specific term)</i> A replica creation process in which copies of the selected source volumes are created using storage virtualization technology. Such a replica is considered to be created at a particular point in time, and is immediately available for use. However, with certain snapshot types, a background data copying process continues to run on the disk array after the moment of the replica creation.<br>See also snapshot.                  |
| <b>source (R1) device</b>                                                | <i>(EMC Symmetrix specific term)</i> An EMC Symmetrix device that participates in SRDF operations with a target (R2) device. All writes to this device are mirrored to a target (R2) device in a remote EMC Symmetrix unit. An R1 device must be assigned to an RDF1 group type.<br>See also target (R2) device.                                                                                                                                                                                                                                            |
| <b>source volume</b>                                                     | <i>(ZDB specific term)</i> A storage volume containing data to be replicated.                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>sparse file</b>                                                       | A file that contains data with portions of empty blocks. Examples are: a matrix in which some or much of the data contains zeros, files from image applications, and high-speed databases. If sparse file processing is not enabled during restore, it might be impossible to restore this file.                                                                                                                                                                                                                                                            |
| <b>split mirror</b>                                                      | <i>(EMC Symmetrix Disk Array and HP P9000 XP Disk Array Family specific term)</i> A type of target volumes created using a specific replication technology. A split-mirror replica provides independent duplicates (clones) of the source volumes.<br>See also replica and split mirror creation.                                                                                                                                                                                                                                                           |
| <b>split mirror backup (EMC Symmetrix specific term)</b>                 | See ZDB to tape.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>split mirror backup (HP P9000 XP Disk Array Family specific term)</b> | See ZDB to tape, ZDB to disk, and ZDB to disk+tape.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>split mirror creation</b>                                             | <i>(EMC Symmetrix and HP P9000 XP Disk Array Family specific term)</i> A replica creation technique, in which a pre-configured set of target volumes (a mirror) is kept synchronized with a set of source volumes until the time at which a replica of the contents of the source volumes is required. Then, the synchronization is stopped (the mirror is split) and a split mirror replica of the source volumes at the time of the split remains in the target volumes.<br>See also split mirror.                                                        |
| <b>split mirror restore</b>                                              | <i>(EMC Symmetrix and HP P9000 XP Disk Array Family specific term)</i> A process in which data backed up in a ZDB-to-tape or a ZDB-to-disk+tape session is first copied from the backup media to a replica, and from the replica to the source volumes afterwards. Individual backup objects or complete sessions can be restored using this method.<br>See also ZDB to tape, ZDB to disk+tape, and replica.                                                                                                                                                |

|                                  |                                                                                                                                                                                                                                                                                                                                                                                               |
|----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>sqlhosts file or registry</b> | <i>(Informix Server specific term)</i> An Informix Server connectivity information file (on UNIX) or registry (on Windows) that contains the names of each of the database servers and any aliases to which the clients on a host computer can connect.                                                                                                                                       |
| <b>SRD file</b>                  | <i>(disaster recovery specific term)</i> A text file in the Unicode (UTF-16) format, generated during CONFIGURATION backup of a Windows or Linux system and stored on the Cell Manager. It contains system information required for installing and configuring the operating system on the target system in the case of a disaster.<br>See also target system.                                |
| <b>SRDF</b>                      | <i>(EMC Symmetrix specific term)</i> The EMC Symmetrix Remote Data Facility is a business continuation process that enables effective, real-time data replication of SLDs between dislocated processing environments. These environments could be situated within the same root computer environment or separated by long distances.                                                          |
| <b>SSE Agent (SSEA)</b>          | See HP StorageWorks P9000 XP Agent.                                                                                                                                                                                                                                                                                                                                                           |
| <b>sst.conf file</b>             | The file <code>/usr/kernel/drv/sst.conf</code> is required on each Data Protector Sun Solaris client to which a multi-drive library device is connected. It must contain an entry for the SCSI address of the robotic mechanism of each library device connected to the client.                                                                                                               |
| <b>st.conf file</b>              | The file <code>/kernel/drv/st.conf</code> is required on each Data Protector Solaris client with a backup device connected. It must contain device information and a SCSI address for each backup drive connected to that client. A single SCSI entry is required for a single-drive device and multiple SCSI entries are required for a multi-drive library device.                          |
| <b>stackers</b>                  | Devices with multiple slots for media storage usually with only one drive. A stacker selects media from the stack sequentially. In contrast, a library can randomly select media from its repository.                                                                                                                                                                                         |
| <b>standalone file device</b>    | A file device is a file in a specified directory to which you back up data.                                                                                                                                                                                                                                                                                                                   |
| <b>Storage Group</b>             | <i>(Microsoft Exchange Server specific term)</i> A collection of mailbox stores and public folder stores that share a set of transaction log files. Exchange Server manages each storage group with a separate server process.                                                                                                                                                                |
| <b>storage volume</b>            | <i>(ZDB specific term)</i> An object that may be presented to an operating system or some other entity (for example, a virtualization mechanism) upon which volume management systems, filesystems, or other objects may exist. The volume management systems, file systems are built on this storage. Typically, these can be created or exist within a storage system such as a disk array. |
| <b>StorageTek ACS library</b>    | <i>(StorageTek specific term)</i> Automated Cartridge System is a library system (also known as Silo) consisting of one Library Management Unit (LMU) and one to 24 Library Storage Modules (LSM) connected to the unit.                                                                                                                                                                      |
| <b>switchover</b>                | See failover.                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Sybase Backup Server API</b>  | <i>(Sybase specific term)</i> An industry-standard interface developed for the exchange of backup and recovery information between a Sybase SQL Server and a backup solution like Data Protector.                                                                                                                                                                                             |
| <b>Sybase SQL Server</b>         | <i>(Sybase specific term)</i> The server in the Sybase "client-server" architecture. Sybase SQL Server manages multiple databases and multiple users, keeps track of the actual location of data on disks, maintains mapping of logical data description to physical data storage, and maintains data and procedure caches in memory.                                                         |
| <b>SYMA</b>                      | See EMC Symmetrix Agent.                                                                                                                                                                                                                                                                                                                                                                      |
| <b>synthetic backup</b>          | A backup solution that produces a synthetic full backup, an equivalent to a conventional full backup in terms of data, without putting stress on the production servers or the network. A synthetic full backup is created from a previous full backup and any number of incremental backups.                                                                                                 |
| <b>synthetic full backup</b>     | The result of an object consolidation operation, where a restore chain of a backup objects is merged into a new, synthetic full version of this object. A synthetic full backup is equivalent to a conventional full backup in terms of restore speed.                                                                                                                                        |
| <b>System Backup to Tape</b>     | <i>(Oracle specific term)</i> An Oracle interface that handles the actions required to load, label, and unload correct backup devices when Oracle issues a backup or restore request.                                                                                                                                                                                                         |

|                                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>system databases</b>                    | <p>(<i>Sybase specific term</i>) The four system databases on a newly installed Sybase SQL Server are the:</p> <ul style="list-style-type: none"> <li>• master database (master)</li> <li>• temporary database (tempdb)</li> <li>• system procedure database (sybsystemprocs)</li> <li>• model database (model).</li> </ul>                                                                                                                                                                                                                                                         |
| <b>System Recovery Data file</b>           | See SRD file.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>System State</b>                        | <p>(<i>Windows specific term</i>) The System State data comprises the Registry, COM+ Class Registration database, system startup files, and the Certificate Services database (if the server is a certificate server). If the server is a domain controller, Active Directory services and the SYSVOL directory are also contained in the System State data. If the server is running the Cluster service, the System State data also includes resource registry checkpoints and the quorum resource recovery log, which contains the most recent cluster database information.</p> |
| <b>system volume/disk/partition</b>        | A volume/disk/partition containing operating system files. Microsoft terminology defines the system volume/disk/partition as the volume/disk/partition containing files required for the initial step of the boot process.                                                                                                                                                                                                                                                                                                                                                          |
| <b>SysVol</b>                              | <p>(<i>Windows specific term</i>) A shared directory that stores the server copy of the domain's public files, which are replicated among all domain controllers in the domain.</p>                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>T</b>                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>tablespace</b>                          | A part of a database structure. Each database is logically divided into one or more tablespaces. Each tablespace has data files or raw volumes exclusively associated with it.                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>tapeless backup (ZDB specific term)</b> | See ZDB to disk.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>target (R2) device</b>                  | <p>(<i>EMC Symmetrix specific term</i>) An EMC Symmetrix device that participates in SRDF operations with a source (R1) device. It resides in the remote EMC Symmetrix unit. It is paired with a source (R1) device in the local EMC Symmetrix unit and receives all write data from its mirrored pair. This device is not accessed by user applications during normal I/O operations. An R2 device must be assigned to an RDF2 group type.</p> <p>See also source (R1) device.</p>                                                                                                 |
| <b>target database</b>                     | <p>(<i>Oracle specific term</i>) In RMAN, the target database is the database that you are backing up or restoring.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>target system</b>                       | <p>(<i>disaster recovery specific term</i>) A system after a computer disaster has occurred. The target system is typically in a non-bootable state and the goal of disaster recovery is to restore this system to the original system configuration. The difference between a faulty system and a target system is that a target system has all faulty hardware replaced.</p>                                                                                                                                                                                                      |
| <b>target volume</b>                       | ( <i>ZDB specific term</i> ) A storage volume to which data is replicated.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Terminal Services</b>                   | <p>(<i>Windows specific term</i>) Windows Terminal Services provide a multi-session environment that allows clients to access a virtual Windows desktop session and Windows-based programs running on the server.</p>                                                                                                                                                                                                                                                                                                                                                               |
| <b>thread</b>                              | <p>(<i>Microsoft SQL Server specific term</i>) An executable entity that belongs to only one process. It comprises a program counter, a user-mode stack, a kernel-mode stack, and a set of register values. Several threads can run at a time within one process.</p>                                                                                                                                                                                                                                                                                                               |
| <b>TimeFinder</b>                          | <p>(<i>EMC Symmetrix specific term</i>) A business continuation process that creates an instant copy of single or multiple Symmetrix Logical Devices (SLDs). The instant copy is created on specially preconfigured SLDs called BCVs and is accessible via a separate device address to the system(s).</p>                                                                                                                                                                                                                                                                          |
| <b>TLU</b>                                 | Tape Library Unit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>TNSNAMES.ORA</b>                        | <p>(<i>Oracle and SAP R/3 specific term</i>) A network configuration file that contains connect descriptors mapped to service names. The file may be maintained centrally or locally, for use by all or individual clients.</p>                                                                                                                                                                                                                                                                                                                                                     |

|                               |                                                                                                                                                                                                                                                                    |
|-------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>transaction</b>            | A mechanism for ensuring that a set of actions is treated as a single unit of work. Databases use transactions to keep track of database changes.                                                                                                                  |
| <b>transaction backup</b>     | Transaction backups generally use fewer resources than database backups, so they can be created more frequently than database backups. By applying transaction backups, you can recover the database to a specific point in time prior to when a problem occurred. |
| <b>transaction backup</b>     | <i>(Sybase and SQL specific term)</i> A backup of the transaction log providing a record of changes made since the last full or transaction backup.                                                                                                                |
| <b>transaction log backup</b> | Transaction log backups generally use fewer resources than database backups so they can be created more frequently than database backups. By applying transaction log backups, you can recover the database to a specific point in time.                           |
| <b>transaction log files</b>  | Files that record transactions of the database modifications, and provide fault tolerance in case of a database disaster.                                                                                                                                          |
| <b>transaction log table</b>  | <i>(Sybase specific term)</i> A system table in which all changes to the database are automatically recorded.                                                                                                                                                      |
| <b>transaction logs</b>       | <i>(Data Protector specific term)</i> Keep track of IDB changes. The archiving of transaction logs should be enabled to prevent you from losing the transaction log files that are created after the last IDB backup and are necessary for IDB recovery.           |
| <b>transportable snapshot</b> | <i>(Microsoft VSS specific term)</i> A shadow copy that is created on the application system and can be presented to the backup system where a backup can be performed.<br>See also Microsoft Volume Shadow Copy Service (VSS).                                    |
| <b>TSANDS.CFG file</b>        | <i>(Novell NetWare specific term)</i> A file that allows you to specify the names of containers where you want backups to begin. It is a text file located in the <code>SYS:SYSTEM\TSA</code> directory on the server where <code>TSANDS.NLM</code> is loaded.     |

## U

|                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>UIProxy</b>                                    | The Java GUI Server ( <code>UIProxy</code> service) runs on the Data Protector Cell Manager. It is responsible for communication between the Java GUI Client and the Cell Manager, moreover, it performs business logic operations and sends only important information to the client. The service is started as soon as Data Protector is installed on the Cell Manager.                                                              |
| <b>unattended operation</b>                       | See lights-out operation.                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>user account (Data Protector user account)</b> | You can use Data Protector only if you have a Data Protector user account, which restricts unauthorized access to Data Protector and to backed up data. Data Protector administrators create this account specifying a user logon name, the systems from which the user can log on, and a Data Protector user group membership. This is checked whenever the user starts the Data Protector user interface or performs specific tasks. |
| <b>User Account Control (UAC)</b>                 | A security component in Windows Vista, Windows 7, and Windows Server 2008 that limits application software to standard user privileges until an administrator authorizes an increase in privilege level.                                                                                                                                                                                                                               |
| <b>user disk quotas</b>                           | NTFS quota-management support enables an enhanced tracking mechanism and control over disk space usage on shared storage volumes. Data Protector backs up user disk quotas on the whole system and for all configured users at a time.                                                                                                                                                                                                 |
| <b>user group</b>                                 | Each Data Protector user is member of a User Group. Each User Group has a set of user rights that are given to every user in that User Group. The number of User Groups with their associated user rights can be defined as desired. Data Protector provides three default user groups: admin, operator, and user.                                                                                                                     |
| <b>user profile</b>                               | <i>(Windows specific term)</i> Configuration information retained on a user basis. This information includes desktop settings, screen colors, network connections, and so on. When the user logs on, the user profile is loaded and the Windows environment is set accordingly.                                                                                                                                                        |
| <b>user rights</b>                                | User rights or access rights are the permissions needed to perform specific Data Protector tasks. Configuring a backup, starting a backup session, or starting a restore session are typical user rights. Users have the access rights of the user group to which they belong.                                                                                                                                                         |

|                                          |                                                                                                                                                                                                                                                                                                                                                                 |
|------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>user_restrictions file</b>            | A file that restricts specific user actions, which are available to Data Protector user groups according to the user rights assigned to them, to be performed only on specific systems of the Data Protector cell. Such restrictions apply only to Data Protector user groups other than <i>admin</i> and <i>operator</i> .                                     |
| V                                        |                                                                                                                                                                                                                                                                                                                                                                 |
| <b>vaulting media</b>                    | The process of storing media to a safe and remote place. The media are brought back to the data center when they are needed for restore or are ready for reuse in subsequent backups. The vaulting procedure depends on your company's backup strategy and policies for data protection/reliability.                                                            |
| <b>verify</b>                            | A function that lets you check whether the Data Protector data on a specified medium is readable. Additionally, consistency within each block can be checked if the backup was performed with the cyclic redundancy check (CRC) option ON.                                                                                                                      |
| <b>Virtual Controller Software (VCS)</b> | ( <i>HP P6000 EVA Disk Array Family specific term</i> ) The firmware that manages all aspects of storage system operation, including communication with HP Command View EVA through the HSV controllers.<br>See also HP Command View (CV) EVA.                                                                                                                  |
| <b>Virtual Device Interface</b>          | ( <i>Microsoft SQL Server specific term</i> ) This is a Microsoft SQL Server programming interface that allows fast backup and restore of large databases.                                                                                                                                                                                                      |
| <b>virtual disk</b>                      | ( <i>HP P6000 EVA Disk Array Family specific term</i> ) A unit of storage allocated from a storage pool of a disk array of the HP P6000 EVA Disk Array Family. A virtual disk is the entity that can be replicated using the snapshot functionality of such disk array.<br>See also source volume and target volume.                                            |
| <b>virtual full backup</b>               | An efficient type of synthetic backup where data is consolidated using pointers instead of being copied. It is performed if all the backups (the full backup, incremental backups, and the resulting virtual full backup) are written to a single file library that uses distributed file medium format.                                                        |
| <b>Virtual Library System (VLS)</b>      | A disk-based data storage device hosting one or more virtual tape libraries (VTLs).                                                                                                                                                                                                                                                                             |
| <b>virtual server</b>                    | A virtual machine in a cluster environment defined in a domain by a network IP name and address. Its address is cached by the cluster software and mapped to the cluster node that is currently running the virtual server resources. This way all requests for a particular virtual server are cached by a specific cluster node.                              |
| <b>virtual tape</b>                      | ( <i>VLS specific term</i> ) An archival storage technology that backs up data to disk drives in the same way as if it were being stored on tape. Benefits of virtual tape systems include improved backup and recovery speed and lower operating costs.<br>See also Virtual Library System (VLS) and Virtual Tape Library (VTL).                               |
| <b>Virtual Tape Library (VTL)</b>        | ( <i>VLS specific term</i> ) An emulated tape library that provides the functionality of traditional tape-based storage.<br>See also Virtual Library System (VLS).                                                                                                                                                                                              |
| <b>VMware management client</b>          | ( <i>VMware (Legacy) integration specific term</i> ) The client that Data Protector uses to communicate with VMware Virtual Infrastructure. This can be a VirtualCenter Server system (VirtualCenter environment) or an ESX Server system (standalone ESX Server environment).                                                                                  |
| <b>volser</b>                            | ( <i>ADIC and STK specific term</i> ) A VOLUME SERIAL number is a label on the medium to identify the physical tape used in very large libraries. A volser is a naming convention specific to ADIC/GRAU and StorageTek devices.                                                                                                                                 |
| <b>volume group</b>                      | A unit of data storage in an LVM system. A volume group can consist of one or more physical volumes. There can be more than one volume group on the system.                                                                                                                                                                                                     |
| <b>volume mountpoint</b>                 | ( <i>Windows specific term</i> ) An empty directory on a volume that can be used to mount another volume. The volume mount point acts as a gateway to the target volume. Provided that the volume is mounted, users and applications can refer to the data on the mounted volume by the full (merged) filesystem path as if both volumes are one integral part. |
| <b>Volume Shadow Copy Service</b>        | See Microsoft Volume Shadow Copy Service (VSS).                                                                                                                                                                                                                                                                                                                 |
| <b>VSS</b>                               | See Microsoft Volume Shadow Copy Service (VSS).                                                                                                                                                                                                                                                                                                                 |



**VSS compliant mode** *(HP P9000 XP Disk Array Family VSS provider specific term)* One of two P9000 XP Array VSS hardware provider operation modes. When the P9000 XP Array provider is in the VSS compliant mode, the source volume (P-VOL) and its replica (S-VOL) are in simplex, unpaired state after a backup. Therefore the number of replicas (S-VOLs per a P-VOL) rotated is not limited. Restore from a backup in such a configuration is possible only by switching the disks.  
*See also* resync mode, source volume, primary volume (P-VOL), replica, secondary volume (S-VOL), and replica set rotation.

**VxFS** Veritas Journal Filesystem.

**VxVM (Veritas Volume Manager)** A Veritas Volume Manager is a system for managing disk space on Solaris platforms. A VxVM system consists of an arbitrary group of one or more physical volumes organized into logical disk groups.

## W

**Wake ONLAN** Remote power-up support for systems running in power-save mode from some other system on the same LAN.

**Web reporting** The Data Protector functionality that allows you to view reports on backup, object copy, and object consolidation status and Data Protector configuration using the Web interface.

**wildcard character** A keyboard character that can be used to represent one or many characters. The asterisk (\*), for example, typically represents one or more characters, and the question mark (?) typically represents a single character. Wildcard characters are often used in operating systems as a means of specifying more than one file by name.

**Windows configuration backup** Data Protector allows you to back up Windows CONFIGURATION, including Windows Registry, user profiles, Event Logs, and WINS and DHCP server data (if configured on a system) in one step.

**Windows Registry** A centralized database used by Windows to store configuration information for the operating system and the installed applications.

**WINS server** A system running Windows Internet Name Service software that resolves Windows networking computer names to IP addresses. Data Protector can back up WINS server data as part of the Windows configuration.

**writer** *(Microsoft VSS specific term)* A process that initiates change of data on the original volume. Writers are typically applications or system services that write persistent information on a volume. Writers also participate in the shadow copy synchronization process by assuring data consistency.

## X

**XBSA interface** *(Informix Server specific term)* ON-Bar and Data Protector communicate with each other through the X/Open Backup Services Application Programmer's Interface (XBSA).

## Z

**ZDB** *See* zero downtime backup (ZDB).

**ZDB database** *(ZDB specific term)* A part of the IDB, storing ZDB-related information such as source volumes, replicas, and security information. The ZDB database is used in zero downtime backup, instant recovery, and split mirror restore sessions.  
*See also* zero downtime backup (ZDB).

**ZDB to disk** *(ZDB specific term)* A form of zero downtime backup where the replica produced is kept on the disk array as a backup of the source volumes at a specific point in time. Multiple replicas, produced using the same backup specification at different times, can be kept in a replica set. A replica from ZDB to disk can be restored using the instant recovery process.  
*See also* zero downtime backup (ZDB), ZDB to tape, ZDB to disk+tape, instant recovery, and replica set rotation.

**ZDB to disk+tape** *(ZDB specific term)* A form of zero downtime backup where the replica produced is kept on the disk array as a backup of the source volumes at a specific point in time, in the same way as with ZDB to disk. However, data in the replica is also streamed to a backup medium, as with ZDB to tape. If this backup method is used, data backed up in the same session can be restored using



the instant recovery process, the standard Data Protector restore from tape, or with specific disk array families, split mirror restore.

See also zero downtime backup (ZDB), ZDB to disk, ZDB to tape, instant recovery, replica, and replica set rotation.

**ZDB to tape**

*(ZDB specific term)* A form of zero downtime backup where data in the replica produced is streamed to a backup medium, typically tape. Instant recovery is not possible from such a backup, so the replica need not be retained on the disk array after backup completion. The backed up data can be restored using standard Data Protector restore from tape. With specific disk array families, split mirror restore can also be used.

See also zero downtime backup (ZDB), ZDB to disk, ZDB to disk+tape, instant recovery, and replica.

**zero downtime backup (ZDB)**

A backup approach in which data replication techniques, provided by a disk array, are used to minimize the impact of backup operations on an application system. A replica of the data to be backed up is created first. All subsequent backup operations are performed on the replicated data rather than the original data, while the application system can be returned to normal operation.

See also ZDB to disk, ZDB to tape, ZDB to disk+tape, and instant recovery.

# Index

## A

- advantages
  - Oracle integration, 19
- architecture
  - MS Exchange 2010 Server integration, 201
- audience, 11

## B

- backing up MS Exchange 2010 Server, 205–217
  - backup specification, modifying, 215
  - backup specifications, creating, 209
  - backup types, 206
  - copy backups, 206
  - differential backups, 206
  - full backups, 206
  - incremental backups, 206
  - previewing backups, 216
  - scheduling backups, 215
  - scheduling backups, example, 215
  - starting backups, 216
- backing up MS Exchange Server 2003
  - backup options, 181
- backing up MS Exchange Server 2010
  - backup options, 214–215
- backing up MS SQL Server, 159, 166
  - backup options, 164
  - backup specifications, creating, 160
  - scheduling backups, 165
- backing up Oracle, 59–61
  - backup options, 55
  - backup set ZDB concepts, 24–27
  - backup set ZDB session flow, 26–27
  - backup specifications, creating, 48
  - backup types, 20
  - proxy-copy ZDB concepts, 28–30
  - proxy-copy ZDB session flow, 29–30
  - scheduling backups, 59
  - starting backups, 60–61
  - starting backups, using CLI, 61
  - starting backups, using GUI, 60
- backing up SAP R/3 , 123–136
  - backup options, 130
  - backup specification, modifying, 130
  - backup specifications, creating, 124
  - backup templates, 125
  - backup types, 104, 123
  - manual balancing, 130, 135
  - previewing backups, 131
  - SAP compliant ZDB, 134
  - scheduling backups, 130
  - scheduling backups, example, 130
  - starting backups, 132
  - verifying backups, 148
  - ZDB concepts, 105–107
  - ZDB flow, 105

- backup flow
  - SAP R/3 integration, 105
- backup flow, Oracle integration, 22–23
  - backup set ZDB session flow, 26–27
  - proxy-copy ZDB session flow, 29–30
- backup options
  - MS Exchange Server 2003 integration, 181
  - MS Exchange Server 2010 integration, 214–215
  - MS SQL Server integration, 164
  - Oracle integration, 55
  - SAP R/3 integration, 130
- backup specifications, creating
  - MS Exchange 2010 Server integration, 209
  - MS SQL Server integration, 160
  - Oracle integration, 48
  - SAP R/3 integration, 124
- backup specifications, modifying
  - MS Exchange 2010 Server integration, 215
  - MS Exchange Server 2003 integration, 186
  - SAP R/3 integration, 130
- backup specifications, scheduling
  - MS Exchange 2010 Server integration, 215
  - MS SQL Server integration, 165
  - Oracle integration, 59
- backup templates
  - SAP R/3 integration, 125
- backup types
  - MS Exchange 2010 Server integration, 206
  - Oracle integration, 20
  - SAP R/3 integration, 104, 123
- BRBACKUP, 106

## C

- checking configuration
  - MS SQL Server integration, 158
  - Oracle integration, 44
  - SAP R/3 integration, 121
- checking database files for consistency
  - MS Exchange Server 2003 integration, 187
- concepts
  - MS Exchange 2010 Server integration, 201
  - MS Exchange Server 2003 integration, 178
  - MS SQL Server integration, 153–154
  - Oracle integration, 21
  - Oracle integration, backup set ZDB concepts, 24–27
  - Oracle integration, proxy-copy ZDB concepts, 28–30
- concepts, ZDB
  - SAP R/3 integration, 105–107
- configuration file
  - SAP R/3 integration, 107
- configuration files
  - MS SQL Server integration, 155
- configuring MS Exchange 2010 Server, 203–205
- configuring MS Exchange Server 2003, 179
- configuring MS SQL Server, 154, 159
  - checking configuration, 158

- configuration files, 155
- configuring Oracle, 31–47
  - backup methods, 47
  - checking configuration, 44
  - example, CLI, 43
  - prerequisites, 33
- configuring SAP R/3, 107–123
  - authentication modes, 117
  - checking configuration, 121
  - configuration file, 107
  - listener, configuring, 113
  - SAP R/3 parameter file, 123
  - verifying configuration, 147
- control files, Oracle integration
  - restore, 65
- conventions
  - document, 16
- copy backups
  - MS Exchange 2010 Server integration, 206
- creating backup specifications
  - MS Exchange 2010 Server integration, 209
  - MS SQL Server integration, 160
  - Oracle integration, 48
  - SAP R/3 integration, 124
- D**
- Data Guard, Oracle integration
  - limitations, 33
- database recovery
  - Oracle integration, after instant recovery, 85
  - Oracle integration, options, 71
  - SAP R/3 integration, 140
- database recovery options
  - SAP R/3 integration, 141
- differential backups
  - MS Exchange 2010 Server integration, 206
- disaster recovery
  - Oracle integration, 63
- document
  - conventions, 16
  - related documentation, 11
- documentation
  - HP website, 11
  - providing feedback, 18
- E**
- examples
  - SAP R/3 integration, starting interactive backups, 132
- examples, MS Exchange 2010 Server integration
  - scheduling backups, 215
  - starting interactive backups, 217
- examples, MS Exchange Server 2003 integration
  - creating backup specifications, 185
- examples, Oracle integration
  - restoring using RMAN, 74
- examples, scheduling backups
  - SAP R/3 integration, 130

- F**
- full backups
  - MS Exchange 2010 Server integration, 206
- H**
- help
  - obtaining, 17
- HP
  - technical support, 17
- I**
- incremental backups
  - MS Exchange 2010 Server integration, 206
- Informix backup
  - backup specifications, creating, 124
- instant recovery
  - MS Exchange Server 2003 integration, 195–196
  - MS SQL Server integration, 170, 172
  - Oracle integration, 85, 89
  - Oracle integration, database recovery after, 87
  - RAC preparation steps, 85
  - reconfiguring an Oracle instance, 269
  - SAP R/3 integration, 140
- instant recovery options
  - SAP R/3 integration, 141
- integrated authentication, MS SQL Server integration, 157
- interactive backups
  - MS Exchange 2010 Server integration, 216
  - Oracle integration, 60
  - SAP R/3 integration, 132
- introduction
  - MS Exchange 2010 Server integration, 200
  - MS Exchange Server 2003 integration, 178
  - MS SQL Server integration, 153
  - Oracle integration, 19
  - SAP R/3 integration, 104
- M**
- manual balancing
  - SAP R/3 integration, 130, 135
- Media Management Library *see* MML
- MML (Data Protector Media Management Library)
  - linking with Oracle, UNIX, 36
- modifying backup specifications
  - MS Exchange 2010 Server integration, 215
  - MS Exchange Server 2003 integration, 186
  - SAP R/3 integration, 130
- monitoring sessions
  - MS Exchange 2010 Server integration, 242
  - MS SQL Server integration, 172
  - SAP R/3 integration, 144
- MS Exchange 2010 Server backup, 205–217
  - backup specification, modifying, 215
  - backup specifications, creating, 209
  - backup types, 206
  - copy backups, 206
  - differential backups, 206
  - full backups, 206
  - incremental backups, 206

- previewing backups, 216
  - scheduling backups, 215
  - scheduling backups, example, 215
  - starting backups, 216
- MS Exchange 2010 Server configuration, 203–205
- MS Exchange 2010 Server integration
  - architecture, 201
  - backup, 205–217
  - concepts, 201
  - configuration, 203–205
  - introduction, 200
  - monitoring sessions, 242
  - restore, 218–242
  - troubleshooting, 243–245
- MS Exchange 2010 Server restore, 218–242
  - finding information, 222
  - restore options, 242
  - using another device, 232
  - using CLI, 229
  - using GUI, 223
- MS Exchange 2010 Server troubleshooting, 243–245
- MS Exchange Server 2003 backup
  - backup options, 181
- MS Exchange Server 2003 configuration, 179–180
- MS Exchange Server 2003 integration
  - concepts, 178
  - configuration, 179–180
  - introduction, 178
  - restore, 189
  - troubleshooting, 197–199
- MS Exchange Server 2003 restore, 189
  - instant recovery, 195–196
  - point in time recovery, 189–190
  - rollforward recovery, 190, 192
  - to the application system, 189
- MS Exchange Server 2003 troubleshooting, 197–199
- MS Exchange Server 2010 backup
  - backup options, 214–215
- MS SQL Server backup, 159, 166
  - backup options, 164
  - backup specifications, creating, 160
  - scheduling backups, 165
- MS SQL Server configuration, 154, 159
  - checking configuration, 158
  - configuration files, 155
- MS SQL Server integration
  - backup, 159, 166
  - concepts, 153–154
  - configuration, 154, 159
  - introduction, 153
  - monitoring sessions, 172
  - restore, 166, 172
  - troubleshooting, 172, 176
- MS SQL Server restore, 166, 172
  - instant recovery, 170, 172
  - restore options, 168–169
  - restoring, tail log backup, 166
- MS SQL Server troubleshooting, 172, 176

O

- OB2\_MIRROR\_COMP, SAP R/3 integration, 134
- OB2RMANSAVE, Oracle integration, 100
- omnirc variables, 271–275
- online backups
  - MS Exchange 2010 Server integration, 200
- Oracle backup, 59–61
  - backup concepts, scheme, 24
  - backup set ZDB concepts, 24–27
  - backup set ZDB session flow, 26–27
  - backup specifications, creating, 48
  - backup types, 20
  - proxy-copy concepts, 28–30
  - proxy-copy ZDB session flow, 29–30
  - scheduling backups, 59
  - starting backups, 60–61
  - starting backups, using CLI, 61
  - starting backups, using GUI, 60
- Oracle configuration
  - backup methods, 47
  - checking configuration, 44
  - example, CLI, 43
  - prerequisites, 33
- Oracle integration
  - advantages, 19
  - backup, 59–61
  - backup set ZDB concepts, 24–27
  - concepts, 21
  - configuration, 31–47
  - instant recovery, 85–89
  - introduction, 19
  - proxy-copy ZDB concepts, 28–30
  - troubleshooting, 89–103
- Oracle restore
  - control files, 65
  - database items, 61
  - database objects, 66
  - database recovery after instant recovery, 85
  - editing RMAN scripts, 100
  - examples, using RMAN, 74
  - instant recovery, 85–89
  - preparing databases for restore, 75
  - recovery catalog, 64
  - restorable items, 61
  - restore flow, 23
  - restore options, 71
  - restore types, 20
  - standard restore procedure, 61–84
  - tablespaces and datafiles, 68
  - to the application system, 63
  - using another device, 84
  - using GUI, 63
  - using RMAN, 74
- Oracle RMAN script, 56
- Oracle troubleshooting, 89–103

## P

- prerequisites, verifying Oracle side
  - SAP R/3 integration, 145

prerequisites, verifying SAP side

SAP R/3 integration, 146

previewing backups

MS Exchange 2010 Server integration, 216

SAP R/3 integration, 131

## R

RAC

preparing for instant recovery, 85

RAC, configuring Oracle Servers

on HP-UX, 36

on other UNIX systems, 36

reconfiguring an Oracle instance for instant recovery , 269

recovery

Oracle database after instant recovery, 85

Oracle integration, options, 71

SAP R/3 integration, 140

recovery catalog, Oracle integration

restore, 64

recovery options

SAP R/3 integration, 141

related documentation, 11

restore methods

SAP R/3 integration, 136

restore options

Informix integration, 169

MS Exchange 2010 Server integration, 242

MS SQL Server integration, 168

restore types

Oracle integration, 20

restoring Informix

restore options, 169

restoring MS Exchange 2010 Server, 218–242

finding information, 222

restore options, 242

using another device, 232

using CLI, 229

using GUI, 223

restoring MS Exchange Server 2003, 189–196

instant recovery, 195–196

point in time recovery, 189–190

rollforward recovery, 190, 192

to the application system, 189

restoring MS SQL Server, 166, 172

instant recovery, 170, 172

restore options, 168

restoring, tail log backup, 166

restoring Oracle

control files, 65

database objects, 66

database recovery after instant recovery, 85

editing RMAN scripts, 100

instant recovery, 85–89

recovery catalog, 64

restore flow, 23

restore types, 20

standard restore procedure, 61–84

tablespaces and datafiles, 68

to the application system, 63

using another device, 84

using GUI, 63

using RMAN, 74

restoring SAP R/3, 136–142

database recovery, 140

instant recovery, 140

restore methods, 136

standard restore, 137

using another device, 144

verifying restores, 148

RMAN, Oracle integration

restore, 74

running backups see starting backups

## S

SAP compliant ZDB

SAP R/3 integration, 134

SAP R/3 backup, 123–136

backup options, 130

backup specification, modifying, 130

backup templates, 125

backup types, 104, 123

manual balancing, 130, 135

previewing backups, 131

SAP compliant ZDB, 134

scheduling backups, 130

scheduling backups, example, 130

starting backups, 132

verifying backups, 148

ZDB concepts, 105–107

ZDB flow, 105

SAP R/3 configuration, 107–123

authentication modes, 117

checking configuration, 121

configuration file, 107

listener, configuring, 113

SAP R/3 parameter file, 123

verifying configuration, 147

SAP R/3 integration

backup, 123–136

configuration, 107–123

introduction, 104

monitoring sessions, 144

restore, 136–142

troubleshooting, 144–152

ZDB concepts, 105–107

SAP R/3 restore, 136–142

database recovery, 140

instant recovery, 140

restore methods, 136

standard restore, 137

using another device, 144

verifying restores, 148

SAP R/3 troubleshooting, 144–152

verifying backups, 148

verifying configuration, 147

verifying prerequisites on Oracle side, 145

verifying prerequisites on SAP side, 146

- verifying restores, [148](#)
- SBT\_LIBRARY, Oracle integration, [36](#), [76](#)
- scheduling backups
  - MS Exchange 2010 Server integration, [215](#)
  - MS SQL Server integration, [165](#)
  - Oracle integration, [59](#)
  - SAP R/3 integration, [130](#)
- SQL Server authentication, MS SQL Server integration, [157](#)
- starting backups
  - MS Exchange 2010 Server integration, [216](#)
  - SAP R/3 integration, [132](#)
- starting backups, Oracle integration, [60–61](#)
  - using CLI, [61](#)
  - using GUI, [60](#)
- Subscriber's Choice, HP, [17](#)

## T

- technical support
  - HP, [17](#)
  - service locator website, [17](#)
- troubleshooting MS Exchange 2010 Server, [243–245](#)
- troubleshooting MS Exchange Server 2003, [197–199](#)
- troubleshooting MS SQL Server, [172](#), [176](#)
- troubleshooting Oracle, [89–103](#)
- troubleshooting SAP R/3 , [144–152](#)
  - verifying backups, [148](#)
  - verifying configuration, [147](#)
  - verifying prerequisites on Oracle side, [145](#)
  - verifying prerequisites on SAP side, [146](#)
  - verifying restores, [148](#)

## U

- users, configuring
  - Oracle integration, [36](#)

## V

- verifying backups
  - SAP R/3 integration, [148](#)
- verifying configuration
  - SAP R/3 integration, [147](#)
- verifying prerequisites, Oracle side
  - SAP R/3 integration, [145](#)
- verifying prerequisites, SAP side
  - SAP R/3 integration, [146](#)
- verifying restores
  - SAP R/3 integration, [148](#)

## W

- websites
  - HP , [17](#)
  - HP Subscriber's Choice for Business, [17](#)
  - product manuals, [11](#)
- Windows authentication, MS SQL Server integration, [157](#)

## Z

- ZDB flow
  - SAP R/3 integration, [105](#)
- ZDB session flow

- Oracle integration, backup set, [26–27](#)
- Oracle integration, proxy-copy, [29–30](#)