

HP Data Protector A.06.10

Integration guide for Oracle and SAP

[Build 500]



T B D

Part number: TBD
First edition: TBD Month 2008



i n v e n t

Legal and notice information

© Copyright 2004, 2008 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 US trademark of Sun Microsystems, Inc.

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

UNIX is a registered trademark of The Open Group.

Printed in the US

Contents

Publication history	13
About this guide	15
Intended audience	15
Documentation set	15
Guides	15
Online help	18
Documentation map	19
Abbreviations	19
Map	20
Integrations	21
Document conventions and symbols	22
Data Protector graphical user interface	23
General Information	24
HP technical support	24
Subscription service	25
HP websites	25
Documentation feedback	25
1 Integrating Oracle and Data Protector	27
Introduction	27
Integration concepts	28
Configuring the integration	33
Prerequisites	33
Limitations	34
Before you begin	35
Cluster-aware clients	35
Linking Oracle Server with the Data Protector MML	36
Linking on UNIX systems	36
Linking on Windows systems	38
Linking on OpenVMS systems	38
Configuring Oracle users on UNIX and OpenVMS	44
Configuring Oracle databases	47

Using the Data Protector GUI	47
Using the Data Protector CLI	51
What happens after the configuration?	54
Checking the configuration	55
Using the Data Protector GUI	55
Using the Data Protector CLI	55
Backup	56
Creating new templates	57
Creating backup specifications	58
Examples of pre-exec and post-exec scripts on UNIX	67
Editing the Oracle RMAN script	68
Creating copies of backed up objects	71
Testing the integration	72
Testing using the Data Protector GUI	73
Testing using the CLI	74
Starting backup sessions	75
Scheduling backup specifications	78
Running an interactive backup	80
Starting a backup using the GUI	80
Starting a backup using the CLI	81
Starting Oracle backup using RMAN	82
Examples of the RMAN scripts	87
Restore	91
Prerequisites	93
Restoring Oracle using the Data Protector GUI	93
Restoring database items in a disaster recovery	94
Changing the database state	94
Restoring the recovery catalog database	95
Restoring the control file	97
Restoring Oracle database objects	99
Restoring tablespaces and datafiles	104
Restoring and recovering an Oracle database in Oracle Data Guard environment	105
Restoring and recovering a primary database	105
Restoring and recovering a standby database	105
Duplicating an Oracle database	106
Restore, recovery, and duplicate options	110
Restore action options	110
General options	110
Duplicate options	112
Restore and recovery options	112
Restoring Oracle using RMAN	114
Preparing the Oracle database for restore	115

Connection strings used in the examples	117
SBT_LIBRARY parameter	117
Example of full database restore and recovery	117
Example of point-in-time restore	119
Example of tablespace restore and recovery	121
Example of datafile restore and recovery	124
Example of archive log restore	128
Restoring Oracle using CLI	130
Restoring the recovery catalog	130
Restoring using another device	131
Disaster recovery	132
Monitoring sessions	133
Monitoring current sessions	133
Viewing previous sessions	134
Using Oracle after removing the Data Protector Oracle integration	134
Removing the link on HP-UX systems	135
Removing the link on Solaris and other UNIX systems	135
Removing the link on OpenVMS systems	135
Removing the link on Windows systems	136
Oracle RMAN metadata and Data Protector Media Management Database synchronization	136
Troubleshooting	137
Before you begin	137
Checks and verifications	137
Problems	146

2 Integrating SAP R/3 and Data Protector 153

Introduction	153
Integration concepts	154
Backup flow	158
Restore flow	160
Data Protector SAP R/3 configuration file	160
Setting, retrieving, listing, and deleting Data Protector SAP R/3 configuration file parameters using the CLI	163
Configuring the integration	166
Prerequisites	167
Before you begin	167
Cluster-aware clients	168
Configuring user accounts	168
Checking the connection	169
Authentication password file	170
Enabling archived logging	170

Linking Oracle Server with the Data Protector MML	172
Choosing authentication mode	173
Configuring SAP R/3 databases	173
Before you begin	174
Using the Data Protector GUI	174
Using the Data Protector CLI	177
Checking the configuration	179
Using the Data Protector GUI	179
Using the Data Protector CLI	180
Backup	181
Considerations	183
Creating backup specifications	183
Modifying backup specifications	189
Scheduling backup specifications	189
Scheduling example	189
Previewing backup sessions	190
Using the Data Protector GUI	190
Using the Data Protector CLI	190
What happens during the preview?	191
Starting backup sessions	191
Backup methods	191
Using the Data Protector GUI	191
Using the Data Protector CLI	192
Using the SAP BRTOOLS	192
Configuring SAP compliant ZDB sessions	193
Using the Data Protector GUI	193
Using the Data Protector CLI	194
Backing up using Oracle Recovery Manager	195
Manual balancing	196
Restore	196
Considerations	197
Restoring using the Data Protector GUI	197
Restoring using the Data Protector CLI	200
Restoring using the SAP commands	201
Restoring using another device	201
Using the Data Protector GUI	201
Using the Data Protector CLI or SAP commands	202
Localized SAP R/3 objects	202
Sparse files	203
Disaster recovery	203
Restoring the control file	203
Monitoring sessions	204
Troubleshooting	204

Before you begin	204
General troubleshooting	205
Troubleshooting on Windows systems	205
Prerequisites concerning the Oracle side of the integration	205
Prerequisites on the SAP side of the integration	208
Configuration problems	209
Backup problems	212
Restore problems	214
Troubleshooting on UNIX systems	217
Using Oracle after removing the Data Protector Oracle integration	217
Prerequisites concerning the Oracle side of the integration	217
Prerequisites on the SAP side of the integration	221
Configuration problems	223
Backup problems	226
Restore problems	229

3 Integrating SAP DB/MaxDB and Data Protector 235

Introduction	235
Integration concepts	236
Backup flow	238
Restore flow	238
Configuring the integration	239
Prerequisites	239
Limitations	239
Before you begin	239
Cluster-aware clients	240
Configuring SAP DB users	240
Configuring SAP DB instances	240
Before you begin	241
Using the Data Protector GUI	241
Using the Data Protector CLI	243
Handling errors	245
Checking the configuration	245
Using the Data Protector GUI	245
Using the Data Protector CLI	245
Backup	246
Creating backup specifications	246
Modifying backup specifications	249
Scheduling backup specifications	249
Scheduling example	250
Previewing backup sessions	250
Using the Data Protector GUI	251

Using the Data Protector CLI	251
What happens during the preview?	251
Starting backup sessions	252
Backup methods	252
Using the Data Protector GUI	252
Using the Data Protector CLI	252
Using SAP DB utilities	253
Restore	256
Restore and recovery overview	256
Before you begin	260
Restoring using the Data Protector GUI	260
Restoring using the Data Protector CLI	263
Restoring using SAP DB utilities	264
Finding information needed for restore	264
SAP DB restore and recovery	265
SAP DB migration	269
SAP DB restore options	270
Restoring using another device	274
Monitoring sessions	274
Troubleshooting	274
Before you begin	274
Problems	275
SAP DB cluster-related troubleshooting	277

Glossary 279

Index 335

Figures

1 Data Protector graphical user interface	24
2 Data Protector Oracle integration concept	32
3 Finding the Oracle user	44
4 Configuring Oracle - General (Windows)	48
5 Configuring Oracle - General (UNIX)	48
6 Configuring Oracle - Primary	49
7 Configuring Oracle - Catalog	50
8 Configuring Oracle - Standby	51
9 Specifying an Oracle Server system (UNIX)	62
10 Oracle specific options	64
11 Saving the backup specification	65
12 Previewing a backup	74
13 Scheduling backups	80
14 Starting an interactive backup	81
15 Taking the Oracle resource group offline	92
16 Checking properties	93
17 Recovery catalog settings dialog	96
18 Source page	100
19 Options page	102
20 Devices page	103
21 Oracle duplicate options	109
22 Checking the status of the Oracle listener	139
23 Output of the ldd command on other UNIX systems:	143
24 SAP R/3 architecture	155

25	SAP R/3 architecture: backint mode	156
26	SAP R/3 architecture: RMAN mode	158
27	Specifying an SAP R/3 system and Oracle instance	175
28	Configuring an SAP R/3 database on a UNIX system (operating system authentication mode)	176
29	Configuring an SAP R/3 database on a Window system (database authentication mode)	177
30	Checking the SAP R/3 configuration	180
31	Selecting backup objects	185
32	Application specific options	186
33	Scheduling backups	190
34	Setting environmental variables	194
35	Selecting objects for restore	198
36	Selecting the target client	199
37	Checking the status of the Oracle listener	207
38	Checking the lnet start-up parameters	211
39	SAP DB integration architecture	237
40	Specifying an SAP DB instance	242
41	SAP DB configuration	243
42	Selecting SAP DB objects	247
43	Application specific options	248
44	Scheduling the backup specification	250
45	SAP DB restore process	259
46	SAP DB archive logs restore process—redo logs details	259
47	Selecting objects for restore	261
48	Properties for data	262
49	SAP DB restore and recovery options	270

Tables

1	Edition history	13
2	Document conventions	22
3	Filenames for the MML on different platforms	37
4	Oracle backup options	66
5	MML filenames on different platforms	85
6	Required database states	94
7	Backup types	153
8	SAP backup and restore utilities	154
9	Backup types	181
10	What is backed up	181
11	Two alternatives of specifying backup options	182
12	Backup templates	184
13	SAP R/3 backup options	187
14	Backup types	235
15	What is backed up	246
16	SAP DB backup options	249

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
B6960-90109	October 2004	Data Protector Release A.05.50
B6960-96008	July 2006	Data Protector Release A.06.00
TBD	TBD	Data Protector Release A.06.10

About this guide

This guide describes how to configure and use Data Protector with Oracle, SAP R/3, and SAP DB/MaxDB.

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.

Documentation set

Other documents and online Help provide related information.

Guides

Data Protector guides are available in printed format and in PDF format. Install the PDF files during the Data Protector setup procedure by selecting the *English documentation and 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 Business Support Center website:

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

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 four guides:

- *HP Data Protector integration guide for Microsoft applications: SQL Server, SharePoint Portal Server, Exchange Server, and Volume Shadow Copy Service*

This guide describes the integrations of Data Protector with the following Microsoft applications: Microsoft Exchange Server, Microsoft SQL Server, and Volume Shadow Copy Service.

- *HP Data Protector integration guide for Oracle and SAP*

This guide describes the integrations of Data Protector with Oracle, SAP R3, and SAP DB.

- *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, and Lotus Notes/Domino Server.

- *HP Data Protector integration guide for Sybase, Network Node Manager, Network Data Management Protocol, and VMware*

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

- *HP Data Protector integration guide for HP Service Information Portal*

This guide describes how to install, configure, and use the integration of Data Protector with HP Service Information Portal. It is intended for backup administrators. It discusses how to use the OpenView applications for Data Protector service management.

- *HP Data Protector integration guide for HP Reporter*

This manual describes how to install, configure, and use the integration of Data Protector with HP Reporter software. It is intended for backup administrators. It discusses how to use the applications for Data Protector service management.

- *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 software 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 software and HP Service Navigator on Windows.

There are two versions of the guide:

- for OVO 7.1x, 7.2x
- for OVO 7.5

- *HP Data Protector software integration guide for HP Performance Manager software and HP Performance Agent software*

This guide provides information about how to monitor and manage the health and performance of the Data Protector environment with HP Performance Manager (PM) software and HP Performance Agent (PA) software on Windows, HP-UX, Solaris and Linux.

- *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 StorageWorks Virtual Array, HP StorageWorks Enterprise Virtual Array, EMC Symmetrix Remote Data Facility and TimeFinder, and HP StorageWorks Disk Array XP. 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, SAP R/3, Microsoft Exchange Server, and Microsoft SQL Server databases. The guide also describes how to configure and use Data Protector to perform backup and restore using the Microsoft Volume Shadow Copy Service.
- *HP Data Protector MPE/iX system user guide*
This guide describes how to configure MPE/iX clients and how to back up and restore MPE/iX data.
- *HP Data Protector Media Operations user guide*
This guide provides tracking and management of offline storage media. It is intended for network administrators responsible for maintaining and backing up systems. 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 A.06.10. It also provides information on supported configurations (devices, platforms and online database integrations, SAN, and ZDB), required patches, and limitations, as well as known problems and workarounds. An updated version of the supported configurations is available at <http://www.hp.com/support/manuals>
There are also four other *Product announcements, software notes and references*, which serve a similar purpose for the following:
 - HP Operations Manager software software UNIX integration
 - HP Operations Manager software software Windows integration
 - HP Service Information Portal and HP Reporter software
 - HP Performance Manager software and HP Performance Agent integration
 - HP Media Operations

Online help

Data Protector provides context-sensitive (F1) Help and Help Topics for Windows and UNIX platforms.

You can access the online help from the top-level directory on the installation DVD without installing Data Protector:

- **Windows:** Unzip `DP_help.zip` and 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
Help	Online Help
IG-IBM	Integration guide—IBM applications
IG-MS	Integration guide—Microsoft applications
IG-O/S	Integration guide—Oracle, SAP R/3, and SAP DB/MaxDB
IG-OMU	Integration guide—HP Operations Manager software software, UNIX
IG-OMW	Integration guide—HP Operations Manager software software, Windows
IG-PM/PA	Integration guide—Performance Manager and Performance Agent software
IG-Report	Integration guide—HP Reporter software
IG-SIP	Integration guide—HP Service Information Portal
IG-Var	Integration guide—Sybase, Network Node Manager, NDMP and VMware
Install	Installation and licensing guide
MO GS	Media Operations getting started guide

Abbreviation	Guide
MO RN	Media Operations product announcements, software notes, and references
MO UG	Media Operations user guide
MPE/iX	MPE/iX system user guide
PA	Product announcements, software notes, and references
Trouble	Troubleshooting guide
ZDB Admin	ZDB administrator's guide
ZDB Concpt	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.

Integrations

Look in these guides for details of the following integrations:

Integration	Guide
HP Operations Manager software software	IG-OMU, IG-OMW
HP Performance Manager software	IG-PM/PA
HP Performance Agent software	IG-PM/PA
HP Reporter Light	IG-OMW
HP Reporter software	IG-R
HP Service Information Portal	IG-SIP
HP StorageWorks Disk Array XP	all ZDB
HP StorageWorks Enterprise Virtual Array (EVA)	all ZDB
HP StorageWorks Virtual Array (VA)	all ZDB
IBM DB2 UDB	IG-IBM
Informix	IG-IBM
Lotus Notes/Domino	IG-IBM
Media Operations	MO User
MPE/iX System	MPE/iX
Microsoft Exchange Server	IG-MS, ZDB IG
Microsoft Exchange Single Mailbox	IG-MS
Microsoft SQL Server	IG-MS, ZDB IG
Microsoft Volume Shadow Copy Service (VSS)	IG-MS, ZDB IG
NDMP Server	IG-Var
Network Node Manager (NNM)	IG-Var

Integration	Guide
Oracle	IG-O/S
Oracle ZDB	ZDB IG
SAP DB	IG-O/S
SAP R/3	IG-O/S, ZDB IG
Sybase	IG-Var
Symmetrix (EMC)	all ZDB
VMware	IG-Var

Document conventions and symbols

Table 2 Document conventions

Convention	Element
Blue text: Table 2 on page 22	Cross-reference links and e-mail addresses
Blue, underlined text: http://www.hp.com	website addresses
<i>Italic text</i>	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 text</i>	<ul style="list-style-type: none"> • Code variables • Command variables
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 or the Data Protector Java GUI. Refer to the online Help for information about the Data Protector graphical user interface.




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/service_locator
- <http://www.hp.com/support/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 AppRM.DocFeedback@hp.com. All submissions become the property of HP.

1 Integrating Oracle and Data Protector

Introduction

Data Protector offers offline as well as online backup of the Oracle Server instances. To enable database recovery from an online backup, the respective Oracle Server instance must operate in the ARCHIVELOG mode.

The online backup concept is widely accepted. It addresses the business requirements for high application availability, as opposed to the offline concept. During an online backup, a database remains available for use, while during an offline backup, the database cannot be used by an application.

Backup types

Using the Data Protector Oracle integration, you can perform the following types of backups:

- Online backup of a whole database or parts of it
- Online incremental backup (*Oracle* differential incremental backup 1 to 4)
- Offline backup of a whole database
- Backup of Archived Redo Logs only
- Backup of the Oracle database recovery catalog
- Backup of the Oracle control files
- With Oracle 10g, backup of **recovery files** residing in the **flash recovery area**.

The following recovery files in the flash recovery area are backed up:

- full and incremental backup sets
- control file autobackup (SPFILE included if used)
- archived redo logs
- datafile copies, control file copies

Flashback logs, the current control file, and online redo logs are not backed up.

- In **Oracle Data Guard** environment, backup of **standby database**.

Restore types

Using the Data Protector Oracle integration, you can restore the following:

- The whole database or parts of it
- The database to a specific point in time
- From incremental backup
- To a host other than the one where the database originally resided
- A datafile to a location other than its original one
- A catalog before restoring the database
- From a chain of incremental backups

Duplicating a database

Using the Data Protector Oracle integration, you can perform duplication of a production database.

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:

- 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.

Oracle backup types handled by the integration

Using this integration, you can perform the *Oracle* full and incremental (up to incremental level 4) backup types.

With Oracle full and incremental level 0 backups all data blocks per datafile are backed up. With Oracle incremental backup (level 1 or higher), only the data blocks that have changed since a previous backup are backed up.

The difference between a full backup and an incremental level 0 backup is that the incremental 0 is a base for subsequent incremental backups. Therefore, Data Protector always performs Oracle incremental 0 when you select the full backup type in a backup specification.

The full backup type is not related to the number of datafiles included in the backup, and can therefore be performed per single datafile. The data being backed up,

regardless of the backup type (full or incremental), is selected and controlled by Oracle.

Oracle incremental backups can be differential or cumulative. By default, Data Protector performs **Oracle differential incremental** backups. By changing the default RMAN script created by Data Protector, you can specify also a cumulative backup. For information on differential and cumulative Oracle backups, see the *Oracle Recovery Manager User's Guide*.

 **NOTE:**

Regardless of the Oracle backup type specified, Data Protector always marks the Oracle backups as full in the Data Protector database, since the Data Protector incremental backup concept is different from the Oracle incremental backup concept.

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, the RMAN utility, or the Oracle Enterprise Manager 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.

Backup 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 a specific user. This user must

be defined as the owner of the Data Protector Oracle 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
- Oracle Enterprise Manager GUI

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 Figure 2 on page 32, and the related terms are explained in the following table.

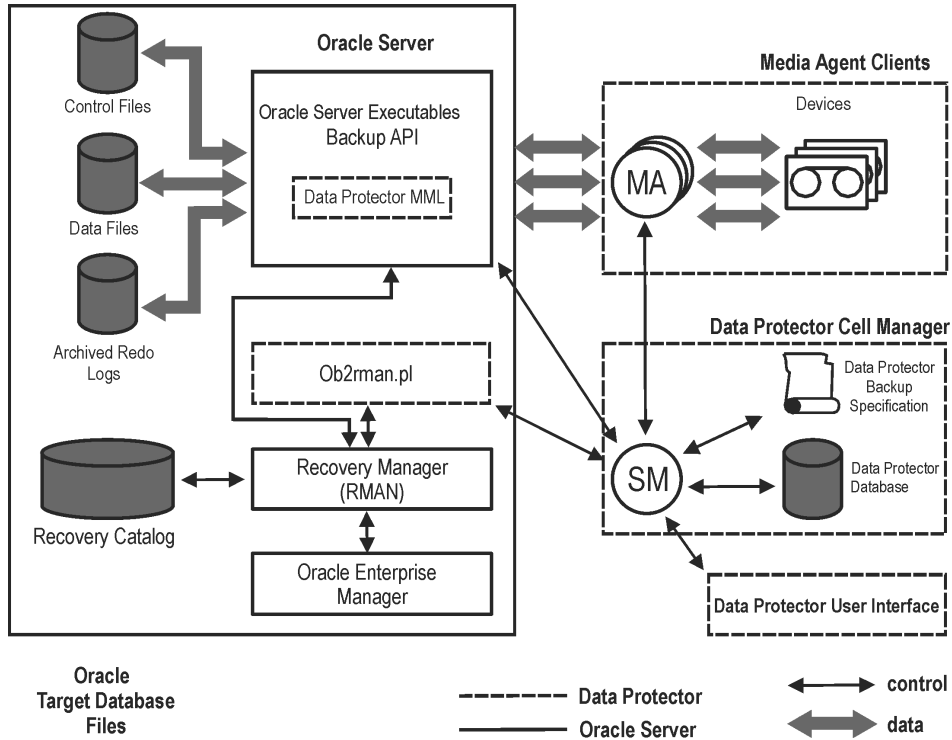


Figure 2 Data Protector Oracle integration concept

Oracle 10g database files can also be part of ASM configuration. They can reside in the flash recovery area.

Legend

<i>SM</i>	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.
<i>RMAN</i>	The Oracle Recovery Manager.
<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.
<i>Backup API</i>	The Oracle-defined application programming interface.
<i>IDB</i>	The IDB where all the information about Data Protector sessions, including session messages, objects, data, used devices, and media is written.
<i>MA</i>	The Data Protector General Media Agent, which reads and writes data from and to media devices.

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 Oracle integration. See the *HP Data Protector installation and licensing guide* for information about licensing.
- Before you begin, ensure that you have correctly installed and configured the Oracle Server and Data Protector systems. See the:
 - *HP Data Protector product announcements, software notes, and references* or <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 the Data Protector Oracle integration.
 - *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 about backup and recovery with the Oracle Enterprise Manager, as well as information about SQL*Plus.
- The Oracle Server software must be installed and the Oracle target database must be open or mounted.

- If the Oracle recovery catalog database is used, ensure that it is properly configured and open.
- Oracle net services must be properly configured and running for the Oracle target database and the recovery catalog, if you use it.
See the *Oracle Recovery Manager User's Guide and References* for more information about different connection options.
See "[Troubleshooting](#)" on page 137 for details about how to check the prerequisites listed above.
- On Windows, if the Oracle target database and the Oracle recovery catalog are installed on two different systems, the `Data Protector Inet` service account on the system with the Oracle target database installed must be configured as a *domain* account that is a member of the Administrators group on both systems. For information on how to change the `Data Protector Inet` service account, see the online Help index: "changing Data Protector Inet account".
- On OpenVMS, check the network alias names of the client. It is recommended to provide the full client name (together with the alias) to avoid non-detection of the Data Protector Oracle Integration agent.
- To successfully back up the recovery files residing in the flash recovery area (Oracle 10g only), ensure that you have correctly configured the flash recovery area.
- 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](#) on page 148.

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 Data Protector Oracle integration does not support the RMAN disk backup of a target database *to* the flash recovery area. The Data Protector Oracle integration supports only backups *from* the flash recovery area to a backup device. However, you can create an RMAN script that backs up the target database to the flash recovery area before or after the Data Protector backs up files from the flash recovery area to a backup device. The script can be set up using the `Pre-exec` or `Post-exec` option when creating a backup specification.
- On an OpenVMS client, you can only configure a Data Protector `admin` user with the username `<Any>` and the group name `<Any>`. This limitation is due to the lack of the user group name concept on OpenVMS.

- **Oracle Data Guard:**
 - You cannot configure only a standby database (without configuring primary database).
 - Only physical standby database backup is supported.
 - Recovery catalog database is required for standby configurations.
 - The Oracle database identifier (DBID) must be unique for all databases within a Data Protector cell.
 - For other limitations regarding RMAN backup, restore, recovery, and duplication in Oracle Data Guard environment, see the Oracle documentation.

Before you begin

- Configure devices and media for use with Data Protector.
- 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.

Cluster-aware clients

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: `set OB2BARHOSTNAME=virtual_server_name`

UNIX: `export OB2BARHOSTNAME=virtual_server_name`

RAC: Configure an Oracle database on every node from where you want to run backups and restores.

HP-UX with RAC: If you want to use virtual hostname, 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 client on which an Oracle instance is running.

- On Windows and UNIX clients with Oracle8i, and on OpenVMS clients with any Oracle version, link Oracle Server with the Data Protector MML manually. For details, see the following sections.

IMPORTANT:

After uninstalling the Data Protector Oracle integration on an Oracle server system, the Oracle server software is still linked to MML. You must re-link the Oracle binary to remove this link. If this is not done, the Oracle server cannot be started after the integration has been removed. See [“Using Oracle after removing the Data Protector Oracle integration”](#) on page 134 for information on removing the integration link.

- On Windows and UNIX clients with Oracle9i/10g, 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.

MC/ServiceGuard: When linking Oracle with MML, link it on all nodes.

Linking on UNIX systems

On Oracle Server systems, MML is located in the directory:

HP-UX, Solaris, and Linux: `/opt/omni/lib`

Other UNIX: `/usr/omni/lib`

The filename for MML depends on the platform:

Table 3 Filenames for the MML on different platforms

Platforms	32-bit	64-bit
HP-UX	libob2oracle8.sl	libob2oracle8_64bit.sl
HP-UX on IA-64	libob2oracle8.so	libob2oracle8_64bit.so
Solaris	libob2oracle8.so	libob2oracle8_64bit.so
AIX	libob2oracle8.a	libob2oracle8_64bit.a
Other UNIX	libob2oracle8.so	libob2oracle8_64bit.so

Proceed as follows:

1. Change to the `ORACLE_HOME/lib` directory:
32-bit Oracle: `cd ORACLE_HOME/lib`
64-bit Oracle 8i: `cd ORACLE_HOME/lib64`
64-bit Oracle 9i/10g: `cd ORACLE_HOME/lib`
2. Perform this step only if the `libobk.sl` (HP-UX) or `libobk.so` (Solaris and other UNIX) file is already created in the `ORACLE_HOME/lib` directory. Otherwise, skip this step.

Run:

HP-UX: `mv libobk.sl libobk.sl.orig`

Solaris and other UNIX: `mv libobk.so libobk.so.orig`



IMPORTANT:

If you intend to uninstall the Data Protector Oracle integration and to continue using Oracle on the same system after the integration is removed, do not delete `libobk.sl.orig` (HP-UX) or `libobk.so.orig` (Solaris and other UNIX).

3. Run:

HP-UX:

- 32-bit:
`ln -s /opt/omni/lib/libob2oracle8.sl libobk.sl`
- 64-bit:
`ln -s /opt/omni/lib/libob2oracle8_64bit.sl libobk.sl`

Solaris:

- 32-bit:
`ln -s /optS/omni/lib/libob2oracle8.so libobk.so`
- 64-bit:
`ln -s /opt/omni/lib/libob2oracle8_64bit.so libobk.so`

Other UNIX:

- 32-bit:
`ln -s /opt/omni/lib/libob2oracle8.so libobk.so`
- 64-bit:
`ln -s /opt/omni/lib/libob2oracle8_64bit.so libobk.so`

Linking on Windows systems

On Oracle Server systems, copy the file

`Data_Protector_home\bin\orasbt.dll` to the directory `Oracle_home\bin`.

Linking on OpenVMS systems

On Oracle Server systems running on OpenVMS, link the MML

`SYS$SHARE:LIBOBK2SHR32_8I.EXE` (Oracle8i) or

`SYS$SHARE:LIBOBK2SHR32.EXE` (Oracle9i) with the Oracle Server.

Linking Oracle8i

1. Run `ORAUSER.COM` under `$ORACLE_HOME/UTIL`.

2. Edit the following files:

- `ORA_UTIL:RDBMS_RMAN_NOSHARE.OPT`

Example

```
!rdbmsm libraries
ora_olb:libvsn8/lib
!ora_rman_mml/lib COMMENT OUT THIS LINE
ora_olb:libwtc8/lib
ora_olb:libclient8/lib
ora_olb:libcommon8/lib
ora_olb:libgeneric8/lib
ora_olb:libclient8/lib
ora_olb:libcommon8/lib
generic8/libgeneric8/lib
```

- `ORA_RDBMS:LORACLE_64.COM`

Example

```
ora_olb:libclient8_64/lib/incl=(kgu) ,-
`rdbmslib$$`-
`plsqllib$$`-
`rdbmslib$$`-
!ora_rman_mml_64/lib,- COMMENT OUT THIS LINE
ora_olb:libnro8_64/lib,-
`network$$`-
ora_olb:libtrace8_64/lib,-
`oracoreSS`-
`cart64$$`-
ora_olb:libslax8_64/lib,-
`utl$$`-
`oracore$$`-
sys$input/options
SYS$SHARE:LIBOBK2SHR32_8I.EXE/SHARE,- ADD THIS LINE
```

- `ORA_UTIL:LOUTL.COM`

Example

```
$nonSharedLink:
`loutl_link_cmd$$`/alpha/nouserlibrary`dotrace$$`map$$`
mapextra$$`
image$$`=
`filename$$`switch$$`userlink$$`/sysexe -
`p2`,-
ora_olb:libclient8/lib,-
```

```

ora_olb:libsql8/lib,-
'ocis$$'-
'fastupi$$'-
'network$$'-
rdbmslib_noshare$$'-
'oracore$$'-
'network$$'-
'rdbmslib_noshare$$'-
'otracelib$$'-
'oracore$$'-
'rdbmslib_noshare$$'-
'oracore$$'-
'useroption$$'-
sys$input/opt
SYS$SHARE:LIBOBK2SHR32_8I.EXE/SHARE,- ADD THIS LINE
sys$share:decc$shr/share
!Temporary: fixup readonly attributes between compiler versions
psect_attr = $readonly$,pic,shr

```

3. Shut down the Oracle database instance on the Oracle Server system.

4. Re-link ORA_RDBMS : executables by invoking ORA_INSTALL:ORACLEINS:

```
$@ORA_INSTALL:ORACLEINS
Oracle Installation Startup Menu
Options:
1. Create a new ORACLE system.
2. Upgrade your system from the Oracle distribution tape.
3. Reconfigure existing products, manage the database, or load demo tables.
4. Exit.
```

Choose option 3.

NOTE:

Before upgrading, configuring, or managing the database, or loading demo tables, run `ORA_UTIL:ORAUSER.COM`. If you created an instance, run:

```
ORA_DB:ORAUSER_DB_NAME.COM SID setup_node.
```

When you are prompted for the root directory, enter `DISK$ORADISK_ODS5:[ORACLE8.HOME1]`.

NOTE:

If loading products from savesets, enter the drive or directory where savesets are located. If loading from a remote device, do not include username and password. For more information, see the Oracle documentation for OpenVMS.

When you are back at the main menu, select the option `Software Installation and Upgrade Menu`. The following appears:

```
Software Installation and Upgrade Menu
1. Select Licensed Products to Load
2. Select Build Configuration Options
3. Load and Build Selected Licensed Products
4. Build Selected Licensed Products
```

Enter 1.

Select the licensed products from the list by entering the number assigned to RDBMS.

Exit the menu. You are taken to Software Installation and Upgrade Menu.

Enter 2 to select build configuration options. You are now at the Select Configuration Options menu.

Enter the number assigned to RDBMS. Select RDBMS configuration options as follows:

1. System or Group Installation? [S/G] S
2. ORACLE Image Identifier? [@6] V817
3. Include Distributed database option? [Y/N] Y
4. Include Context option? [Y/N] Y
5. Include Object Support option? [Y/N] Y
6. Include Spatial Data option? [Y/N] Y
7. Include Data Partitioning option? [Y/N] Y
8. Include Parallel Server option? [Y/N] Y
9. Include Java Aurora external option? [Y/N] N

The options marked by Y will be selected.

Exit the menu to return to Select Configuration Options. Enter the number of the product you want to configure (18 corresponds to RDBMS). In Software Installation and Upgrade Menu, enter 4 to build the selected licensed products (RDBMS). That will initiate the relinking process.



NOTE:

To create known file entries for the linked products using the `VMS INSTALL` utility, run `ORA_INSTALL:ORA_INSUTL.COM`. For details, see the Oracle documentation for OpenVMS.

After relinking

1. Start the Oracle database.

2. Configure ORACLE8I using the GUI (see “Configuring Oracle databases” on page 47), and then execute the following RMAN script to test the MML (SBT) interface:

```
run {  
  allocate channel 'dummy' type 'SBT_tape';  
  release channel 'dummy';  
}
```

If the channel allocation through SBT succeeds, relinking was performed successfully.

Linking Oracle9i

1. Make sure Oracle RMAN is set up and you are able to access it. This can be achieved by performing a test backup using the following RMAN script:

```
{  
  allocate channel d1 type disk;  
  backup tablespace system;  
  release channel d1;  
}
```

You can skip this step if you are already using RMAN for backing up Oracle.

2. Check the presence of the MML `LIBOBK2SHR32.EXE` in the `SYS$SHARE :` directory.

 **NOTE:**

The logical definition for `SYS$SHARE : LIBOBK2SHR32 . EXE` is
`$DEFINE / SYSTEM DP_SBT SYS$SHARE : LIBOBK2SHR32 . EXE.`

You are now ready to use the MML with RMAN to perform backups. For information on how to use RMAN, see the Oracle documentation.

After relinking

To test the MML (SBT) interface, configure Oracle 9i using the GUI (see “Configuring Oracle databases” on page 47).

Configuring Oracle users on UNIX and OpenVMS

On UNIX and OpenVMS, to start an Oracle backup session, a user needs to perform an operating system logon to the system where an Oracle Server is running.

If properly configured, this user is allowed to back up or restore an Oracle database. To start a backup of an Oracle database using Data Protector, the user must also become the owner of the Data Protector backup specification.

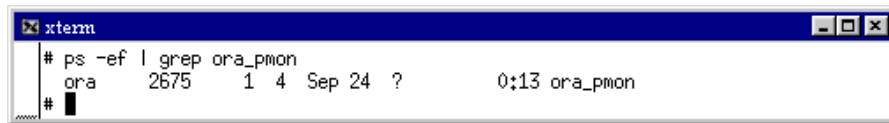
As the owner of the backup specification, the Oracle user must be added to the Data Protector `admin` or `operator` user group. On OpenVMS, configure a Data Protector `admin` user with the username `<Any>` and the group name `<Any>`.

On UNIX, you can identify this user by running the following command on the Oracle Server system:

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

or

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



```
xterm
# ps -ef | grep ora_pmon
ora      2675      1  4  Sep 24  ?        0:13 ora_pmon
#
```

Figure 3 Finding the Oracle user

The example above states that the user `ora` has sufficient privileges within the Oracle database to back up and restore the database. Therefore, this user must be added to the corresponding Data Protector user group (`admin` or `operator`) and must also become the owner of the backup specification to be able to back up the Oracle database using Data Protector.

 **IMPORTANT:**

Additionally, the user `root` on the Oracle Server has to be added to the Data Protector `admin` or `operator` user group.

For information on how to add a user to a user group, see the online Help index: “adding users”.

After the two users are added to the Data Protector `admin` or `operator` user group, Data Protector sessions can be started under the user account with all the necessary privileges required to perform an Oracle database backup with Data Protector.

MC/ServiceGuard: In a cluster environment, add both users (the Oracle user and the user `root`) to the Data Protector `admin` or `operator` group on the virtual server and on every physical and virtual node in the cluster.

If two or more Oracle users have the same user ID, all of them must be added to the Data Protector `admin` or `operator` user group.

OpenVMS

To configure an Oracle user on OpenVMS, proceed as follows:

1. Oracle 9i

Modify the location of `ORAUSER.COM` and `ORATAB` files.

- `ORAUSER.COM`

Depending on the current location of `ORAUSER.COM`, modify `$PIPE@DKA0:[ORACLE]ORAUSER.COM > NLA0:` accordingly. For example, if `ORAUSER.COM` is located in `DKC0:[ORACLE9i]`, the changes will be:

```
$PIPE@DKC0:[ORACLE9i]ORAUSER.COM > NLA0:
```

- `ORATAB`

Depending on the current location of `ORATAB`, modify `$DEFINE/NOLOG/JOB ORATAB_LOC DKA0:[ORACLE]ORATAB` accordingly. For example, if `ORATAB` is located in `DKC0:[ORACLE9i]`, the changes will be:

```
$DEFINE/NOLOG/JOB ORATAB_LOC DKCF0:[ORACLE9i]ORATAB
```

Oracle 8i

Execute the `OMNI$ROOT:[BIN]DP_ORA8I_RENAME.COM` command. This will update the required Oracle8i executables.

2. Oracle 8i/9i

Uncomment the following lines in `OMNI$ROOT:[LOG]LOGIN.COM`:

```
$DEFINE /NOLOG /SYSTEM DP_SBT SYS$SHARE:LIBOBK2SHR32.EXE
$@OMNI$ROOT:[BIN]OMNI$CLI_SETUP.COM
$@OMNI$ROOT:[BIN.PERL]PERL_SETUP.COM$DEFINE /process
PERL_ENV_TABLES "LNM$PROCESS", "LNM$JOB", "LNM$SERVER",
"LNM$GROUP", "LNM$SYSTEM"
```

3. Oracle 8i/9i

If you run the Media Agent and Data Protector Oracle integration agents on the same OpenVMS system, modify the group ID of the `omniadmin` user as DBA using the `MCR AUTHORIZE` utility:

a. Log in as a privileged user.

b. Execute:

```
$set def sys$system $mcr authorizeUAF>show  
omniadminUAF>show oracle
```

c. Compare the accounts for Oracle and `omniadmin` users. If the accounts are different, execute:

```
UAF>modify omniadmin/UIC=UID show
```

d. Verify the changes of the group ID.

4. Oracle 8i/9i

If you use CLI commands for Oracle integration agents, execute `OMNI$ROOT:[LOG]LOGIN.COM`.

5. Oracle 8i/9i

Verify that the `-key Oracle8` entry is present in `OMNI$ROOT:[CONFIG.CLIENT]OMNI_INFO`, for example:

```
-key oracle8 -desc "Oracle Integration" -nlset 159 -nlsetId  
12172 -flags 0x7 -ntpath "" -uxpath "" -version A.06.00
```

If the entry is not present, copy it from

`OMNI$ROOT:[CONFIG.CLIENT]OMNI_FORMAT`. Otherwise, the Oracle integration will not be shown as installed on the OpenVMS client.

TIP:

To determine the status of processes (`OMNI$I*`) and subprocesses (`OMNI$ADMIN_*`) on your OpenVMS system, use the following command procedure:

```
$@OMNI$ROOT:[BIN]OMNI$DIAGNOSE.COM
```

This command procedure displays the active parent processes, the session of job name, and the logfile name.

Configuring Oracle databases

Configuring an Oracle database involves preparing the environment for starting a backup. The environment parameters such as the Oracle home directory and the connection string to the database are saved in the Data Protector Oracle configuration files on the Cell Manager. The database must be open during the configuration procedure. The configuration must be done for each Oracle database.

If a recovery catalog has been created and the Oracle target database has not yet been registered in the recovery catalog database, this will occur during the configuration procedure.

To configure an Oracle database, use the Data Protector GUI or CLI.

Using the Data Protector GUI

Configure an Oracle database when you create first backup specification for the database. Start with the procedure described in [“Creating backup specifications”](#) on page 58 and at [Step 5](#) on page 62 proceed as follows:

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

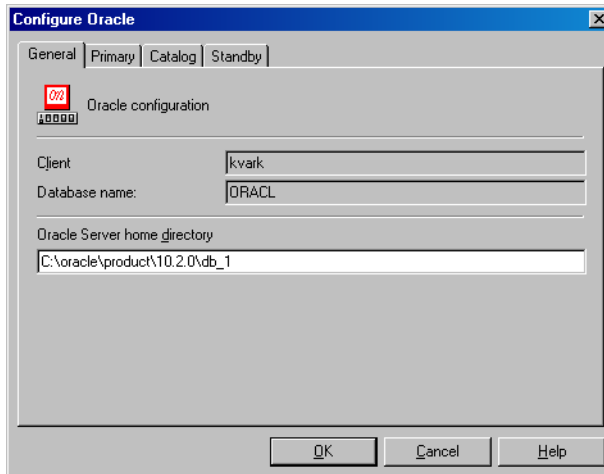


Figure 4 Configuring Oracle - General (Windows)

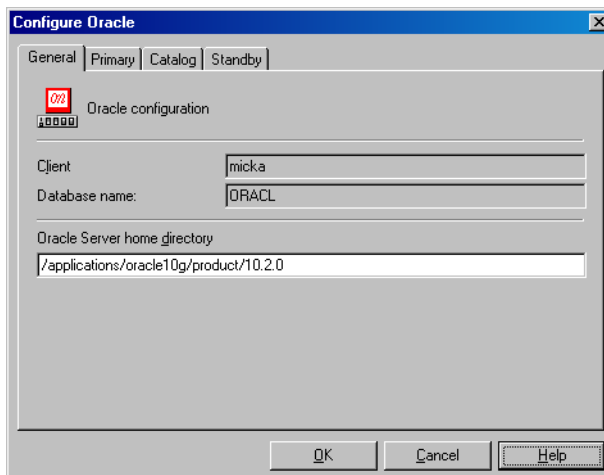


Figure 5 Configuring Oracle - General (UNIX)

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.

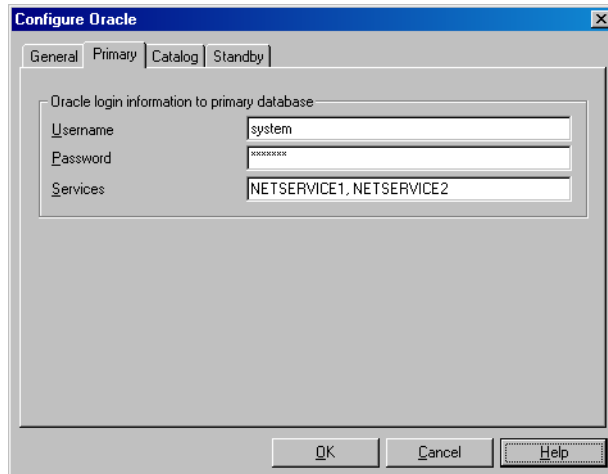


Figure 6 Configuring Oracle - Primary

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.

Oracle Data Guard: If you intend to back up a standby database, 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.

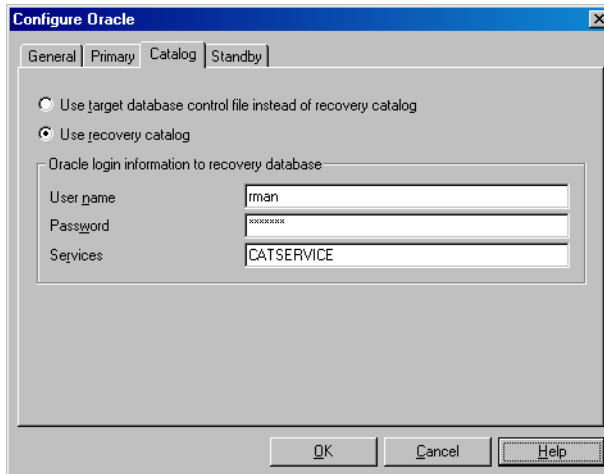


Figure 7 Configuring Oracle - Catalog

- Oracle Data Guard:** 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.

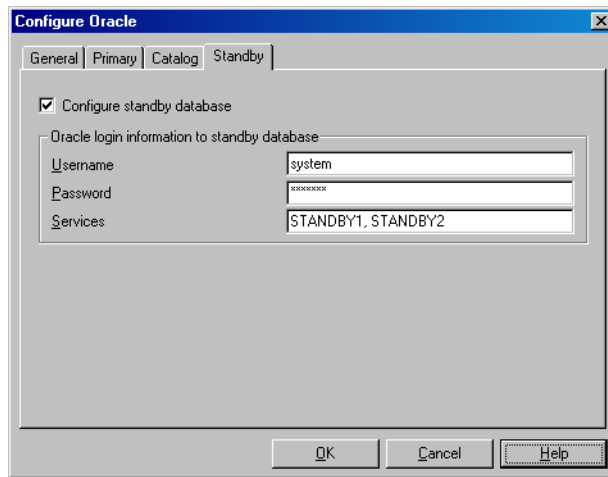


Figure 8 Configuring Oracle - Standby

- Click **OK**.

The Oracle database is configured. Exit the GUI or proceed with creating the backup specification at [Step 6](#) on page 63.

Using the Data Protector CLI

 **NOTE:**

On OpenVMS, to invoke the Data Protector CLI, run:
`$@OMNI$ROOT: [BIN]OMNI$CLI_SETUP.COM`

- UNIX only:** Log in to the Oracle Server system as user `root` or as the Oracle user that is identified as described in [“Configuring Oracle users on UNIX and OpenVMS”](#) on page 44.

2. On the Oracle Server system, from the directory:

Windows: *Data_Protector_home\bin*

HP-UX, Solaris, and Linux: */opt/omni/lbin*

Other UNIX: */usr/omni/bin/*

OpenVMS: *OMNI\$ROOT: [BIN]*

run:

Windows:

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

UNIX and OpenVMS:

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

where:

PRIMARY_DB_LOGIN is:

-prmuser *PRIMARY_USERNAME*

-prmpasswd *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, ...]

Oracle Data Guard: If you intend to back up a standby database, you must provide the *STANDBY_DB_LOGIN* information. For standby database backup, a recovery catalog must be used. Therefore, you must also provide the *CATALOG_DB_LOGIN* information.

Parameter description

CLIENT_NAME

Name of the Oracle Server system with the database to be configured. It needs to be specified only in a cluster environment.

RAC: Name of the node or the virtual server of the Oracle resource group. The latter can only be used on HP-UX.

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.

Example

The following example represents configuration on HP-UX or Solaris of an Oracle database and its recovery catalog in Oracle Data Guard environment.

The following names are used in the example:

- database name: `oracl`
- primary user name: `system`
- primary password: `manager`
- primary net service name 1: `net service1`
- primary net service name 2: `net service2`
- recovery catalog user name: `rman`
- recovery catalog password: `manager`
- recovery catalog net service name: `cat service`
- standby user name: `system`
- standby password: `manager`
- standby net service name 1: `net servicesb1`
- standby net service name 2: `net servicesb2`

Syntax

```
/opt/omni/lbin/util_oracle8.pl -config -dbname oracl -orahome
/app10g/oracle10g/product/10.1.0 -prmsuser system -prmpasswd
manager -prmservice net service1,net service2 -rcuser rman
-rcpasswd manager -rcservice cat service -stbuser system
-stbpasswd manager -stbservice net servicesb1,net servicesb2
```

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.

What happens after the configuration?

The `util_oracle8.pl` command is started on the Oracle server system. It saves the configuration parameters in the Data Protector Oracle configuration files.

If the recovery catalog was selected, `util_oracle8.pl` starts the Oracle RMAN command, which registers the target database in the recovery catalog.

Information about the Oracle database's structure is transferred to the recovery catalog from the Oracle database's control files.

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:

On UNIX, it is possible that although the GUI check is successful, the backup still fails. This can happen if the backup owner is not the Oracle user `root` or the Oracle user that is identified as described in [“Configuring Oracle users on UNIX and OpenVMS”](#) on page 44.

Using the Data Protector CLI

1. **UNIX only:** Log in to the Oracle server system as the Oracle user or as user `root`.

2. From the directory:

Windows: *Data_Protector_home\bin*

HP-UX, Solaris, and Linux: */opt/omni/lbin*

Other UNIX: */usr/omni/bin/*

OpenVMS: *OMNI\$ROOT: [BIN]*

run:

Windows:

```
perl -I..\lib\perl util_oracle8.pl -CHKCONF -dbname DB_NAME
```

UNIX and OpenVMS:

```
util_oracle8.pl -CHKCONF -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: *Data_Protector_home\bin\omnigetmsg 12 error_number*

HP-UX, Solaris, and Linux: */opt/omni/lbin/omnigetmsg 12 error_number*

Other UNIX systems: */usr/omni/bin/omnigetmsg 12 error_number*

OpenVMS: *\$_OMNI\$ROOT: [BIN]OMNI\$CLI_SETUP.COM \$_OMNIGETMSG 12 error_number*

IMPORTANT:

On UNIX, it is possible that although you receive a **RETVAL*0*, the backup still fails.

This can happen if the backup owner is not the Oracle user *root* or the Oracle user that is identified as described in

[“Configuring Oracle users on UNIX and OpenVMS”](#) on page 44.

Backup

To configure an Oracle backup, perform the following steps:

1. Configure the devices you plan to use for a backup. See the online Help index: “configuring devices” for instructions.
2. Configure media pools and media for a backup. See the online Help index: “creating media pools” for instructions.
3. Create a Data Protector Oracle backup specification. See “[Creating backup specifications](#)” on page 58.

OpenVMS

On OpenVMS, before performing Data Protector tasks using the CLI, execute:

```
$@OMNI$ROOT: [BIN]OMNI$CLI_SETUP.COM
```

This command procedure defines the symbols needed to invoke the Data Protector CLI. It gets installed when you chose the CLI option during the installation. Execute this command procedure from `LOGIN.COM` for all CLI users.

Creating new templates

You can use backup templates to apply same set of options to a number of backup specifications. By creating your own template, you can specify the options exactly as you want them to be.

This allows you to apply all the options to a backup specification with a few mouse clicks, rather than having to specify all the options over and over again. This task is optional, as you can use one of the default templates as well.

If you prefer using predefined templates, see “[Creating backup specifications](#)” on page 58 for a detailed explanation.

To create a new backup template, proceed as follows:

1. In the `Data Protector Manager`, switch to the **Backup** context.
2. In the Scoping Pane, expand **Backup** and then **Templates**, and then right-click **Oracle Server**.
3. Click **Add Template**. Follow the wizard to define the appropriate backup options in your template.

Creating backup specifications

Cluster-aware clients

Before you perform an *offline* backup in a cluster environment, take the Oracle Database resource offline and bring it back online after the backup. 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.

To create an Oracle 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, double-click **Blank Oracle Backup** to create a backup specification without predefined options, or use one of the pre-defined templates given below

Archive	Backs up the Archived Redo Logs.
Archive_Delete	Backs up the Archived Redo Logs, then deletes them after the backup.
Whole_Online	Backs up the database instance and the Archived Redo Logs.
Whole_Online_Delete	Backs up the database instance and the Archived Redo Logs, and then deletes the Archived Redo Logs.
Database_Archive	Backs up the database instance and the Archived Redo Logs.
Database_Switch_Archive	Backs up the database instance, switches the Online Redo Logs and backs up the Archived Redo Logs.
Database_Switch_ArchiveDel	Backs up the database instance, switches the Online Redo Logs, backs up the Archived Redo Logs and then deletes the Archived Redo Logs.
Direct_Database	Backs up the database instance and controlfile.
SMB_Proxy_Database	Backs up the database instance and control file in the ZDB (split mirror or snapshot) mode using the proxy-copy method.
SMB_BackupSet_Database	Backs up the database instance and control file in the ZDB (split mirror or snapshot) mode using the backup set method.

Click **OK**.

4. In the **Client**, select the Data Protector Oracle integration client. In a cluster environment, select the virtual server.

RAC: Select either the node or the virtual server of the Oracle resource group. The latter can only be selected on HP-UX.

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

Specify the application that you want to back up. 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;
```

RAC: Note that the database name is the same for all instances.

UNIX only: Type the username and user group of the Oracle user. See [“Configuring Oracle users on UNIX and OpenVMS”](#) on page 44 for information on how to identify that user.

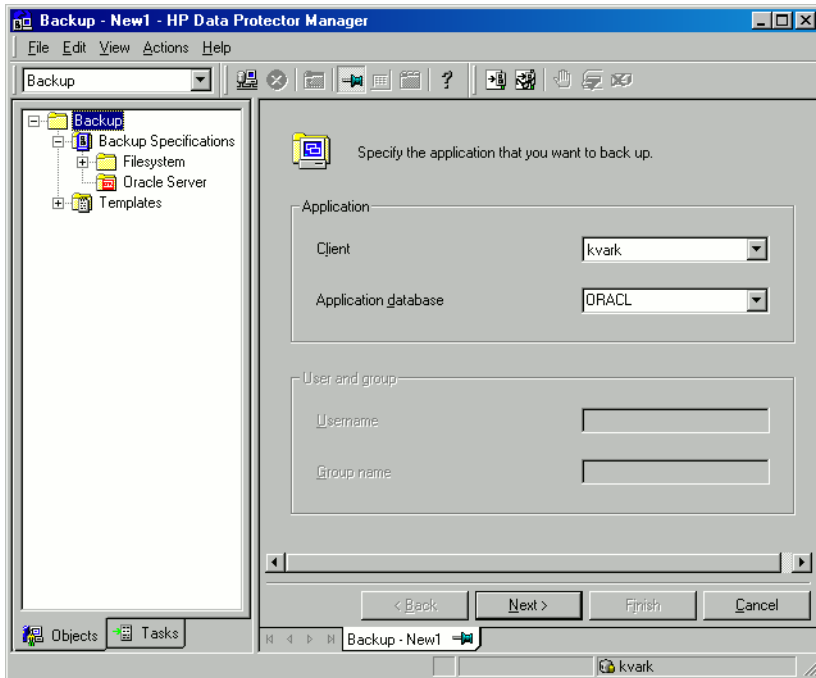


Figure 9 Specifying an Oracle Server system (UNIX)

Click **Next**.

5. 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”](#) on page 47.

6. Select the Oracle database objects to be backed up.

For example, a single tablespace can be separately selected for backup, but for a complete online backup of the database, the **ARCHIVELOGS** must also be selected.

Oracle 10g: The archived logs can reside in the flash recovery area. In this case, if you select the **FLASH RECOVERY AREA** to be backed up, you do not need to select also **ARCHIVELOGS**.

Oracle Data Guard (10g): If the database is configured with standby connection, you can back up a control file for the standby database, which can be used when restoring the standby database.

 **NOTE:**

If your database uses a recovery catalog, it is backed up by default after each database backup, unless otherwise specified in the backup specification.

Click **Next**.

7. 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".

Click **Next** to proceed.

8. Set the backup options.

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

Oracle Data Guard: To back up a standby database, you must select **Back up standby database** in the Application Specific Options dialog box.

For information on the Application Specific Options (Figure 10 on page 64), see Table 4 on page 66 or press **F1**.

 **TIP:**

When backing up data from the Oracle 10g flash recovery area to tape, you can specify the location of the RMAN script that performs backups to the flash recovery area in the **Pre-exec** or **Post-exec** text box. The script will be executed every time before (**Pre-exec**) or after (**Post-exec**) the Data Protector Oracle integration backup to tape.

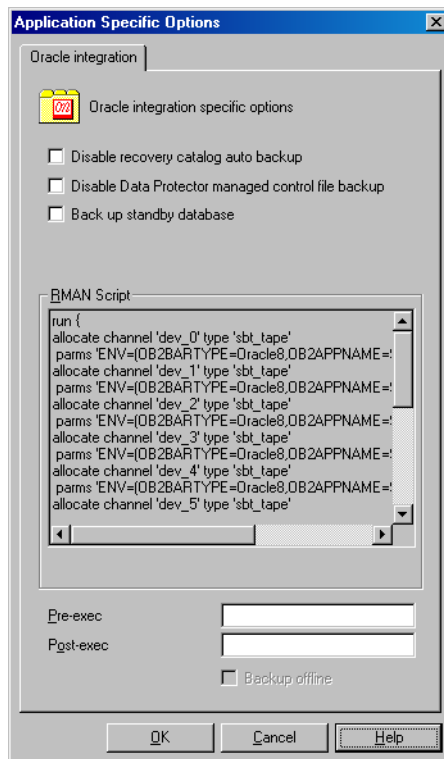


Figure 10 Oracle specific options

Click **Next**.

9. Optionally, schedule the backup. For more details, see [“Scheduling backup specifications”](#) on page 78.

Click **Next**.

10. 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.

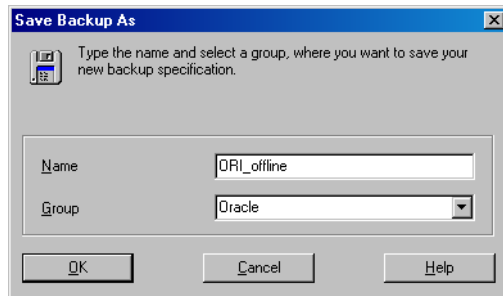


Figure 11 Saving the backup specification

Click **OK**.

To start the backup, see [“Starting backup sessions”](#) on page 75.

11. On UNIX, after the backup specification is saved, verify that the owner of the backup specification is the specified Oracle user. See [“Configuring Oracle users on UNIX and OpenVMS”](#) on page 44 for details about this user.

- 12.** You can examine the newly-created and saved backup specification in the **Backup** context, under the specified group of backup specifications. The backup specification is stored in the following file on the Cell Manager:

Windows:

Data_Protector_home\Config\server\Barlists\Oracle8\Backup_Specification_Name

UNIX: /etc/opt/omni/server/barlists/oracle8/Backup_Spec_Name

- 13.** It is recommended to test the backup specification. See “[Testing the integration](#)” on page 72 for details.

Table 4 Oracle backup options

Disable recovery catalog auto backup	By default, Data Protector backs up the recovery catalog in every backup session. 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 in every backup session. Select this option to disable backup of the Data Protector managed control file.
Back up standby database	Oracle Data Guard: This option is applicable if the database is configured with the standby connection. By default, RMAN backs up the database files and archived redo logs on the primary system. Select this option to enable backup of the database files and archive logs on standby system. However, only the archive logs created after the standby database was configured can be backed up at standby site. Archive logs created before the standby database was configured must be backed up on the primary database. Note that the current control file or the control file for standby will still be backed up from the primary system.
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 ” on page 68.

Pre-exec, Post-exec

Specify a command or RMAN script that will be started by `ob2rman.pl` on the Oracle Server system before the backup (`pre-exec`) or after it (`post-exec`). RMAN scripts must have the `.rman` 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 UNIX, see [“Examples of pre-exec and post-exec scripts on UNIX”](#) on page 67.

Provide the pathname of the command or RMAN script.

OpenVMS: Provide the pathname of the command (`OMNI$ROOT:[BIN]`).

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

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 (introduced by Oracle 9i).
- 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)
 - With Oracle 10g, the flash recovery area (the third 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 [Figure 17](#) on page 96), 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.
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)';
```

- 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, datafile, or the flash recovery area**. 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;
```

 **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 %s:%t:%p substitution variables and *DB_NAME*, which are obligatory.

- 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**.

If an appropriate template was selected, or if the statement was manually added, the RMAN sql statement to switch the Online Redo Logs before backing up the Archived Redo Logs:

```
sql 'alter system archive log current';
```

The backup statement consists of the following:

- The Oracle format of the backup file in the following format:

```
format 'Backup_Specification_NameDB_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*.

- The RMAN archivelog *all* command.
- If an appropriate template was selected, or if the statement was manually added, the RMAN statement to delete the Archived Redo Logs after they are backed up:

```
archivelog all delete input;
```

- 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;
```

 **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 %s:%t:%p substitution variables and *DB_NAME*, which are obligatory.

- The RMAN `current controlfile` command.

Example of the RMAN script

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

```
run {
allocate channel 'dev_0' type 'sbt_tape' parms
'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DIPSI,OB2BARLIST=New1)';
allocate channel 'dev_1' type 'sbt_tape' parms
'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DIPSI,OB2BARLIST=New1)';
allocate channel 'dev_2' type 'sbt_tape' parms
'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DIPSI,OB2BARLIST=New1)';
backup incremental level <incr_level>
format 'New1<DIPSI_%s:%t:%p>.dbf'
database
;
backup format 'New1<DIPSI_%s:%t:%p>.dbf' archivelog all;
backup format 'New1<DIPSI_%s:%t:%p>.dbf' current controlfile
;
}
```

Creating copies of backed up objects

Oracle duplex mode

Oracle supports the duplex mode, which allows you to create copies of every backed up object to a separate backup device. To enable the duplex feature, perform the following steps:

1. Add the following command to the RMAN script before any allocate channel command:

```
set duplex=<on | 2 | ... >
```

 **IMPORTANT:**

If more than one allocated channel is used, it may happen that some original and copied objects are backed up to the same medium. To prevent this, you should use only one allocated channel when backing up using the duplex mode.

2. Add the following parameter to every format string used for backup:

```
%c
```

3. Set the concurrency of each device used for backup to 1.
4. Set the `MIN` and `MAX` load balancing parameters according to the following formula:

*(number of duplex copies) * (number of allocated channels)*

Example

If the duplex is set to 2 and the backup runs with 1 allocated channel, then the `MIN` and `MAX` parameters should be set to 2.

 **IMPORTANT:**

If the `MIN` and `MAX` load balancing parameters are set to lower values, the backup will hang.

If the `MIN` and `MAX` load balancing parameters are set to higher values, it may happen that the original and copied objects are backed up to the same medium.

Testing the integration

Once you have created and saved a backup specification, you should test it before running a backup. The test verifies both parts of the integration, the Oracle side and the Data Protector side. In addition, the configuration is tested as well.

The procedure consists of checking both the Oracle and the Data Protector parts of the integration to ensure that communication between Oracle and Data Protector is established, that the data transfer works properly, and that the transactions are recorded either in the recovery catalog (if used) or in the control file.

Details of the test backup, such as media protection, backup user and backup status are registered in the Data Protector database and in the Oracle control files. Set the **Protection** option of your test backup specification to **None**.

Testing using the Data Protector GUI

Follow the procedure below to test the backup of an Oracle backup specification:

1. In the **Data Protector Manager**, switch to the **Backup** context.
2. In the **Scoping Pane**, expand **Backup**, then **Backup Specifications**. Expand **Oracle Server** and right-click the backup specification you want to preview.
3. Click **Preview Backup**.

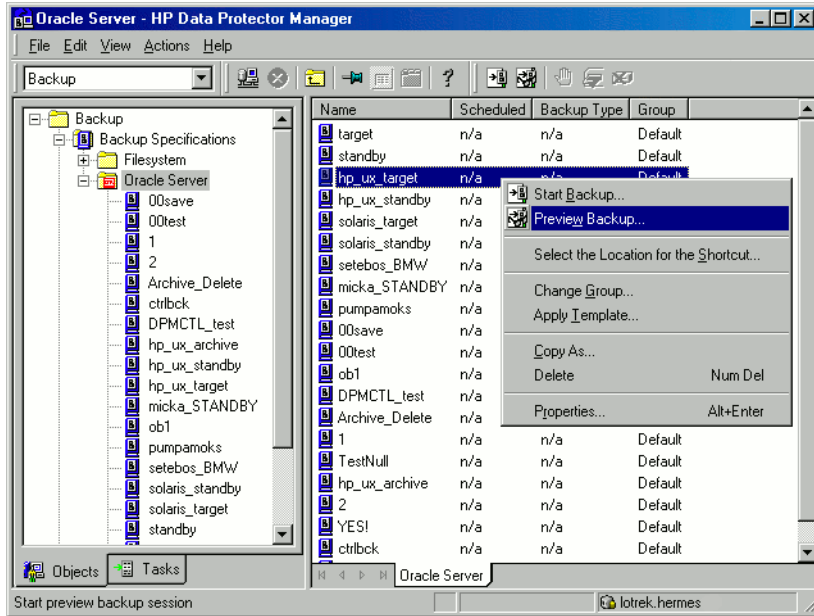


Figure 12 Previewing a backup

Testing using the CLI

A test can be executed from the command line on the Oracle Server system or on any Data Protector client system within the same Data Protector cell, provided that the system has the Data Protector User Interface installed.

NOTE:

On OpenVMS, to invoke the Data Protector CLI, execute:
`.$@OMNI$ROOT:[BIN]OMNI$CLI_SETUP.COM`

Run the omnib command with the `-test_bar` option as follows:

- On Windows: `Data_Protector_home\bin\omnib -oracle8_list backup_specification_name -test_bar`
- On HP-UX, Solaris, and Linux: `/opt/omni/bin/omnib -oracle8_list \ backup_specification_name -test_bar`
- On other UNIX systems: `/usr/omni/bin/omnib -oracle8_list \ backup_specification_name -test_bar`

- On OpenVMS: `$omnib -oracle8_1 qist backup_specification_name -test_bar`

The `ob2rman.pl` command is started, which then starts the `BACKUP VALIDATE DATABASE RMAN` command.

Starting backup sessions

There are two strategies for backing up a database. These are an **offline** or **consistent** database backup, and an **online** or inconsistent database backup. The latter is also known as a **hot** backup. Special attention is required to reach a consistent state with an online backup.

A decision about your database backup strategy depends on a number of factors. If the database must be open and available all the time, then online backup is your only choice. If you can afford to have the database offline at a certain time, then you are more likely to make periodic offline backups of the entire database, supplementing them with online backups of the dynamically changing tablespaces.

Oracle offline

An offline backup of a database is a backup of the datafiles and control files which are consistent at a certain point in time. The only way to achieve this consistency is to cleanly shut down the database and then back up the files while the database is either closed or mounted.

If the database is closed, the offline backup of an Oracle target database can be performed using a Data Protector filesystem backup specification. In this case, the Data Protector Disk Agent is used.

If the database is mounted, a Data Protector Oracle backup specification, based on which Data Protector automatically generates and executes the RMAN script, can be used. In this case, the Data Protector Oracle integration software component is used.

Typically, you would perform an offline backup of the entire database, which must include all datafiles and control files, while the parameter files may be included optionally.

The whole offline database backup is performed as follows:

1. Shut down the database cleanly.

A clean shutdown means that the database is not shut down using the `ABORT` option.

2. Mount the database if you are backing it up using RMAN.

3. Back up all datafiles, control files and, optionally, parameter files.
4. Restart the database in the normal online mode.

Oracle online

As opposed to an offline backup, an online backup is performed when a database is open.

The backup of an open database is inconsistent, because portions of the database are being modified and written to disk while the backup is progressing. Such changes to the database are entered into the online redo logs as well. A database running in the `ARCHIVELOG` mode enables the archiving of the online redo logs. In the case of a restore, this feature is essential to bring a database to a consistent state as part of the entire restore process.

When using an online backup, the following must be done in order to bring the database to a consistent state:

1. Restore the database files (which are inconsistent) to disk.
2. Perform database recovery, which requires applying the Archived Redo Logs. This is an Oracle operation.

An Oracle online database backup can be performed using the Oracle RMAN utility or Data Protector GUI. In the latter case, Data Protector creates and executes the RMAN script automatically based on data entered in the Data Protector GUI. During an Oracle online backup, the Oracle target database is open, while tablespaces, datafiles, control files, and archived redo logs are being backed up.

The database must operate in the `ARCHIVELOG` mode so that the current Online Redo Logs are archived to the Archived Redo Logs.



IMPORTANT:

Before you run an Oracle online backup, make sure that the database is really operating in `ARCHIVELOG` mode. This can be done on the Oracle server system by starting SQL*Plus and issuing the following command:

```
archive log list;
```

If the Oracle target database is not operating in the `ARCHIVELOG` mode, proceed as follows:

If SPFILE is used:

1. Shut down the database.
2. Mount the database.
3. Start SQL*Plus and type:

```
alter database archivelog;  
alter database open;  
alter system archive log start SCOPE=SPFILE;
```

If PFILE is used:

1. Shut down the database.
2. Change PFILE to enable log archiving by setting:

```
log_archive_start = true
```

3. Mount the database.
4. Start SQL*Plus and type:

```
alter database archivelog;  
alter database open;
```

Oracle Data Guard: The archive logs generated after an archive log backup must be manually cataloged so that they are known to RMAN for future backups when:

- The primary or standby control file is re-created. The archive logs must be re-cataloged because RMAN uses the control file to determine which archive logs must be backed up.
- The primary database role changes to standby after a failover. The archive logs must be re-cataloged because a change in database role resets the version time of the mounted control file.

Use the RMAN command `CATALOG ARCHIVELOG 'archive_log_file_name';` to manually catalog the archived redo logs.

Now you are ready to run an online backup of the Oracle database, using any of the following methods:

Backup methods

- Schedule a backup of an existing Oracle backup specification using the Data Protector Scheduler. See [“Scheduling backup specifications”](#) on page 78.

- Start an interactive backup of an existing Oracle backup specification using the Data Protector GUI or the Data Protector CLI. See [“Running an interactive backup”](#) on page 80.
- Start a backup on the Oracle server using either Oracle Recovery Manager or Oracle Enterprise Manager. See [“Starting Oracle backup using RMAN”](#) on page 82.

Backup procedure

The following happens when you start a backup using the Data Protector user interface:

- 1.** Data Protector executes `ob2rman.pl` on the client. This command starts RMAN and sends the Oracle RMAN Backup Command Script to the standard input of the RMAN command.
- 2.** The Oracle RMAN contacts the Oracle Server, which contacts Data Protector via the MML interface and initiates a backup.
- 3.** During the backup session, the Oracle Server reads data from the disk and sends it to Data Protector for writing to the backup device.

Messages from the Data Protector backup session and messages generated by Oracle are logged to the Data Protector database.

A backup of the Oracle recovery catalog is performed automatically following each Oracle target database backup, unless otherwise specified in the backup specification. Using the standard Oracle export utility, the Data Protector `ob2rman.pl` starts an export of the Oracle recovery catalog to a file which is then backed up by Data Protector.

Deleting data from the recovery catalog

When backing up an Oracle database using the recovery catalog database, all information about the backup, restore, and database recovery is stored in the recovery catalog. This information is used by RMAN during the restore. If you overwrite or format the media on which this data is backed up, Data Protector exports the object from the Data Protector database. You must manually delete the data from the recovery catalog while logged on to RMAN. See the *Oracle Recovery Manager User's Guide and References* for detailed information about deleting data from the recovery catalog.

Scheduling backup specifications

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

A backup schedule can be tailored according to your business needs. If you have to keep the database online continuously, then you should back it up frequently, including the backup of the Archived Redo Logs, which is required in case you need database recovery to a particular point in time.

For example, you may decide to perform daily backups and make multiple copies of the online redo logs and the Archived Redo Logs to several different locations.

An example of scheduling backups of production databases:

- Weekly full backup
- Daily incremental backup
- Archived Log backups as needed

To schedule an Oracle 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 the backup type can be full or incremental, with the incremental level as high as Incr 4. See [Figure 13](#) on page 80. See the RMAN documentation for details on incremental backup levels.

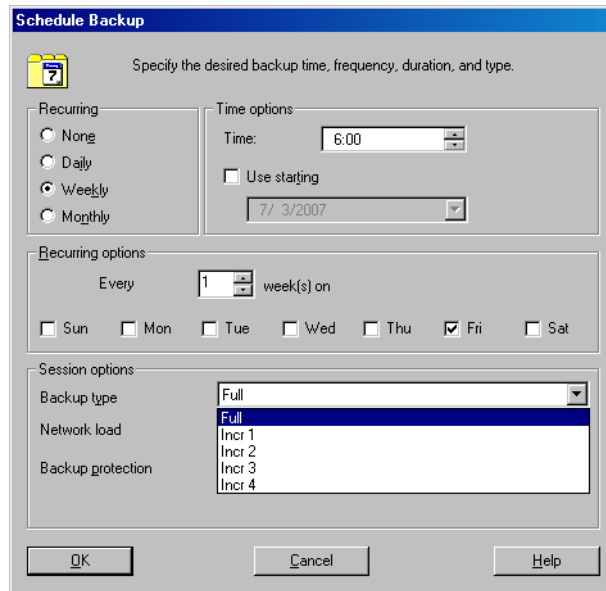


Figure 13 Scheduling backups

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 backup 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 start and click **Start Backup**.

3. In the **Start Backup** dialog box, select the **Backup type** and **Network load** options. For information on these options, click **Help**.

Note that the backup type can be full or incremental, with the incremental level as high as Incr 4. See [Figure 13](#) on page 80. See the RMAN documentation for details on incremental backup levels.

Click **OK**.

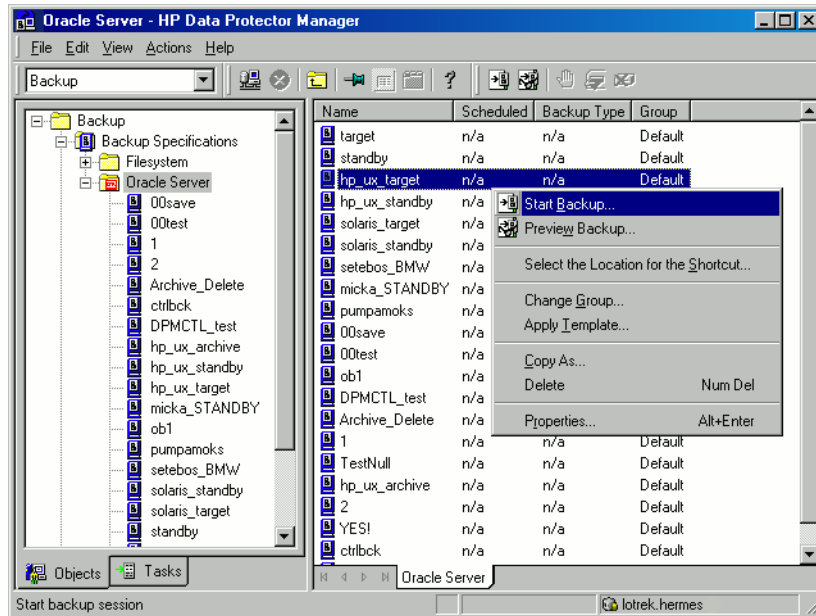


Figure 14 Starting an interactive backup

Starting a backup using the CLI

1. On an Oracle Server, switch to the directory:

Windows: `Data_Protector_home\bin`

HP-UX, Solaris, and Linux: `/opt/omni/bin`

Other UNIX: `/usr/omni/bin`

OpenVMS: To set up the CLI, run:

```
$@OMNI$ROOT: [BIN] OMNI$CLI_SETUP.COM
```

2. Run:

```
omnib -oracle8_list backup_specification_name [-barmode  
Oracle8Mode] [list_options]
```

You can select among the following *list_options*:

```
-protect {none | weeks n | days n | until date | permanent}
```

```
-load {low | medium | high}
```

```
-crc
```

```
-no_monitor
```

```
Oracle8Mode = {-full | -incr1 | -incr2 | -incr3 | -incr4}
```

See the omnib man page for details.

Example

To start a backup using an Oracle backup specification called `RONA`, run the following command:

```
omnib -oracle8_list RONA
```

Starting Oracle backup using RMAN

To start an Oracle backup using RMAN, an Oracle backup specification must be created.

See “[Backup](#)” on page 56 for information on how to create an Oracle backup specification.

To start an Oracle backup using RMAN:

1. Connect to the Oracle target database specified in the backup specification:

If you use the recovery catalog, run:

Oracle 9i/10g:

- On Windows: `ORACLE_HOME\bin\rman target Target_Database_Login catalog Recovery_Catalog_Login`
- On UNIX: `ORACLE_HOME/bin/rman target Target_Database_Login catalog Recovery_Catalog_Login`
- On OpenVMS:
 - a. Run ORAUSER.COM using `$(OMNI$ROOT:[LOG]LOGIN.COM`.
 - b. Execute `$rman target target_connect_string catalog catalog_connect_string`.

Oracle 8i:

Use `rcvcat` instead of `catalog` in the above syntax.

If you *do not* use the recovery catalog:

- On Windows: `ORACLE_HOME\bin\rman target Target_Database_Login nocatalog`
- On UNIX: `ORACLE_HOME/bin/rman target Target_Database_Login nocatalog`
- On OpenVMS:
 1. Run ORAUSER.COM using `$(OMNI$ROOT:[LOG]LOGIN.COM`.
 2. Execute `$rman target target_connect_string nocatalog`.

Target database login

The format of the *target database login* is `user_name/password@service`, where:

user_name 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 been granted Oracle SYSDBA or SYSOPER rights.

password must be the same as the password specified in the Oracle password file (`orapwd`), which is used for authentication of users performing database administration.

service is the name used to identify an SQL*Net server process for the target database.

Recovery catalog login

The format of the Recovery Catalog Database login is

user_name/password@service,

where the description of the user name and password is the same as for the login information to the target database. Note that the Oracle user specified here has to be the owner of the Oracle Recovery Catalog.

service is the name used to identify SQL*Net server process for the Recovery Catalog Database.

2. Allocate the Oracle channels.

Allocating a channel tells RMAN to initiate an Oracle Server process for backup, restore, or recovery on the Oracle target database. For example:

```
allocate channel 'dev_0' type 'disk';
```

or

```
allocate channel 'dev_1' type 'sbt_tape';
```

where you specify the backup directly to disk in the first case and directly to tape in the second case.

To use Data Protector backup media, specify the channel type `SBT_TAPE`. For this channel type, RMAN needs the Data Protector MML:

- On Windows and UNIX clients with Oracle8i, and on OpenVMS clients with any Oracle version, ensure that a symbolic link to the Data Protector MML exists.
- On Windows and UNIX clients with Oracle9i/10g, specify the path to the Data Protector MML at run time by setting the `SBT_LIBRARY` RMAN script parameter. For details, see [Step 3](#) on page 85.

If you specify more than a single `allocate channel` command, RMAN will establish multiple logon sessions and conduct multiple backup sets in parallel. This “parallelization” of backup and restore commands is handled internally by RMAN.

IMPORTANT:

On Windows, a maximum of 32 or 64 (if device is local) channels can be allocated.

3. Specify the parms operand:

```
parms 'SBT_LIBRARY=Path_to_Data_Protector_MML,  
ENV(OB2BARTYPE=Oracle8,  
OB2APPNAME=DB_NAME,OB2BARLIST=backup_specification_name)';
```

Note that the RMAN script will not work without the above parameters being specified in this form.

- On UNIX and Windows clients with Oracle8i, and on OpenVMS clients with any Oracle version, do not specify the `SBT_LIBRARY` parameter. Otherwise, you will receive a syntax error.
- On Windows and UNIX clients with Oracle9i/10g, set the `SBT_LIBRARY` parameter to point to the correct platform-specific Data Protector MML. The location and the filename of the Data Protector MML depend on the platform:

HP-UX, Solaris, and Linux: `/opt/omni/lib`

Other UNIX: `/usr/omni/lib`

Windows: `Data_Protector_home\bin`

Table 5 MML filenames on different platforms

Platforms	32-bit	64-bit
HP-UX	<code>libob2oracle8.sl</code>	<code>libob2oracle8_64bit.sl</code>
HP-UX on IA-64	<code>libob2oracle8.so</code>	<code>libob2oracle8_64bit.so</code>
Solaris	<code>libob2oracle8.so</code>	<code>libob2oracle8_64bit.so</code>
AIX	<code>libob2oracle8.a</code>	<code>libob2oracle8_64bit.a</code>
Other UNIX	<code>libob2oracle8.so</code>	<code>libob2oracle8_64bit.so</code>
Windows	<code>orasbt.dll</code>	<code>orasbt.dll</code>

For example, on 32-bit Solaris client, set

`SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so.`

4. Specify `format`:

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

Note that `%s:%t:%p` and the Oracle database name are required, whereas the backup specification is recommended.

For example, if you have created and saved a backup specification named `bspec1` for backing up an Oracle database identified by the Oracle instance called `inst1`, you would enter the following string:

```
format 'bspec1<inst1_%s:%t:%p>.dbf'
```

See the *Oracle Recovery Manager User's Guide and References* for information on substitution variables. The Oracle channel format specifies which Oracle backup specification to use for the backup.

5. Optionally, specify `backup incremental level`.

Note that a Data Protector full backup performs the same operation as an incremental level 0 backup type in the Oracle RMAN scripts. They both back up all the blocks that have ever been used.

This option is required if you want to use the backup as a base for subsequent incremental backups.

To run a backup using RMAN, start RMAN by running the following command from the `ORACLE_HOME` directory (if you use the recovery catalog):

Oracle 9i/10g:

- On Windows: `bin\rman target Target_Database_Login catalog Recovery_Catalog_Login`
- On UNIX: `bin/rman target Target_Database_Login catalog Recovery_Catalog_Login`
- On OpenVMS:
 1. Run `ORAUSER.COM` using `$(OMNI$ROOT): [LOG] LOGIN.COM`.
 2. Execute `$rman target target_connect_string catalog catalog_connect_string`.

Oracle 8i:

Use `rcvcat` instead of `catalog` in the above syntax.

Examples of the RMAN scripts

Some examples of RMAN scripts that must be executed from the `RMAN>` prompt are listed below:



NOTE:

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 clients with Oracle9i/10g.

Backing up a single channel

To back up the Oracle instance `ORACL`, using a backup specification named `ora1`, enter the following command sequence:

```
run {
allocate channel 'dev_0' type 'sbt_tape' parms
'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';
backup
incremental level 0
format 'oracl1<ORACL_%s:%t>.dbf' database;
}
```

Backing up three channels in parallel

The RMAN backup script for backing up the database by using three parallel channels for the same backup specification would look like this:

```
run {
allocate channel 'dev_0' type 'sbt_tape' parms
'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';
allocate channel 'dev_1' type 'sbt_tape' parms
'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';
allocate channel 'dev_2' type 'sbt_tape' parms
'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';
backup
incremental level 0
format 'ora1<ORACL_%s:%t>.dbf' database;
}
```

Backing up all archived logs and tablespaces

If you want to back up the Archived Redo Logs and the tablespace SYSTEM and RONA of the previous database using three parallel channels and a backup specification named `ora1`, the RMAN script should look like this:

```
run {
allocate channel 'dev_0' type 'sbt_tape' parms
'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';
allocate channel 'dev_1' type 'sbt_tape' parms
'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';
allocate channel 'dev_2' type 'sbt_tape' parms
'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';
backup
incremental level 0
format 'ora1<ORACL_%s:%t>.dbf'
tablespace SYSTEM, RONA
sql 'alter system archive log current'
format 'ora1<ORACL_%s:%f:%p>.dbf'
archivelog all;
}
```

Backing up particular archived logs

To back up all Archived Redo Logs from sequence #5 to sequence #105 and delete the Archived Redo Logs after backup of the instance named `ora1` is complete, run the following script:

```
run {
allocate channel 'dev_0' type 'sbt_tape' parms
'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';
allocate channel 'dev_1' type 'sbt_tape' parms
'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';
allocate channel 'dev_2' type 'sbt_tape' parms
'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';
backup
(archivelog sequence between 5 and 105 delete input
format 'ora1<ORACL_%s:%t:%p>.dbf');
}
```

If the backup fails, the logs are not deleted.

Backing up the flash recovery area

If you want to back up the Oracle 10g Flash Recovery Area using three parallel channels and a backup specification named `ora1`, the RMAN script should look like this:

```
run {
allocate channel 'dev_0' type 'sbt_tape' parms
'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';
allocate channel 'dev_1' type 'sbt_tape' parms
'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';
allocate channel 'dev_2' type 'sbt_tape' parms
'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';
backup
format 'ora1<ORACL_%s:%t>.dbf'
recovery area;
}
```

Including control file in a backup specification

The current control file is automatically backed up when the first datafile of the system tablespace is backed up. The current control file can also be explicitly included in a backup, or backed up individually. To include the current control file after backing up a tablespace named `COSTS`, run the following script:

```
run {
allocate channel 'dev_0' type 'sbt_tape' parms
'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';
allocate channel 'dev_1' type 'sbt_tape' parms
'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';
allocate channel 'dev_2' type 'sbt_tape' parms
'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';
backup
format 'ora1<ORACL_%s:%t>.dbf'
(tablespace COSTS current controlfile);
}
```

Backing up while allowing for some corrupted blocks

The set `maxcorrupt` command determines the number of corrupted blocks per datafile that can be tolerated by RMAN before a particular backup will fail.

If a backup specification named `ora1` backs up the database and allows for up to 10 corrupted blocks per datafile `/oracle/data1.dbs` (UNIX systems) or `C:\oracle\data1.dbs` (Windows systems), then the appropriate RMAN script would be:

On UNIX

```
run {
set maxcorrupt for datafile
'/oracle/data1.dbs' to 10;
allocate channel 'dev_0' type 'sbt_tape' parms
'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';
allocate channel 'dev_1' type 'sbt_tape' parms
'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';
allocate channel 'dev_2' type 'sbt_tape' parms
'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';
backup
incremental level 0
format 'ora1<ORACL_%s:%t>.dbf'
database;
}
```

On Windows

```
run {
set maxcorrupt for datafile
'C:\oracle\data1.dbs' to 10;
allocate channel 'dev_0' type 'sbt_tape' parms
'SBT_LIBRARY=Oracle_home\bin\orasbt.dll,
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';
allocate channel 'dev_1' type 'sbt_tape' parms
'SBT_LIBRARY=Oracle_home\bin\orasbt.dll,
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';
allocate channel 'dev_2' type 'sbt_tape' parms
'SBT_LIBRARY=Oracle_home\bin\orasbt.dll,
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';
backup
incremental level 0
format 'ora1<ORACL_%s:%t>.dbf'
database;
}
```

Restore

You can restore the database objects using:

- Data Protector GUI. See [“Restoring Oracle using the Data Protector GUI”](#) on page 93.
- RMAN. See [“Restoring Oracle using RMAN”](#) on page 114.

Restorable items

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

- Control files
- Datafiles
- Tablespaces
- Databases
- Recovery Catalog Databases

Duplicating databases

Using the Data Protector GUI, you can also **duplicate** a production database. See [“Duplicating an Oracle database”](#) on page 106.

Microsoft Cluster Server clients

Before you start restoring a cluster-aware Oracle server, take the Oracle Database resource offline using, for example, the **Cluster Administrator** utility. See [Figure 15](#) on page 92.

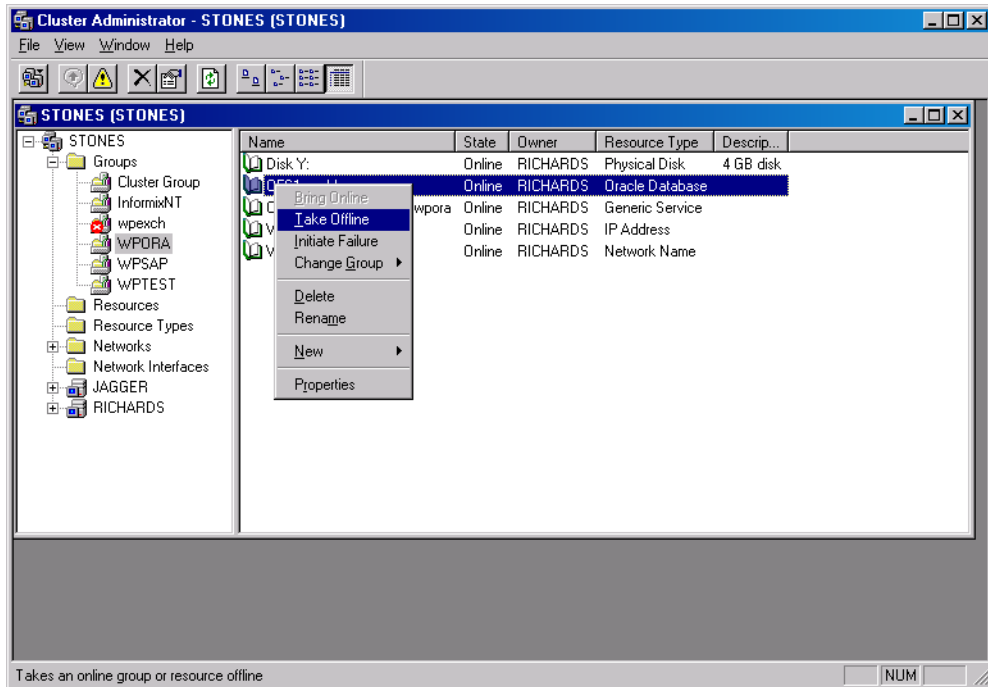


Figure 15 Taking the Oracle resource group offline

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

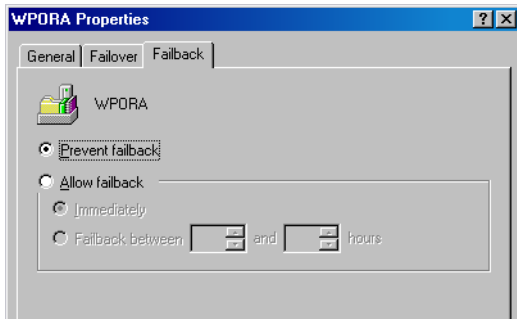


Figure 16 Checking properties

MC/ServiceGuard clients

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/ora9i/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.

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.

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.

If the recovery catalog was used:

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

If the recovery catalog was *not* used:

- **Oracle 8i:**
See problem [“The Recovery Catalog was lost and the control file cannot be restored from Data Protector managed backup” on page?](#)
- **Oracle 9i/10g:**
 1. Restore the control file from automatic backup.
If no automatic backup of the control file is available, see problem [“The Recovery Catalog was lost and the control file cannot be restored from Data Protector managed backup” on page?](#)
 2. Restore the 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 6 Required database states

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

¹When 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. In the Data Protector GUI, switch to the **Restore** context.
3. Under **Restore Objects**, expand **Oracle Server**, expand the client on which the database, for which you want to restore the recovery catalog, resides, and then click the database.

4. 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.

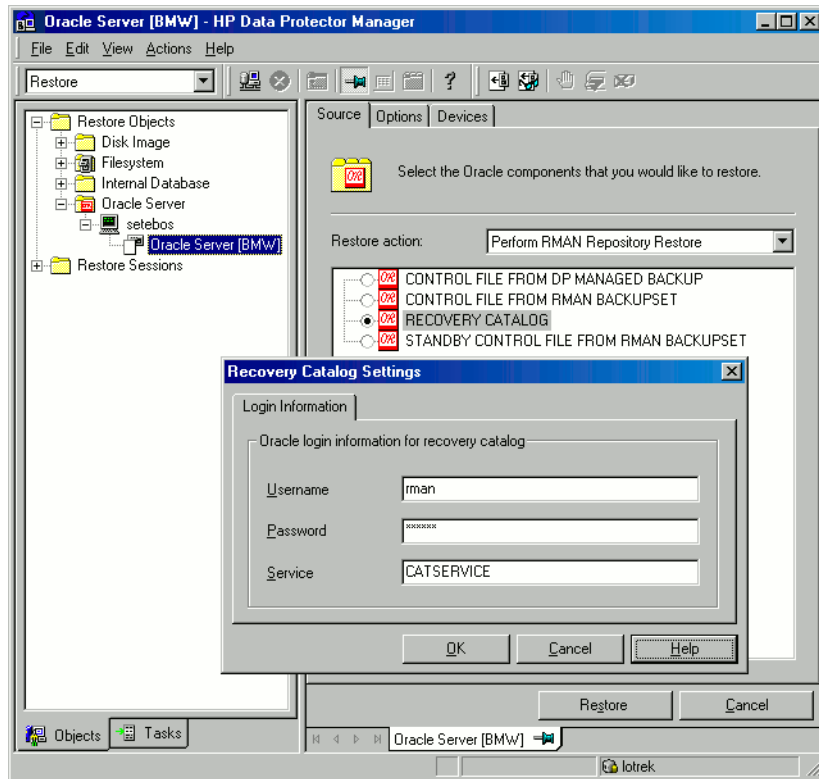


Figure 17 Recovery catalog settings dialog

5. 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”](#) on page 110.

6. 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: `Data_Protector_home\tmp`

HP-UX, Solaris, and Linux: `/var/opt/omni/tmp`

Other UNIX: `/usr/opt/omni/tmp`

OpenVMS: `OMNI$ROOT:[TMP]`

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 autobackup (`CONTROLFILE FROM RMAN AUTOBACKUP`)

This type of restore is *not* available with Oracle 8i.

The control file was automatically backed up by RMAN and the recovery catalog is *not* available.

 **IMPORTANT:**

Ensure that you have properly configured the RMAN autobackup and that the correct backup version is available. If the RMAN autobackup session is not found during the restore, the procedure is aborted. See the Oracle 9i/10g documentation on how to set up RMAN AUTOBACKUP.

- Restoring from RMAN backup set (`CONTROLFILE FROM RMAN BACKUPSET`)
The recovery catalog *is* required.
- **Oracle Data Guard (10g only):** Restoring standby control file from RMAN backup set (`STANDBY CONTROL FILE FROM RMAN BACKUPSET`)
If you restore a *standby* database (not using duplication), you must restore this type of control file.
This type of restore is available only in Oracle 10g standby configurations and if you selected the **CONTROL FILE FOR STANDBY** database object in the backup specification.

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”](#) on page 94.
2. In the Data Protector GUI, switch to the **Restore** context.
3. Under **Restore Objects**, expand **Oracle Server**, expand the client 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 client 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. See [“Restore, recovery, and duplicate options”](#) on page 110 for information.
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](#)” on page 95 and “[Restoring the control file](#)” on page 97.

To restore Oracle database objects:

1. **Oracle Data Guard:** If you restore a *standby* database, stop the managed recovery process (log apply services):

```
SQL> ALTER DATABASE RECOVER MANAGED STANDBY DATABASE  
CANCEL;
```

2. Put the database in the mount state. See “[Changing the database state](#)” on page 94.
3. In the Data Protector GUI, switch to the **Restore** context.
4. Under **Restore Objects**, expand **Oracle Server**, expand the client on which the database, for which you restore the database objects, resides, and then click the database.

5. 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”](#) on page 110.

 **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”](#) on page 114.

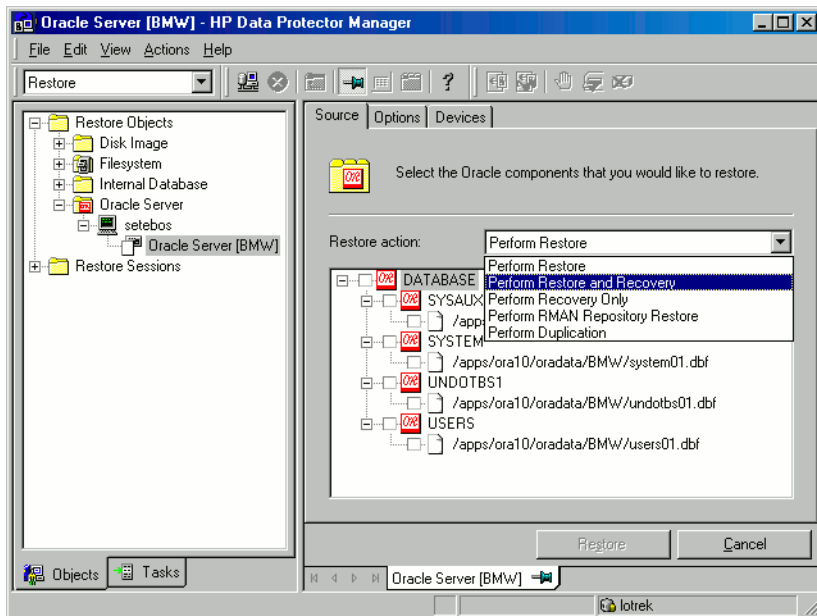


Figure 18 Source page

6. 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.

Oracle Data Guard: If you restore a *primary* database from a standby database backup or if you restore a *standby* database from a primary database backup, the location of datafiles can be different. In the **Restore as** dialog box, specify the appropriate location for each datafile.

 **TIP:**

The same can be done if you set the `DB_FILE_NAME_CONVERT` initialization parameter. This parameter captures all the target datafiles and converts them appropriately.

7. In the **Options** page, from the **Client** drop-down list, select the client 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.

Oracle Data Guard: If you restore the primary database, specify the login information for the primary database. If you restore the standby database, specify the login information for the standby database. Otherwise, the login information of the selected database will be used.

Set the other restore options. See [“Restore, recovery, and duplicate options”](#) on page 110 for information.

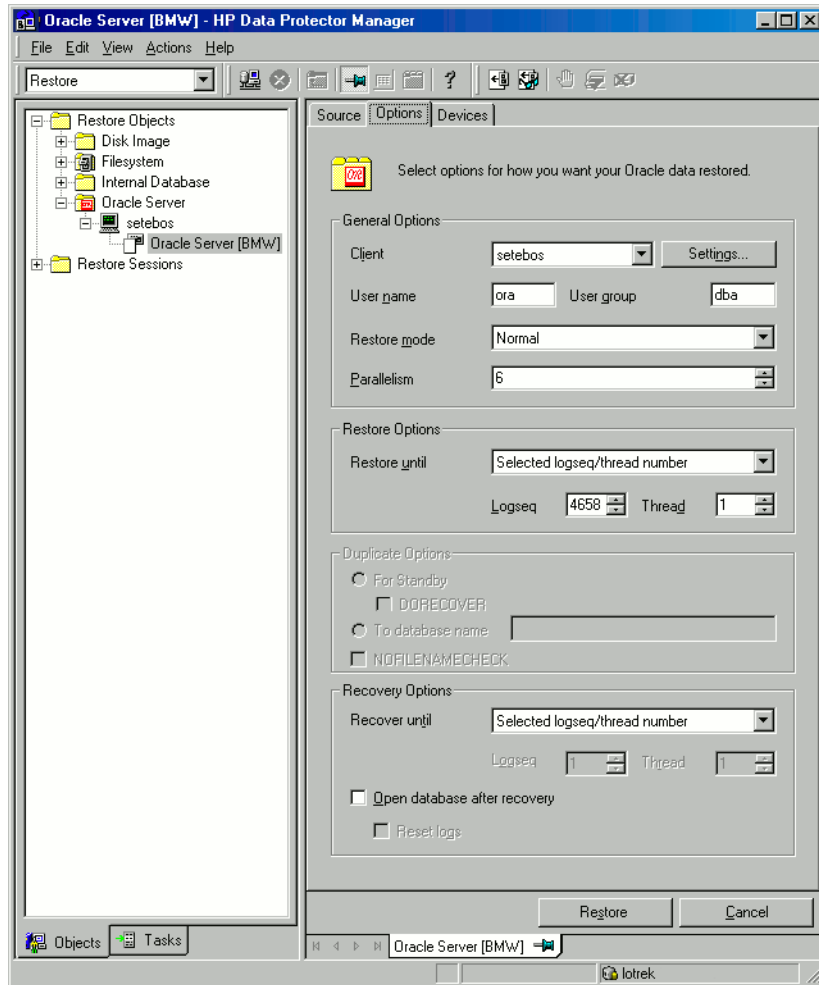


Figure 19 Options page

8. In the **Devices** page, select the devices to be used for the restore. You can restore using a device other than that used for backup, although Data Protector defaults to the original device on which the backup was made. To change the device from which an item is restored, select your desired device and click **Change**.

For more information on the **Devices** page, press **F1**.

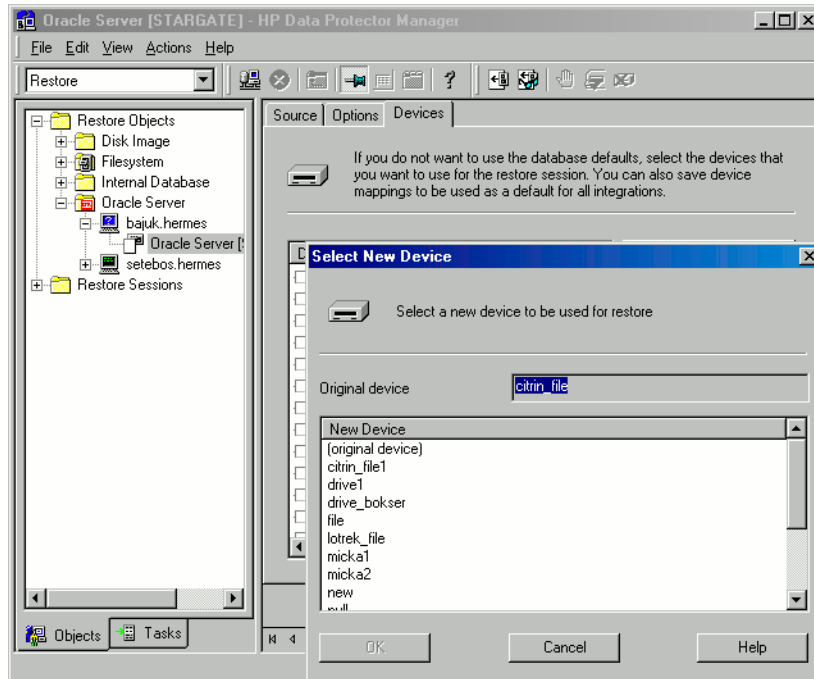


Figure 20 Devices page

9. 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:

Oracle 9i/10g:

```
rman target Target_Database_Login catalog  
Recovery_Catalog_Login
```

```
RMAN> RESET DATABASE;
```

```
RMAN> exit
```

Oracle 8i:

Use `rcvcat` instead of `catalog` in the above syntax.

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;
```

4. **Oracle Data Guard:** If you restored a *standby* database and if you have all archived redo logs on disk, restart the managed recovery process (log apply services):

```
SQL> ALTER DATABASE RECOVER MANAGED STANDBY DATABASE  
DISCONNECT;
```

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;
```

Restoring and recovering an Oracle database in Oracle Data Guard environment

Restoring and recovering a primary database

You can restore and recover a primary database from backups done on either a primary or standby database. The restore and recover is almost the same as restore and recover of a database in a standalone configuration. For information, see [“Restoring Oracle using the Data Protector GUI”](#) on page 93.

Restoring and recovering a standby database

You can restore and recover a standby database from backups of either a primary or standby database. The restore and recover is almost the same as restore and recover of a database in a standalone configuration. For information, see [“Restoring Oracle using the Data Protector GUI”](#) on page 93.

If the archived redo log files required for recovery are not accessible on disk, but only on tape, use RMAN to recover the restored datafiles to an SCN/log sequence greater than the last log applied to the standby database.

Obtain UNTIL_SCN:

```
SQL> SELECT MAX(NEXT_CHANGE#)+1 UNTIL_SCN FROM V$LOG_HISTORY
LH, V$DATABASE DB WHERE
LH.RESETLOGS_CHANGE#=DB.RESETLOGS_CHANGE# AND
LH.RESETLOGS_TIME = DB.RESETLOGS_TIME;
```

If the archived redo logs required for recovery are accessible on disk, restore only damaged datafiles and restart redo apply process.

If you have lost the entire standby database, it is better to perform **duplication** of the database (unless only a few damaged datafiles or tablespaces need to be restored).

Perform duplication of the database also when:

- Primary database control file was restored or recreated.
- Point-in-time recovery was performed on the primary database.
- Failover of database roles occurred.

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.

Limitation

- Database duplication is not supported using proxy copy backups of the primary database.

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

- 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 client where the selected database will be duplicated, put the Oracle auxiliary database instance in the nomount state. See [“Changing the database state”](#) on page 94.
2. In the Context List of the Data Protector GUI, click **Restore**.
3. Under **Restore Objects**, expand **Oracle Server**, expand the client on which the production database resides, and then click the production database which you want to duplicate. If there are several such clients, select the client 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 client 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](#)” on page 112 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.

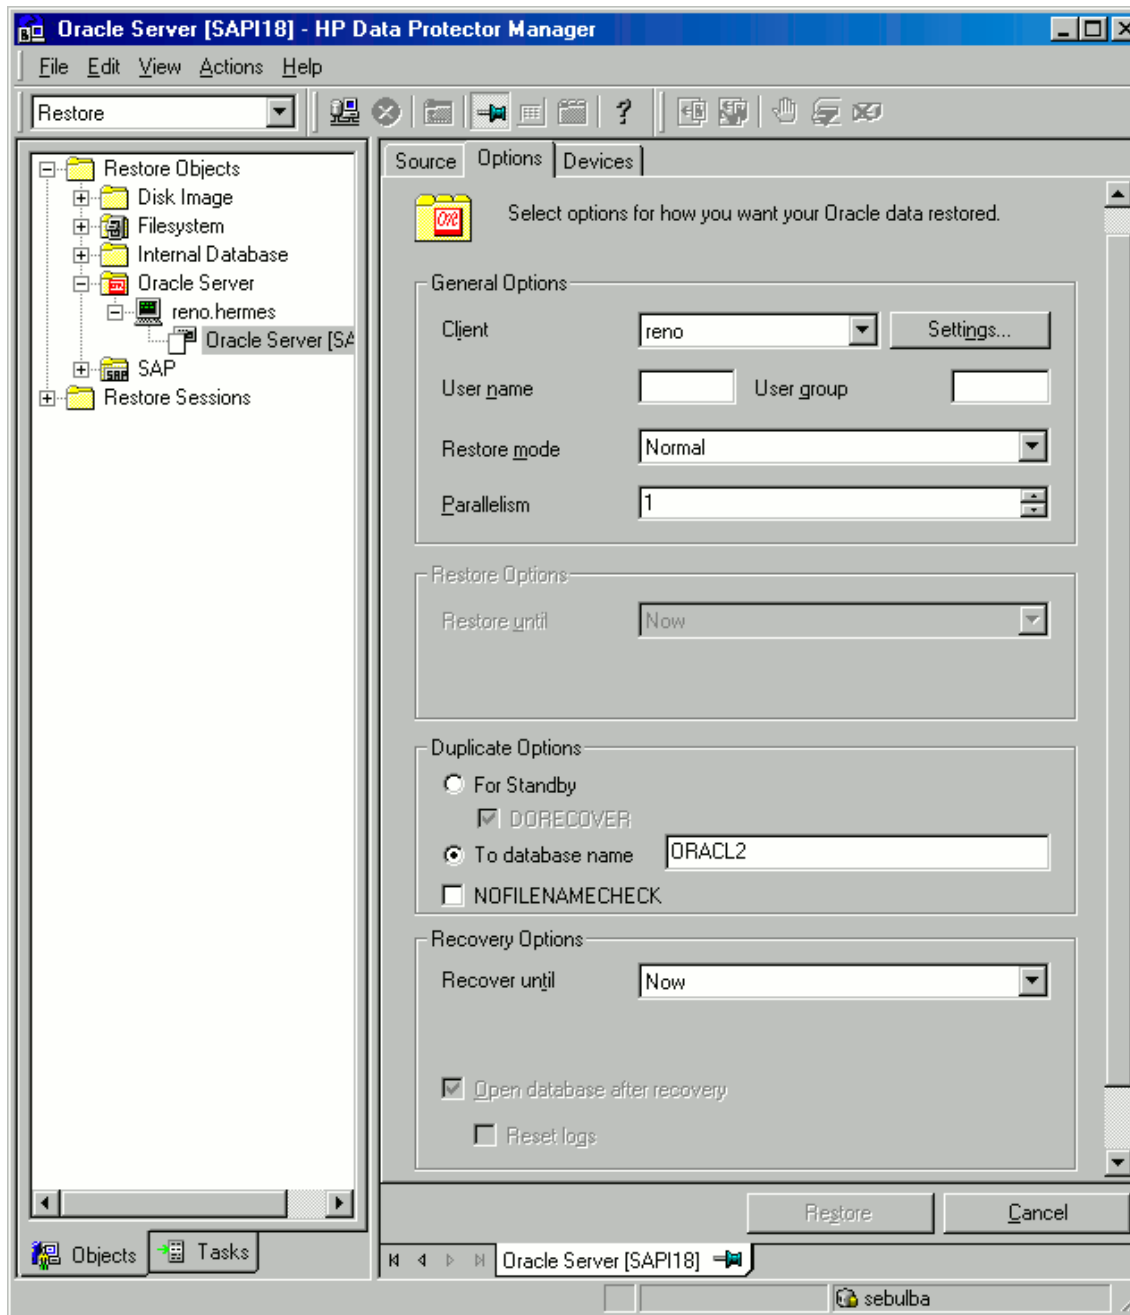


Figure 21 Oracle duplicate options

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”](#) on page 114.

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 client 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 client will be used.

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

User name (UNIX systems only)

Use this field to enter the Oracle user name. The user needs to be a member of the Oracle `DBA` group.

User group (UNIX systems only)

The User group the user in the **User name** field belongs to. This has to be the Oracle `DBA` group.



NOTE:

The user name and the user group must be the same as defined in the backup ownership. See [“Configuring Oracle users on UNIX and OpenVMS”](#) on page 44 for more information on this user and on how to identify it.

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, such as ZDB of Oracle 8i/9i.

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.

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 resides 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:

Oracle 9i/10g:

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

Oracle 8i:

Use `rcvcat` instead of `catalog` in the above syntax.

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”](#) on page 117
- [“Example of point-in-time restore”](#) on page 119
- [“Example of tablespace restore and recovery”](#) on page 121
- [“Example of datafile restore and recovery”](#) on page 124

- “[Example of archive log restore](#)” on page 128

The restore and recovery procedure of Oracle control files is a very delicate operation, which depends on whether you are using the recovery catalog or control file as a central repository and 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:

- If you use the recovery catalog 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](#)” on page 91 for details on how to restore the recovery catalog database.
- 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 or control files, 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 or control 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 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 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'
```

OpenVMS example

```
ORACLE_HOME=DKA400:[ORACLE9I]
ORACLE_TERM=HP
DB_NAME=PROD
```

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

Windows: PROD:Oracle_home\product\10.1.0:N

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

OpenVMS:

- **Oracle 9i:**

```
Oracle_home/oratab
```

```
TEST:/DKA400/ORACLE9I:N CAT:/DKA400/ORACLE9I:N
```

- **Oracle 8i:**

```
Oracle_home/rdbms/ORA_RDBMS_SIDS.DAT
```

```
VMS1 TEST TEST
```

```
VMS1 CAT CAT
```

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 clients with Oracle9i/10g, 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 [Step 3](#) on page 85.

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 clients.

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:

If you use the recovery catalog database, run:

Oracle 9i/10g:

- On Windows: `ORACLE_HOME\bin\rman target sys/manager@PROD catalog rman/rman@CATAL`
- On UNIX: `ORACLE_HOME/bin/rman target sys/manager@PROD catalog rman/rman@CATAL`
- On OpenVMS: `rman target sys/manager@PROD sys/manager@PROD catalog rman/rman@CAT`

Oracle 8i:

Use `rcvcat` instead of `catalog` in the above syntax.

If you do not use the recovery catalog database, run:

- On Windows: `ORACLE_HOME\bin\rman target sys/manager@PROD nocatalog`
- On UNIX: `ORACLE_HOME/bin/rman target sys/manager@PROD nocatalog`
- On OpenVMS: `rman target sys/manager@PROD nocatalog`

2. Start the full database restore and recovery:

```
run{
allocate channel 'dev1' type 'sbt_tape' parms
'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DE_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` directory.

2. Start the full database restore:

If you use the recovery catalog database, run:

Oracle 9i/10g:

- On Windows: `ORACLE_HOME\bin\rman target sys/manager@PROD catalog rman/rman@CATAL cmdfile=Data_Protector_home\tmp\restore_datafile`
- On UNIX: `ORACLE_HOME/bin/rman target sys/manager@PROD catalog rman/rman@CATAL cmdfile=/var/opt/omni/tmp/restore_datafile`

Oracle 8i:

Use `rcvcat` instead of `catalog` in the above syntax.

If you do not use the recovery catalog database, run:

- On Windows: `ORACLE_HOME\bin\rman target sys/manager@PROD nocatalog cmdfile=Data_Protector_home\tmp\restore_datafile`
- On UNIX: `ORACLE_HOME/bin/rman target sys/manager@PROD nocatalog 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:

If you use the recovery catalog database, run:

Oracle 9i/10g:

- On Windows: `ORACLE_HOME\bin\rman target sys/manager@PROD catalog rman/rman@CATAL`
- On UNIX: `ORACLE_HOME/bin/rman target sys/manager@PROD catalog rman/rman@CATAL`
- On OpenVMS: `rman target sys/manager@PROD sys/manager@PROD catalog rman/rman@CAT`

Oracle 8i:

Use `rcvcat` instead of `catalog` in the above syntax.

If you do not use the recovery catalog, run:

- On Windows: `ORACLE_HOME\bin\rman target sys/manager@PROD nocatlog`
- On UNIX: `ORACLE_HOME/bin/rman target sys/manager@PROD nocatlog`
- On OpenVMS: `rman target sys/manager@PROD nocatlog`

2. Start the point-in-time restore:

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

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:

If you use the recovery catalog database, run:

Oracle 9i/10g:

- On Windows: `ORACLE_HOME\bin\rman target sys/manager@PROD catalog rman/rman@CATAL cmdfile=Data_Protector_home\tmp\restore_PIT`
- On UNIX: `ORACLE_HOME/bin/rman target sys/manager@PROD catalog rman/rman@CATAL cmdfile=/var/opt/omni/tmp/restore_PIT`

Oracle 8i:

Use `rcvcat` instead of `catalog` in the above syntax.

If you do not use the recovery catalog, run:

- On Windows: `ORACLE_HOME\bin\rman target sys/manager@PROD nocatalog cmdfile=Data_Protector_home\tmp\restore_PIT`
- On UNIX: `ORACLE_HOME/bin/rman target sys/manager@PROD nocatalog 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:

If you use the recovery catalog database, run:

Oracle 9i/10g:

- On Windows: `ORACLE_HOME\bin\rman target sys/manager@PROD catalog rman/rman@CATAL`
- On UNIX: `ORACLE_HOME/bin/rman target sys/manager@PROD catalog rman/rman@CATAL`
- On OpenVMS: `rman target sys/manager@PROD sys/manager@PROD catalog rman/rman@CAT`

Oracle 8i:

Use `rcvcat` instead of `catalog` in the above syntax.

If you do not use the recovery catalog, run:

- On Windows: `ORACLE_HOME\bin\rman target sys/manager@PROD nocatalog`
- On UNIX: `ORACLE_HOME/bin/rman target sys/manager@PROD nocatalog`
- On OpenVMS: `rman target sys/manager@PROD nocatalog`

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:

```
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;
}
```

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

```
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;
}
```

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.

If you use the recovery catalog database, run:

Oracle 9i/10g:

- On Windows: `ORACLE_HOME\bin\rman target sys/manager@PROD catalog rman/rman@CATAL cmdfile=Data_Protector_home\tmp\restore_TAB`
- On UNIX: `ORACLE_HOME/bin/rman target sys/manager@PROD catalog rman/rman@CATAL cmdfile=/var/opt/omni/tmp/restore_TAB`

Oracle 8i:

Use `rcvcat` instead of `catalog` in the above syntax.

If you do not use the recovery catalog, run:

- On Windows: `ORACLE_HOME\bin\rman target sys/manager@PROD nocatalog cmdfile=Data_Protector_home\tmp\restore_TAB`
- On UNIX: `ORACLE_HOME/bin/rman target sys/manager@PROD nocatalog 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.

If you use the recovery catalog database, run:

Oracle 9i/10g:

- On Windows: `ORACLE_HOME\bin\rman target sys/manager@PROD catalog rman/rman@CATAL`
- On UNIX: `ORACLE_HOME/bin/rman target sys/manager@PROD catalog rman/rman@CATAL`
- On OpenVMS: `rman target sys/manager@PROD sys/manager@PROD catalog rman/rman@CAT`

Oracle 8i:

Use `rcvcat` instead of `catalog` in the above syntax.

If you do not use the recovery catalog database, run:

- On Windows: `ORACLE_HOME\bin\rman target sys/manager@PROD nocatalog`
- On UNIX: `ORACLE_HOME/bin/rman target sys/manager@PROD nocatalog`
- On OpenVMS: `rman target sys/manager@PROD nocatalog`

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

```
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;
}
```

Windows

```
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;
}
```

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

UNIX

```
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;
}

```

Windows

```

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;
}

```

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

1. Create a file `restore_dbf` the `/var/opt/omni/tmp` or `Data_Protector_home\tmp` (Windows systems) directory.

2. Start the datafile restore:

If you use the recovery catalog database, run:

Oracle 9i/10g:

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

Oracle 8i:

Use `rcvcat` instead of `catalog` in the above syntax.

If you do not use the recovery catalog database, run:

- On Windows: `ORACLE_HOME/bin\rman target sys/manager@PROD nocatalog cmdfile=Data_Protector_home\tmp\restore_dbf`
- On UNIX: `ORACLE_HOME/bin/rman target sys/manager@PROD nocatalog cmdfile=/var/opt/omni/tmp/restore_dbf`

Example of archive log restore

To restore an archive log:

1. Login to the Oracle RMAN:

If you use the recovery catalog database, run:

Oracle 9i/10g:

- On Windows: `ORACLE_HOME\bin\rman target sys/manager@PROD catalog rman/rman@CATAL`
- On UNIX: `ORACLE_HOME/bin/rman target sys/manager@PROD catalog rman/rman@CATAL`
- On OpenVMS: `rman target sys/manager@PROD sys/manager@PROD catalog rman/rman@CAT`

Oracle 8i:

Use `rcvcat` instead of `catalog` in the above syntax.

If you do not use the recovery catalog database, run:

- On Windows: `ORACLE_HOME\bin\rman target sys/manager@PROD nocatalog`
- On UNIX: `ORACLE_HOME /bin/rman target sys/manager@PROD nocatalog`
- On OpenVMS: `rman target sys/manager@PROD nocatalog`

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:

If you use the recovery catalog database, run:

Oracle 9i/10g:

- On Windows: `ORACLE_HOME\bin\rman target sys/manager@PROD catalog rman/rman@CATAL cmdfile=Data_Protector_home\tmp\restore_arch`
- On UNIX: `ORACLE_HOME/bin/rman target sys/manager@PROD catalog rman/rman@CATAL cmdfile=/var/opt/omni/tmp/restore_arch`

Oracle 8i:

Use `rcvcat` instead of `catalog` in the above syntax.

If you do not use the recovery catalog database, run:

- On Windows: `ORACLE_HOME\bin\rman target sys/manager@PROD nocatalog cmdfile=Data_Protector_home\tmp\restore_arch`
- On UNIX: `ORACLE_HOME/bin/rman target sys/manager@PROD nocatalog cmdfile=/var/opt/omni/tmp/restore_arch`

Restoring Oracle using CLI

Restoring the recovery catalog

Data Protector can restore the binary file which contains the logical backups of the Oracle recovery catalog. This file is made using the Oracle Export utility, which creates it by reading the Oracle database and writing the output to the binary file, which is then backed up by Data Protector.

This file can be restored back to the disk and then imported to the Oracle database by the Oracle Import utility.

To restore the Oracle recovery catalog, proceed as follows:

1. Login to the Oracle Recovery Catalog Database. Ensure that the recovery catalog database exists and that the recovery catalog is *not* present. If necessary, remove the recovery catalog using the RMAN command `DROP CATALOG`.

Identify the Oracle recovery catalog owner. If necessary, create the Oracle user.

On UNIX, Data Protector determines the Oracle login information for the recovery catalog from the Data Protector Oracle configuration files.

2. Set the `OB2APPNAME` environment variable. Its value must be set to the name of the target database (`DB_NAME`), not of the Oracle recovery catalog:

Windows: `set OB2APPNAME=DB_NAME`

UNIX:

- if you are using an `sh` - like shell, run:
`OB2APPNAME="DB_NAME"`
`export OB2APPNAME`
- if you are using a `csh` - like shell, run:
`setenv OB2APPNAME "DB_NAME"`

OpenVMS: `$DEFINE/log/process ob2appname DB_NAME`

3. Run:

Windows: From the `Data_Protector_home\bin` directory:

```
perl -I..\lib\perl ob2rman.pl -restore_catalog -session  
Session_ID [-apphost application_hostname]
```

HP-UX, Solaris, and Linux:

```
/opt/omni/lbin/ob2rman.pl -restore_catalog -session  
session_ID [-apphost application_hostname]
```

Other UNIX:

```
/usr/omni/bin/ob2rman.pl -restore_catalog -session  
session_ID [-apphost application_hostname]
```

OpenVMS:

```
$_OMNIROOT: [BIN]OMNIROOT.COM$ob2rman-restore_catalog  
-session session_ID [-apphost application_hostname]
```

Provide the `Session_ID` of the backup session. In case of object copies, do not use the copy session ID, but the object's backup ID, which equals the object's backup session ID.

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, this file must be in 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"
```

Disaster recovery

Disaster recovery is a very complex process that involves products from several vendors. As such, successful disaster recovery depends on all the vendors involved. The information provided here is intended to be used as a guideline.

Check the instructions from the database/application vendor on how to prepare for a disaster recovery. Also see the *HP Data Protector disaster recovery guide* for instructions on how to approach system disaster recovery using Data Protector.

This is a general procedure for recovering an application:

1. Complete the recovery of the operating system.
2. Install, configure, and initialize the database/application so that data on the Data Protector media can be loaded back to the system. Consult the documentation from the database/application vendor for a detailed procedure and the steps needed to prepare the database.
3. Ensure that the database/application server has the required Data Protector client software installed and is configured for the database/application. Follow the procedures in this chapter and in the section. See also the section of this manual about the Data Protector Restore GUI for Oracle for information about using this to restore database items, ["Restoring Oracle using the Data Protector GUI"](#) on page 93.

4. Start the restore. When the restore is complete, follow the instructions from the database/application vendor for any additional steps required to bring the database back online.

Monitoring sessions

During a backup, system messages are sent to the Data Protector monitor. You can monitor the backup session from any Data Protector client on the network where the Data Protector User Interface is installed.

Monitoring current sessions

To monitor a currently running session using the Data Protector GUI:

1. In the Context List, click **Monitor**.
In the Results Area, all currently running sessions are listed.
2. Double-click the session you want to monitor.

Clearing sessions

To remove all completed or aborted sessions from the Results Area of the **Monitor** context:

1. In the Scoping Pane, click **Current Sessions**.
2. In the **Actions** menu, select **Clear Sessions**. Or click the **Clear Sessions** icon on the toolbar.

To remove a particular completed or aborted session from the current sessions list, right-click the session and select **Remove From List**.



NOTE:

All completed or aborted sessions are automatically removed from the Results Area of the **Monitor** context if you restart the Data Protector GUI.

Monitoring tools

The progress of backups and restores can also be monitored by querying the Oracle target database using the following SQL statement:

```
select * from v$SESSION_LONGOPS where  
compnam='dbms_backup_restore';
```

For detailed information on a completed or aborted session, see “[Viewing previous sessions](#)” on page 134.

Viewing previous sessions

To view a previous session using the Data Protector GUI, proceed as follows:

1. In the Context List, click **Internal Database**.
2. In the Scoping Pane, expand **Sessions** to display all the sessions stored in the IDB.

The sessions are sorted by date. Each session is identified by a session ID consisting of a date in the YY/MM/DD format and a unique number.

3. Right-click the session and select **Properties** to view details on the session.
4. Click the **General**, **Messages**, or **Media** tab to display general information on the session, session messages, or information on the media used for this session, respectively.

Details about Oracle backup and restore sessions are also written in the following logs on the Oracle Server system:

- Data Protector writes the logs in:
 - Windows:** `Data_Protector_home\log\oracle8.log`
 - HP-UX, Solaris, and Linux:** `/var/opt/omni/log/oracle8.log`
 - Other UNIX:** `usr/omni/log/oracle8.log`
 - OpenVMS:** `OMNI$ROOT:[LOG]ORACLE8.LOG`
- Oracle writes the logs in the `Oracle_user_dump_directory\sbtio.log` file.

Using Oracle after removing the Data Protector Oracle integration

When you remove the Data Protector Oracle integration from a UNIX or Windows client with Oracle8i, or from an OpenVMS client with any Oracle version, you also need to release the symbolic link between the Data Protector MML and Oracle Server. If this is not done, the Oracle Server cannot be started.

Removing the link on HP-UX systems

To remove the Data Protector Oracle integration link on HP-UX systems:

1. Change to the `ORACLE_HOME/lib` directory:

```
cd ORACLE_HOME/lib (32-bit Oracle),  
cd ORACLE_HOME/lib64 (64-bit Oracle 8i) or  
cd ORACLE_HOME/lib (64-bit Oracle 9i/10g).
```

2. If the `libobk.sl.orig` file exists in the `ORACLE_HOME/lib` directory, run:

```
mv libobk.sl.orig libobk.sl
```

where `libobk.sl.orig` is the Oracle soft link as it existed before configuring the integration.

Removing the link on Solaris and other UNIX systems

To remove the Data Protector Oracle integration link on Solaris and other UNIX systems:

1. Change to the `ORACLE_HOME/lib` directory:

```
cd ORACLE_HOME/lib (32-bit Oracle),  
cd ORACLE_HOME/lib64 (64-bit Oracle 8i) or  
cd ORACLE_HOME/lib (64-bit Oracle 9i/10g).
```

2. If the `libobk.so.orig` file exists in the `ORACLE_HOME/lib` directory, execute the following command:

```
mv libobk.so.orig libobk.so
```

where `libobk.so.orig` is the Oracle soft link as it existed before configuring the integration.

Removing the link on OpenVMS systems

Oracle 8i

Relink the Oracle 8i binary using the default linking procedure. See the Oracle documentation for details.

For Oracle 9i running on OpenVMS, re-linking the Oracle 9i binary after uninstalling the Data Protector Oracle integration on an Oracle Server is not required.

Removing the link on Windows systems

Delete the Data Protector `orasbt.dll` file from the `ORACLE_HOME\bin` directory.

Oracle RMAN metadata and Data Protector Media Management Database synchronization

This section describes how to synchronize the Oracle RMAN metadata with the Data Protector Media Management Database.

The RMAN metadata contains information about the target database. RMAN uses this information for all backup, restore and maintenance operations. The metadata can be stored either in the recovery catalog database or in the control files.

Data Protector is the media manager that Oracle needs to perform tape storage backups and restores.

Data Protector has its own data protection policy that is not automatically synchronized with Oracle RMAN metadata. To have both catalogs synchronized, run the following command using RMAN:

```
allocate channel for maintenance type 'sbt_tape' parms
'SBT_LIBRARY=Path_to_Data_Protector_MML,
ENV=(OB2MAINTENANCE=1)';

crosscheck backup completed after "TO_DATE('01/13/06
10:30:00', 'MM/DD/YY HH24:MI:SS')";

release channel;
```

The `SBT_LIBRARY` parameter should be specified only on UNIX and Windows clients with Oracle9i/10g.

RMAN checks every backup piece in the repository and queries the MMDB for the availability of that backup piece. RMAN then mark the backup piece as expired or available, depending on media availability. Note that in the above example, RMAN does not delete backup pieces that are reported as expired by the MMDB, but instead marks them as expired.

In order to delete expired backup objects from the recovery catalog database, run the following command using RMAN:

```
delete expired backup;
```

See the *Oracle Recovery Manager User's Guide and References* for more details on recovery catalog maintenance.

 **TIP:**

It is recommended that synchronization be performed in the following cases:

- after a Data Protector import or export of media with Oracle objects and
 - whenever protection for media with Oracle objects has expired.
-

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](#)” on page 146 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*.

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:

- Verify that you can access the Oracle target database and that it is opened as follows:

1. UNIX: 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"
```

Windows: 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_tablespaces;  
exit
```

If this fails, open the Oracle target database.

- 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](#) on page 138.

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](#) on page 138.
 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.

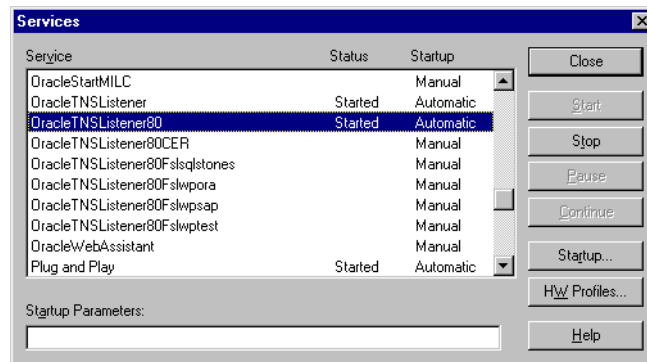


Figure 22 Checking the status of the Oracle listener

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`).

- Verify that the Oracle target database and the recovery catalog database are configured to allow remote connections with the system privileges:

1. Export or set the `ORACLE_HOME` and `DB_NAME` variables as described in [Step 1](#) on page 138.

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

```
sqlplus /nolog
```

3. Start SQL*Plus and type:

```
connect Target_Database_Login as SYSDBA
```

```
exit
```

```
and
```

```
sqlplus connect Recovery_Catalog_Login as SYSDBA
```

```
exit
```

Repeat the procedure using `SYSOPER` instead of `SYSDBA`.

If this fails, see the Oracle documentation for instructions about how to set up the password file and any relevant parameters in the `initDB_NAME.ora` file.

- 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](#) on page 138.

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, 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](#) on page 138.

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

Oracle 9i/10g:

```
rman target Target_Database_Login catalog
Recovery_Catalog_Login cmd_file=rman_script
```

Oracle 8i:

Use `rcvcat` instead of `catalog` in the above syntax.

If you do not use the recovery catalog:

1. Export or set the `ORACLE_HOME` variable as described in [Step 1](#) on page 138.
2. Start RMAN from the `bin` directory in the `ORACLE_HOME` directory:

```
rman target Target_Database_Login nocatalog
cmd_file=rman_script
```

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';release channel 'dev0';
}
```

If this fails, see the Oracle documentation for details on 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.
See the *HP Data Protector installation and licensing guide* for details.
- 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. See the `util_cmd` man page or the *HP Data Protector command line interface reference*

for information on setting the variables in the Data Protector Oracle configuration files.

- On UNIX clients with Oracle8i, verify that the Data Protector Oracle integration MML is linked with the Oracle executable:
 1. Export `ORACLE_HOME` and `DB_NAME` as described in [Step 1](#) on page 138.
 2. Use the following command to check if the `libob2oracle8.sl` (`libob2oracle8_64bit.sl`) file is linked with the Oracle executable. Note that on HP-UX IA-64 and Solaris systems, the extension for MML is `.so`, and on AIX, the extension is `.a`:

HP-UX:

```
/usr/bin/chatr ORACLE_HOME/bin/oracle (32-bit Oracle)
```

```
/usr/ccs/bin/ldd ORACLE_HOME/bin/oracle (64-bit Oracle)
```

Solaris:

```
/usr/bin/ldd -s ORACLE_HOME/bin/oracle
```

Other UNIX:

```
/usr/bin/ldd -s ORACLE_HOME/bin/oracle
```

IBM AIX:

```
/usr/bin/dump -H ORACLE_HOME/bin/oracle (32-bit Oracle)
```

```
/usr/bin/dump -H -X64 ORACLE_HOME/bin/oracle (64-bit Oracle)
```

Linux:

```
/usr/bin/ldd ORACLE_HOME/bin/oracle
```

The output must state that the respective MML is required by the Oracle executable. The following is an extract from the command output on HP-UX:

```
bin/oracle:
shared executable
shared library dynamic path search:
SHLIB_PATH enabled second
embedded path disabled first Not Defined
shared library list:
static /opt/omni/lib/libob2oracle8.sl(libob2oracle8_64bit.sl)
dynamic /usr/lib/librt.2
dynamic /usr/lib/libnss_dns.1
dynamic /usr/lib/libdld.2
```

The line starting with `SHLIB_PATH` should be as presented in the example above. If this line is different, enable MML dynamic path as follows:

```
/usr/bin/chrtr +s enable ORACLE_HOME/bin/oracle
```

On Solaris, HP-UX (64-bit), and other UNIX systems, LD_LIBRARY_PATH is used instead of SHLIB_PATH as on HP-UX (32-bit).

The following is an extract from the command output on other UNIX systems:

```
find library=/usr/omni/lib/libob2oracle8.so; required by /app/oracle8/product/8.0.4/bin/oracle
/usr/omni/lib/libob2oracle8.so
find library=libob2.so.1; required by /app/oracle8/product/8.0.4/bin/oracle
```

Figure 23 Output of the ldd command on other UNIX systems:

- On Windows clients with Oracle8i, verify that the Data Protector Oracle integration MML is loaded:
 - a. Switch to the `ORACLE_HOME\bin` directory and right-click `orasbt.dll`.
 - b. Select **Properties** and click the **Version** tab from the `orasbt.dll Properties` window. In **Description**, you should see the file described as a part of the Data Protector integration.
- 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.
- On Windows, 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 Oracle server system:

HP-UX, Solaris, and Linux: `/var/opt/omni/log/debug.log`
Other UNIX: `/usr/omni/log/debug.log`
Windows: `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:

HP-UX, Solaris, and Linux: `/etc/opt/omni/client/cell_server`

Other UNIX: `/usr/omni/config/cell/cell_server`

Windows: `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
-object:object_name
-version:object_version-bar:backup_specification_name
```

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 owner of the backup specification (in case of failed backup) or the restore session (in case of failed restore) is the Oracle backup owner, and that this user belongs to the Data Protector `operator` or `admin` 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:**
As the owner of the restore session, 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 hangs 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”](#) on page 114.

- 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

Data Protector reports "12:8422" error when using Data Protector Oracle integration after an upgrade of Oracle 8i to Oracle 9i

After Oracle 8i is upgraded to Oracle 9i, the following error is returned during the configuration of Oracle instance or during the backup:

```
*RETVAl*8422
```

Action

Rename the Oracle 8i `svrmgr1` binary to something else so that Data Protector will not find it. The Oracle upgrade process from Oracle 8i to Oracle 9i does not remove the Oracle 8i `svrmgr1` binary, rather it changes its permissions. Once the `svrmgr1` binary is renamed, Data Protector will use Oracle 9i `sqlplus`, as it should, to complete the operations correctly.

Problem

Data Protector reports errors when calling SYS.LT_EXPORT_PKG.schema_inf_exp during Oracle 9i/10g 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 UNIX, 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:

Oracle 9i/10g:

```
rman target Target_Database_Login catalog  
Recovery_Catalog_Login
```

```
RMAN> RESET DATABASE;
```

```
RMAN> exit
```

Oracle 8i:

Use `rcvcat` instead of `catalog` in the above syntax.

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

“Binary util_orarest is missing” error message is displayed when browsing Oracle 9i database for restore on Linux

The following error message is displayed when browsing *Oracle9i* database for restore on Linux:

```
Binary util_orarest is missing. Cannot get information from
the remote host.
```

Action

Two actions are possible:

- Start the following command:

```
/usr/omni/bin/util_orarest.exe -objs0 DB_NAME
```

If the command core dumps, make sure that the `libc` version is 2.3.2-23 or higher. This should eliminate the problem.

- Replace the `util_orarest.exe` utility with the new `util_orarest9.exe` (both located in the `/usr/omni/bin` directory on Linux):

1. Rename the `util_orarest.exe` to `util_orarest.exe.orig`
2. Rename the `util_orarest9.exe` to `util_orarest.exe`

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 (for Oracle 9i/10g), and the control file cannot be restored from Data Protector managed backup. A valid control file backup exists on tape.

Action

- For Oracle 8i, restore the control file from RMAN backup set with the following SQL script:

```
DECLARE
devtype varchar2(256);
done boolean;
BEGIN
devtype:=dbms_backup_restore.deviceallocate('sbt_tape',
params=>'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME,
OB2BARHOSTNAME=hostname) ');
dbms_backup_restore.restoresetdatafile;
dbms_backup_restore.restorecontrolfileto('/tmp/tmp.cf');
dbms_backup_restore.restorebackuppiece('backup piece handle',
```

```
done=>done);  
END;
```

For the *backup piece handle* search the Data Protector internal database and session outputs of previous backup sessions.

Use the following RMAN script to copy the control file, mount and restore the database, and perform a database recovery:

```
run {  
allocate channel 'dev_0' type disk;  
replicate controlfile from '/tmp/foo.cf';  
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 Oracle 9i/10g, 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

Shared library that provides thread local storage cannot be loaded

The problem occurs with Oracle8i on HP-UX 11.11.

When, during restore, Data Protector attempts to dynamically load a shared library that provides thread local storage, an error similar to the following is displayed:

```
Can't dlopen() a library containing Thread Local Storage:  
ORACLE_HOME/JRE/lib/PA_RISC/native_threads/libjava.sl
```

The problem occurs when the Radius Authentication Adapter is installed. In this case, `libc1ntsh.sl` is dynamically linked with the library `libjava.sl` that provides thread local storage.

Action

Uninstall the Radius Authentication Adapter to remove `libjava.sl` from the list of dynamic libraries for `libc1ntsh.sl`. See *Oracle MetaLink, DOC ID: 113395.1* for information on how to uninstall the Radius Authentication Adapter.

Problem

Binary util_orarest is missing

The message below sometimes appears when you are restoring database items to a new host:

```
"Binary util_orarest is missing. Cannot get information from  
the remote host."
```

Action

To resolve the problem:

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.

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 client, restart the Data Protector `Inet` service for the settings to take effect.

2 Integrating SAP R/3 and Data Protector

Introduction

This chapter explains how to configure and use the Data Protector SAP R/3 integration (**SAP R/3 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:

Table 7 Backup types

Full	Backs up all the selected SAP R/3 objects.
Incr	Oracle RMAN backup incremental level 1 (available only if you use Oracle RMAN). Backs up changes made to the selected Oracle data files since the last Full backup.

You can start backups using:

- The Data Protector user interface
- The SAP BRTOOLS interface

Data Protector supports only a filesystem restore. You can restore SAP R/3 files:

- To the original location
- To another client
- To another directory

You can restore Data Protector backups using:

- The Data Protector user interface
- The SAP BRTOOLS user interface

When the restore completes, you can recover the database to a specific point in time using the SAP BRTOOLS interface.

This chapter provides information specific to the Data Protector SAP R/3 integration. For general Data Protector procedures and options, see the *online Help*.

Integration concepts

This integration links SAP backup and restore tools (BR*Tools) with Data Protector. Because the SAP R/3 application runs on top of Oracle databases, the SAP R/3 backup objects are very similar to those of Oracle. The main difference is that SAP backup utilities hide the database from Data Protector, which sees those objects as plain files.

SAP tools can be started using the Data Protector interface or the SAP BRTOOLS interface.

Table 8 SAP backup and restore utilities

BRBACKUP	Backs up control files, data files, and online redo log files. Additionally, saves the profiles and logs relevant for a particular backup session.
BRARCHIVE	Backs up offline (archived) redo logs, written by Oracle to the archiving directory.
BRRESTORE	Restores data backed up with BRBACKUP and BRARCHIVE.

You can back up Oracle data files in two different modes:

backint	Data is backed up using the Data Protector SAP R/3 integration.
RMAN	Data is backed up using the Oracle Recovery Manager (RMAN). The main benefit of the RMAN mode is that you can back up Oracle database incrementally.

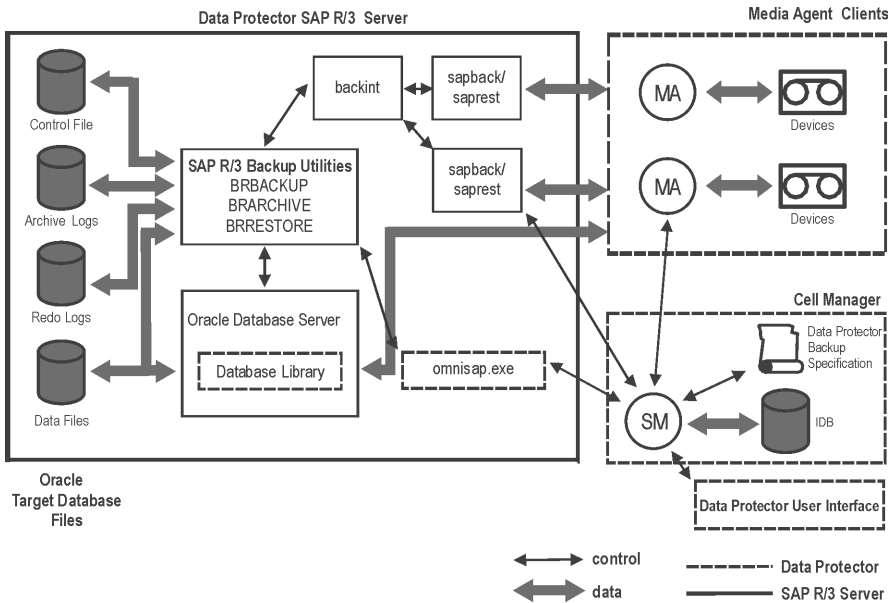


Figure 24 SAP R/3 architecture

Legend	
SM	Data Protector Session Manager: Backup Session Manager during backup and Restore Session Manager during restore.
Database Library	A set of Data Protector executables that enable data transfer between Oracle Server and Data Protector. Required only if Oracle data files are backed up in the RMAN mode.
MA	Data Protector General Media Agent.
Backup Specification	A list of objects to be backed up, backup devices, and options to be used.
IDB	The Data Protector Internal Database.
backint	Backup interface between Data Protector and SAP R/3 application. It is started by SAP tools: BRBACKUP or BRARCHIVE uses BACKINT to pass a backup request to Data Protector. BRRESTORE uses BACKINT to trigger Data Protector to restore the requested files.

Legend	
sapback/saprest	Program that performs the actual backup/restore of files.
omnisap.exe	Data Protector program that starts the SAP backup tools.

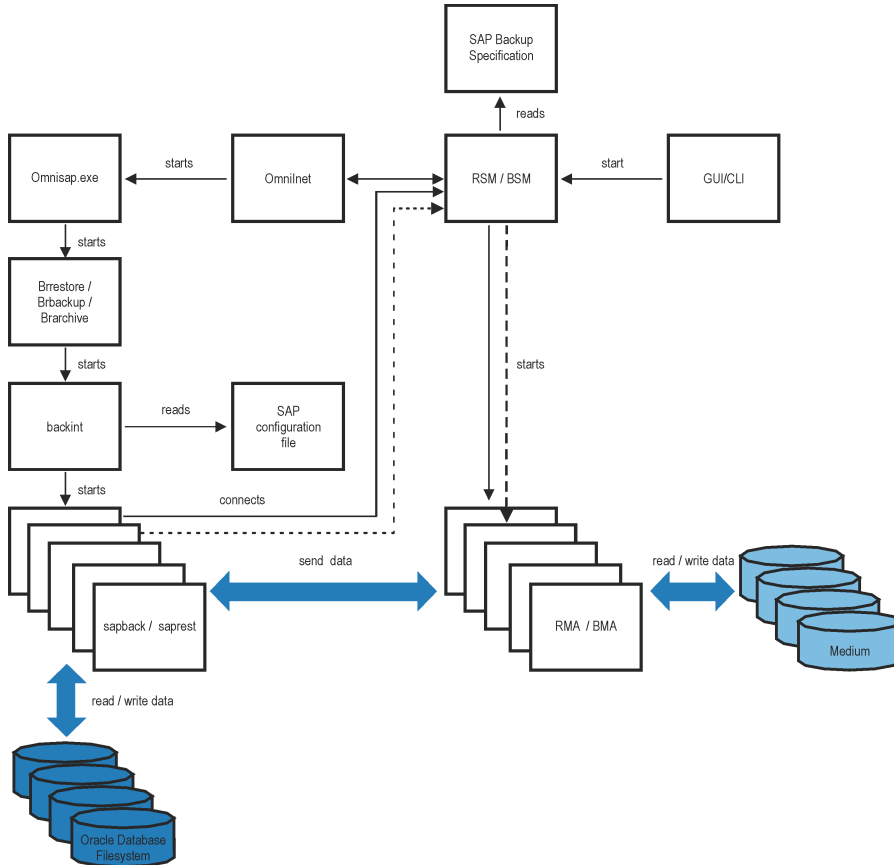


Figure 25 SAP R/3 architecture: backint mode

Legend	
BSM/RSM	Data Protector Backup/Restore Session Manager
BMA/RMA	Data Protector Backup/Restore Media Agent

Legend

GUI/CLI

Data Protector graphical/command-line user interface

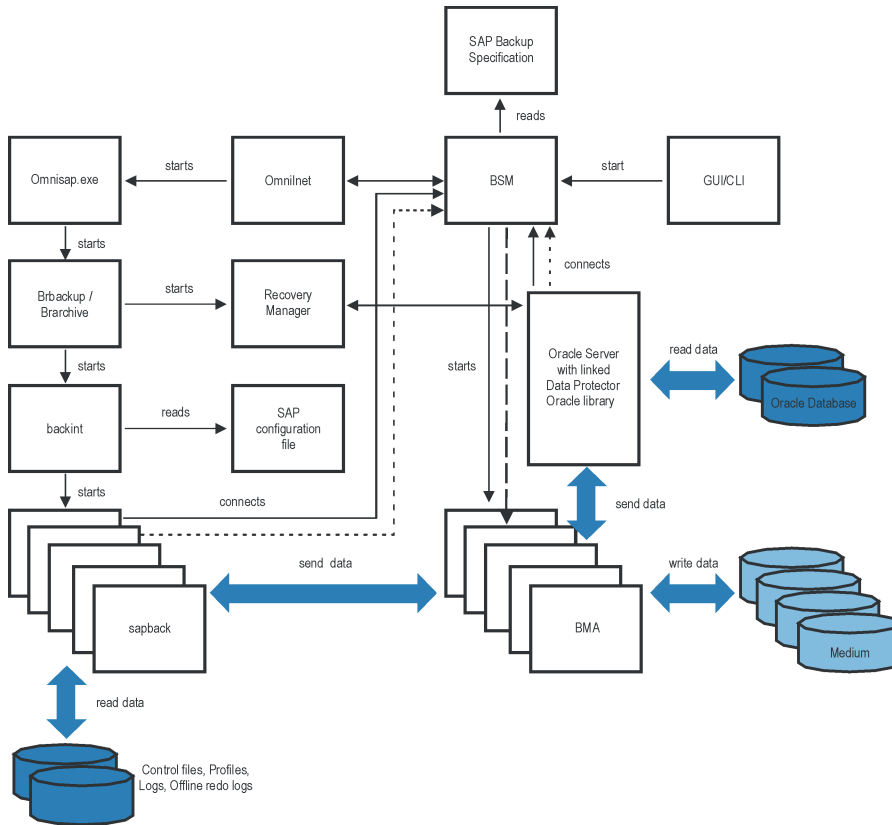


Figure 26 SAP R/3 architecture: RMAN mode

Backup flow

1. If the backup session is started:
 - **Using the Data Protector interface (or the scheduler):** BSM is started, which reads the appropriate Data Protector backup specification, checks if the devices are available, and starts `omnisap.exe` on the SAP R/3 client. The `omnisap.exe` agent exports the appropriate environment variables and starts `BRBACKUP` or `BRARCHIVE`.
 - **Using SAP BRTOOLS interface:** `BRBACKUP` or `BRARCHIVE` are started directly.

2. BRBACKUP does the following:
 - Changes the state of the Oracle Target Database (opened or closed), according to the backup type (online or offline).
 - Switches the Oracle Target Database to the `ARCHIVELOG` mode.
The archived redo log files are written to the archiving directory by Oracle and are backed up later using `BRARCHIVE`.
 - Creates the BRBACKUP log during the backup session, which contains information about backed up files and the backup ID. This information is needed to determine the location of the database files and archived redo log files during restore.
 - Sets the tablespace mode (`BEGIN / END BACKUP`) in the case of online backup using `backint`. In this way, the SAP R/3 application puts a tablespace in the backup mode just before it is backed up and returns it to the normal mode immediately after the backup completes.
3. If BRBACKUP is started:
 - a. BRBACKUP starts a `backint` command (`backint` mode) or `RMAN` (`RMAN` mode), which backs up Oracle data files and control files.
 - b. BRBACKUP starts a `backint` command (in the `backint` and `RMAN` mode), which backs up the SAP parameter file and the SAP R/3 history files that have been created during the backup of Oracle data files and control files.
- If `BRARCHIVE` is started (in the `backint` or `RMAN` mode), `BRARCHIVE` starts a `backint` command, which backs up archived redo log files. In addition, a copy of control files is created, which is also backed up.

 **NOTE:**

Backint divides the files specified for backup into subsets according to the selected balancing type and starts a `sapback` process for each subset (provided that the specified concurrency is large enough). The `sapback` processes read data from disks and send it to General Media Agents.

4. When all the General Media Agents finish with data transfer, the BSM waits for a timeout (`SmWaitForNewClient` omnirc global variable) and completes the backup session if no `backint` command is started within this time frame.

Restore flow

You can start a restore using the Data Protector user interface or SAP BRTOOLS user interface. However, only a standard filesystem restore can be performed using Data Protector.

1. When you select objects to be restored and start a restore using SAP BRTOOLS, the following happens (depending on which mode you use):
 - **Backint mode:** BRRESTORE checks if enough free disk space is available and starts a backint command to restore Oracle data files.
If the backups of files to be restored reside on different media, backint starts a separate `saprest` process for each medium, so that the files are restored in parallel (provided that the specified concurrency is large enough). The first `saprest` process starts a RSM, while the subsequent `saprest` processes connect to the same RSM. RSM checks if the restore devices are available and starts the data flow.
 - **RMAN mode:** BRRESTORE starts RMAN, which connects to Data Protector via Data Protector Database Library and Oracle Server processes and enables data transfer of Oracle data files.
2. When all the General Media Agents finish with data transfer, the RSM waits for a timeout (`SmWaitForNewClient` global variable) and completes the restore session if no backint command is started within this time frame.

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)
- concurrency number and balancing (for each backup specification), and number of channels for RMAN backup
- 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' | '-channels
#_of_RMAN_channels');
}
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 on how to configure SAP R/3 database for use with Data Protector, see [“Configuring SAP R/3 databases”](#) on page 173.

Example

This is an example of the file:

```

ORACLE_HOME='/app/oracle805/product';
ConnStr='EIBBKIBBEIBBFIBBGHBOHBB
QDBBOFBBCFBPFBBFCBBIFFBBGFBBDBBBFBBCFBBDFFBCCFB';
BR_directory='/usr/sap/ABA/SYS/exe/run';
SAPDATA_HOME='/sap';
ORA_NLS_CHARACTERSET='USASCII7';
OSUSER='orasad';
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), `/opt/omni/lbin` (HP-UX, Solaris, and Linux systems), or `/usr/omni/bin` (other UNIX 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, etc.

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, and Linux

```
/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
```

Other UNIX

```
/usr/omni/bin/util_cmd -putopt SAP ICE NLS_LANG \  
'US7ASCII' -sublist Environment
```

```
/usr/omni/bin/util_cmd -putopt SAP TOR 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, and Linux: `/opt/omni/lbin/util_cmd -getopt SAP ICE OB2OPTS \ -sublist Environment`
- On other UNIX: `/usr/omni/bin/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, and Linux: `/opt/omni/lbin/util_cmd -getconf SAP ICE`
- On other UNIX: `/usr/omni/bin/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, and Linux: `/opt/omni/lbin/util_cmd -putopt SAP ICE OB2OPTS "" -sublist Environment`
- On other UNIX: `/usr/omni/bin/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”](#) on page 168.
2. Check the connection to the Oracle database. See [“Checking the connection”](#) on page 169.

3. Enable the use of the authentication password file. See [“Authentication password file”](#) on page 170.
4. Optionally, set the archived logging mode to enable online backups. See [“Enabling archived logging”](#) on page 170.
5. If you plan to do backups and restores in the RMAN mode, link Oracle Server with the Data Protector MML. See [“Linking Oracle Server with the Data Protector MML”](#) on page 172.
6. Configure every SAP R/3 database you intend to back up from or restore to. See [“Configuring SAP R/3 databases”](#) on page 173.

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 (for example, Informix).
 - For supported versions, platforms, devices, and other information, see the *HP Data Protector product announcements, software notes, and references* or <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 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 SAP R/3 integration in various architectures, see the *HP Data Protector installation and licensing guide*.
 - For information on the Data Protector Cell Manager package configuration in the MC/SG cluster, see the online Help index: “MC/ServiceGuard integration”.

Every SAP R/3 application system you intend to back up from or restore to must have the Data Protector `SAP R/3 Integration` component installed.

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 only:** Restart the Data Protector `Inet` service under the Oracle operating system user account described in “[Configuring user accounts](#)” on page 168. For information on changing the Data Protector `Inet` account, see the online Help index: “changing Data Protector `Inet` 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` program on the currently active node. On all the other nodes, do it manually. Run:

```
ln -s /opt/omni/lbin/backint \  
/usr/sap/ORACLE_SID/sys/exe/run
```

Windows: During the configuration, Data Protector copies the Data Protector `backint` program from `Data_Protector_home\bin` to the directory that stores the SAP backup tools. 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: <code>dba</code> and <code>sapsys</code> • Windows systems: <code>ORA_DBA</code> and <code>ORA_SID_DBA</code> local groups <p>For example, user <code>oraSID</code>.</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 <code>root</code> (UNIX systems only)	Default operating system administrator's user account added to the <code>dba</code> user group.
Oracle database user account	<p>Database user account granted at least the following Oracle roles:</p> <ul style="list-style-type: none"> • <code>sysdba</code> • <code>sysoper</code> <p>For example, user <code>system</code>.</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`

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".

Checking the connection

To check the connection to the Oracle instance:

1. Log in to the SAP R/3 client 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.
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;
```

Linking Oracle Server with the Data Protector MML

To use the Data Protector SAP R/3 integration in the RMAN mode, the Oracle Server software needs to be linked with the Data Protector Oracle integration Media Management Library (**MML**) on every client on which an Oracle instance is running:

- On UNIX and Windows clients with Oracle8i, link Oracle Server with the Data Protector MML manually. For details, see [“Linking Oracle Server with the Data Protector MML”](#) on page 36.
- On UNIX and Windows clients with Oracle9i/10g, 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 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 SAP R/3 instance configuration file.

Choosing authentication mode

Data Protector SAP R/3 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 an SAP R/3 database with the new Oracle login information each time the respective 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
- If Oracle 10g Instant Client is installed on the backup system, the directory in which the SQL tools are installed.

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:

HP-UX, Solaris, and Linux: `/opt/omni/lbin`

Other UNIX: `/usr/omni/bin`

On Windows, Data Protector copies the `backint` program from `Data_Protector_home\bin` to the directory that stores the SAP backup tools.

 **IMPORTANT:**

If you plan to do offline backups using RMAN, do not configure the database with the Oracle database user `Internal` because the backup will fail. Configure the database with the user `System`.

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.
Click **OK**.

4. In **Client**, select the SAP R/3 system. In cluster environments, select the virtual server.

In **Application database**, type the Oracle instance name (ORACLE_SID).

UNIX only: In **Username**, type the Oracle OS user described in “Configuring user accounts” on page 168. In **Group name**, type dba.

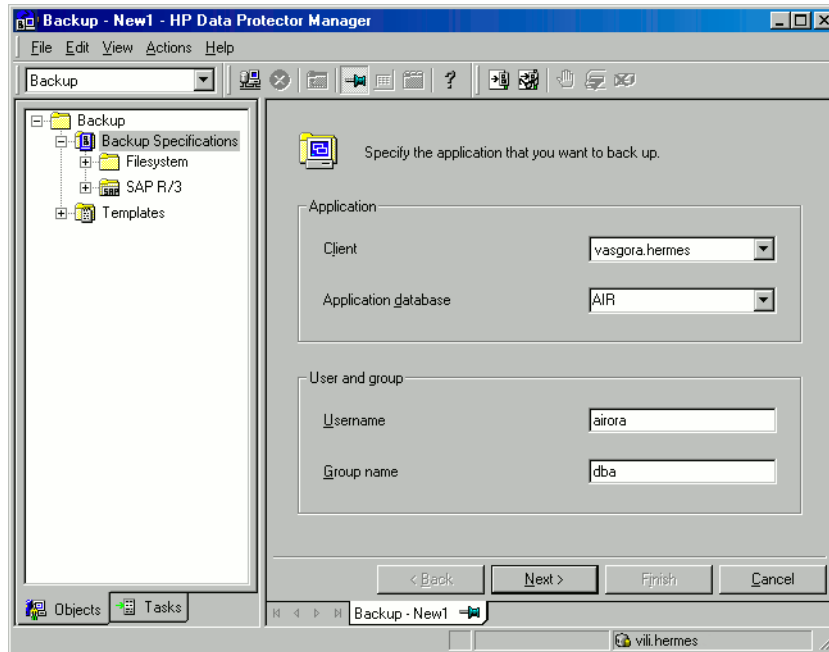


Figure 27 Specifying an SAP R/3 system and Oracle instance

Click **Next**.

5. 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:

- **Username** and **Password**: These two options determine the authentication mode that is used for accessing the Oracle database.
 - To select the database authentication mode, specify the user name and password of the Oracle database user account described in “[Configuring user accounts](#)” on page 168.
 - To select the operating system authentication mode, leave the text boxes empty.
- **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 28 Configuring an SAP R/3 database on a UNIX system (operating system authentication mode)

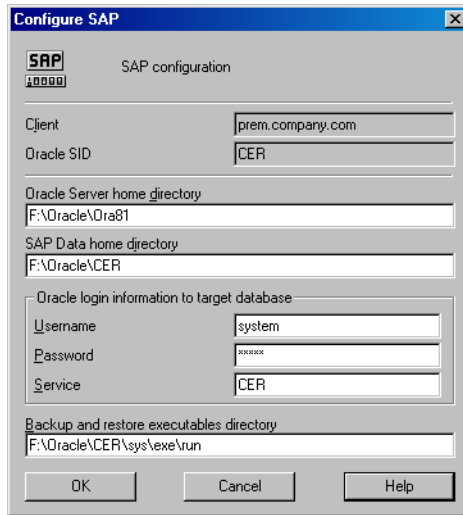


Figure 29 Configuring an SAP R/3 database on a Windows system (database authentication mode)

Click **OK**.

6. The SAP R/3 database is configured. Exit the GUI or proceed with creating the backup specification at [Step 6](#) on page 185.

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, and Linux systems: `/opt/omni/lbin`

Other UNIX systems: `/usr/omni/bin/`

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 operating system authentication mode, specify only the character */*.

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*.

SQL_PATH

Pathname of the directory where the SQLPlus tools are installed. Required if you installed Oracle 10g Instant Client on the backup client. If not specified, *ORACLE_HOME\bin* is used.

A successful configuration displays the message **RETVAL*0*.

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, and Linux: Run:

/opt/omni/lbin/omnigetmsg 12 error_number

Other UNIX: Run:

/usr/omni/bin/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 -OBJS0 ORACLE_SID
```

To get a list of database files of a tablespace, run:

```
util_sap.exe -OBJS1 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 30 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, and Linux: `/opt/omni/lbin`

Other UNIX: `/usr/omni/bin/`

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 `*RETVL*error_number` where `error_number` is different than zero, an error occurred. On how to get the error description, see [“Handling errors”](#) on page 178.

Backup

The integration provides online and offline database backups of the following types:

Table 9 Backup types

Full	Backs up all the selected SAP R/3 objects.
Incr	Oracle RMAN backup incremental level 1 (available only if you are using Oracle RMAN). Backs up changes made to the selected SAP R/3 data files since the last Full backup. Before you run an incremental backup, ensure that a Full backup exists.

For details on these backup types, see the Oracle SAP R/3 documentation.

To configure a backup, create a backup specification.

What is backed up depends on your selection in the backup specification. For details, see [Table 10](#) on page 181.

Table 10 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. See examples in [Table 11](#) on page 182.

Table 11 Two alternatives of specifying backup options

Backup type	<ol style="list-style-type: none">1. BRBACKUP options2. SAP parameter file settings
offline backup using backint	<ol style="list-style-type: none">1. <code>-t offline -d util_file</code>2. <code>backup_type = offline</code> <code>backup_dev_type = util_file</code>
online backup using backint (tablespaces are in the backup mode during the whole backup session)	<ol style="list-style-type: none">1. <code>-t online -d util_file</code>2. <code>backup_dev_type = util_file</code> <code>backup_type = online</code>
online backup using backint (tablespaces are in the backup mode only while being backed up)	<ol style="list-style-type: none">1. <code>-t online -d util_file_online</code>2. <code>backup_dev_type = util_file_online</code> <code>backup_type = online</code>
full backup	<ol style="list-style-type: none">1. <code>-m full</code>2. <code>backup_mode = full</code>
backup using RMAN	<ol style="list-style-type: none">1. <code>-d rman_util</code>2. <code>backup_dev_type = rman_util</code> <code>rman_channels = number_of_channels</code> <code>rman_parms =</code> <code>"ENV=(OB2BARTYPE=SAP,OB2APPNAME=DB_Name,</code> <code>OB2BARLIST=Backup_Specification_Name)"</code>

Backup type	<ol style="list-style-type: none"> 1. BRBACKUP options 2. SAP parameter file settings
	For more information, see “Backing up using Oracle Recovery Manager” on page 195.

 **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.
- Backup sessions that back up the same Oracle instance cannot run simultaneously.
- Generally, restore takes longer than backup. The restore is significantly prolonged if files are backed up with many streams. Note that if you start a backup in the RMAN mode with the Oracle RMAN script option `FILESERSET` set to 1, RMAN creates a separate backup stream (object) for each database file.

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 and click **OK**.

Table 12 Backup templates

Blank SAP Backup	No predefined options.
Brarchive_CopyDeleteSave	Creates a second copy of offline redo logs, saves them, deletes them after the backup, and archives the newly-created redo logs.
Brarchive_Save	Backs up offline redo logs.
Brarchive_SaveDelete	Backs up offline redo logs and deletes them after the backup.
Brarchive_SecondCopyDelete	Creates a second copy of offline redo logs that have already been archived and deletes them after the backup.
Brbackup_Offline	Backs up the shut-down database using <code>backint</code> .
Brbackup_Online	Backs up the active database. The <code>util_file</code> device type is used for backup. Tablespaces are in the backup mode (locked) during the whole backup session. You can back up the entire database or only individual tablespaces or datafiles.
Brbackup_Util_File_Online	Backs up the active database. Tablespaces are in the backup mode only while being backed up. Consequently, the increase in archived log files is smaller compared to backup with the <code>util_file</code> device type. However, if the database consists of a large number of small files, this backup can take longer.
Brbackup_RMAN_Offline	Backs up the shut-down database using Oracle RMAN.
Brbackup_RMAN_Online	Backs up the active database using Oracle RMAN. Tablespaces are in the backup mode during the whole backup session.

4. In **Client**, select the SAP R/3 system on which the backup should be started. In cluster environments , select the virtual server.

In **Application database**, select the Oracle instance (ORACLE_SID) to be backed up.

UNIX only: In **Username**, type the Oracle OS user name described in “Configuring user accounts” on page 168. In **Group name**, type dba.

Click **Next**.

5. 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” on page 173.
6. Select SAP R/3 objects to be backed up. You can select individual tablespaces, data files, or archived logs.

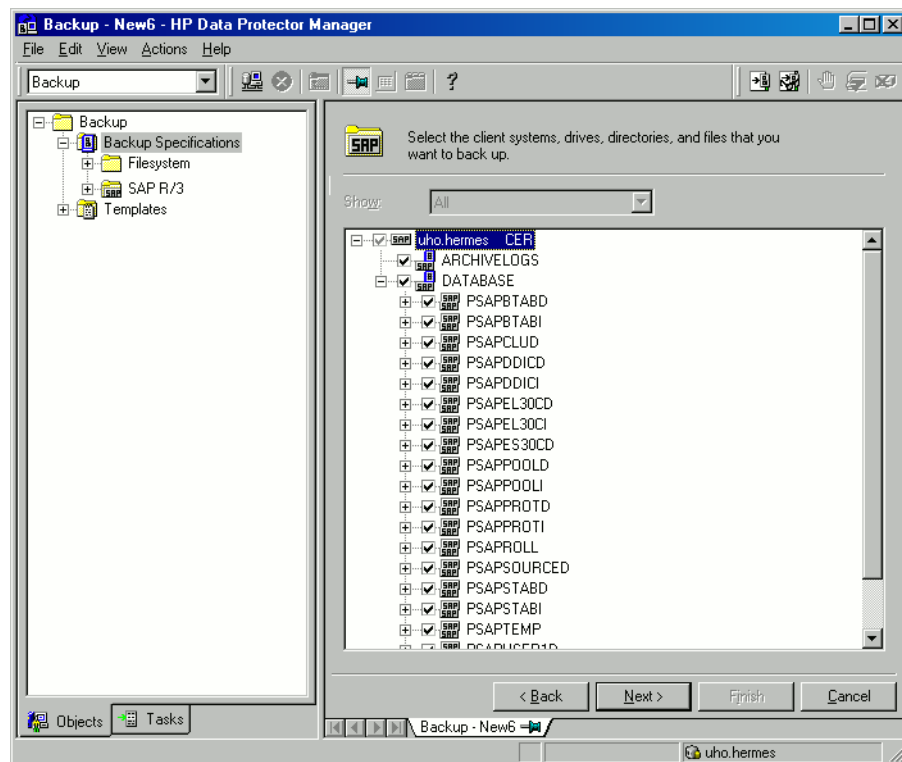


Figure 31 Selecting backup objects

Click **Next**.

7. 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**.

8. Set backup options. For information on the application specific options, see [Table 13](#) on page 187.

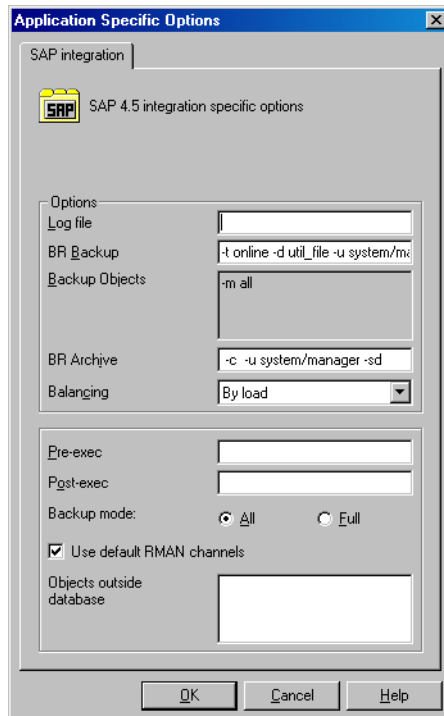


Figure 32 Application specific options

Click **Next**.

9. Optionally, schedule the backup. See [“Scheduling backup specifications”](#) on page 189.

Click **Next**.

10. 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”](#) on page 190.

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. 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.
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” on page 160.
Balancing: By Time	Groups files into subsets that are backed up in approximately equal periods of time. The duration depends on the file types,

Option	Description
	<p>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.</p> <p>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.</p>
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 ” on page 196.
Balancing: None	<p>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:</p> <pre>select * from dba_data_files</pre>
Pre-exec, Post-exec	<p>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:</p> <p>Windows: <code>Data_Protector_home\bin</code> HP-UX, Solaris, and Linux: <code>/opt/omni/bin</code> Other UNIX: <code>/usr/omni/bin</code></p>
Backup mode	<p>Specifies the RMAN backup type to be used. Available only if the whole database is selected for backup.</p> <p>If <code>All</code> is specified, RMAN backs up the whole database.</p> <p>If <code>Full</code> is specified, RMAN performs a Full backup (level 0), thus enabling RMAN incremental backups.</p>
Use default RMAN channels	Specifies the concurrency value for your backup. Applicable only if RMAN is used for backup. This option overrides the settings in the SAP parameter file.
Objects outside database	<p>Specifies non-database files of the Oracle SAP R/3 environment to be saved.</p> <p>Save these files in a separate backup session.</p>



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 **Full** backups 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**. See [Figure 33](#) on page 190.
Click **OK**.
3. Repeat [Step 1](#) on page 189 and [Step 2](#) on page 189 to schedule backups at 13:00 and 18:00.
4. Click **Apply** to save the changes.

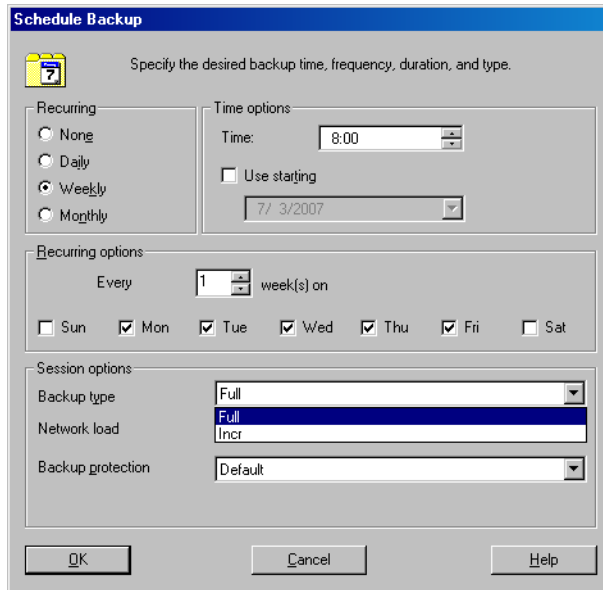


Figure 33 Scheduling backups

Previewing backup sessions

Preview the backup session to test it. You can use the Data Protector GUI or CLI.

Using the Data Protector GUI

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 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

From the directory:

Windows: `Data_Protector_home\bin`

HP-UX, Solaris, and Linux: `/opt/omni/bin/`

Other UNIX: /usr/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:

- Communication between the Oracle instance and Data Protector (only if RMAN is used)
- 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 start 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

From the directory:

Windows: `Data_Protector_home\bin`

HP-UX, Solaris, and Linux: `/opt/omni/bin/`

Other UNIX: `/usr/omni/bin/`

run:

```
omnib -sap_list backup_specification_name [-barmode  
SAP_mode] [List_options]
```

where `SAP_mode` is one of the following:

```
full|incr
```

For details, see the `omnib` man page or the *HP Data Protector command line interface reference*.

Example

To start a full backup using the SAP R/3 backup specification `RONA`, run:

```
omnib -sap_list RONA -barmode full
```

Using the SAP BRTOOLS

1. Log in to the SAP R/3 system as the Oracle OS user.

2. Export/set the following environmental 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
```

Specifies Data Protector devices to 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.

```
[OB2BARHOSTNAME=application_system_name]
```

Optional if you want to specify virtual server name in cluster environments.

3. **Oracle9i/10g only:** If you plan to do backups in the RMAN mode, ensure that the SBT_LIBRARY parameter in the `initSAP_instance.sap` file points to the correct platform-specific Data Protector MML. For details on the Data Protector MML location, see [Step 3](#) on page 85.
4. Run the BRBACKUP command.

```
brbackup -t [online | offline]_[split | mirror] -d
util_file -m all -c -u user/password
```

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` environmental 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` environmental variable 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 **SAP R/3**. Click a backup specification for your SAP R/3 instance.
 3. In the Source page, right click the SAP R/3 instance and click **Set Environmental Variables**.
 4. In the Advanced dialog box, set OB2_MIRROR_COMP to 1. See [Figure 34](#) on page 194.
- Click **OK**.

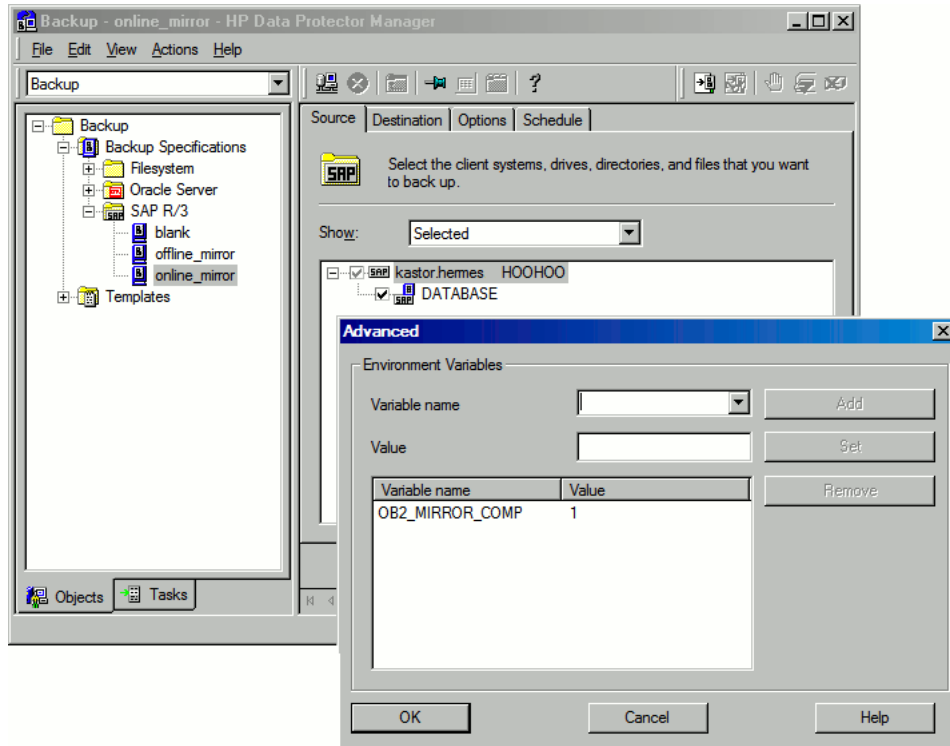


Figure 34 Setting environmental variables

Using the Data Protector CLI

From the directory:

Windows: `Data_Protector_home\bin`

HP-UX, Solaris, and Linux: `/opt/omni/bin/`

Other UNIX: /usr/omni/bin/

run:

```
util_cmd -putopt SAP instance_name OB2_MIRROR_COMP 1 -sublist  
Environment
```

Backing up using Oracle Recovery Manager

If RMAN is used directly, consider the following:

- RMAN stores information about backups in the recovery catalog. For security reasons, keep the catalog in a separate database. This requires more administrative work.
- In a disaster situation (such as the loss of a production database and recovery catalog), the restore and recovery of data is complicated. It may be impossible without the help of Oracle Support. If the Recovery Manager does not have administrative data stored in the recovery catalog, it cannot recover the database only by using the backups that have been made.
- **Oracle9i/10g only:** For each RMAN channel, set the `SBT_LIBRARY` parameter to point to the correct platform-specific Data Protector MML. For details on the Data Protector MML location, see [Step 3](#) on page 85.

If RMAN is used through the BRBACKUP utility, consider the following:

- The recovery catalog is not used. Information about backups is saved in the control file and SAP R/3 log files. After each backup, the control file and SAP R/3 log files are saved. When data is restored, the control file is copied back first, followed by data files. In case of a disaster, restore SAP R/3 log files before you restore any data files.
- Other important files will still be automatically backed up using the backint program.
- All previous SAP R/3 backup strategies can still be used with RMAN. However, RMAN cannot be used for offline redo log backups with BRARCHIVE, or for standby database backups.
- **Oracle9i/10g only:** Ensure that the `SBT_LIBRARY` parameter in the `initSAP_instance.sap` file points to the correct platform-specific Data Protector MML. For details on the Data Protector MML location, see [Step 3](#) on page 85.

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 in any of the following ways:

- Use the Data Protector GUI. See [“Restoring using the Data Protector GUI”](#) on page 197.
- Use the Data Protector CLI. See [“Restoring using the Data Protector CLI”](#) on page 200.
- Use the SAP restore commands. See [“Restoring using the SAP commands”](#) on page 201.

After the restore, you can recover the database to a specific point in time using the SAP BRTOOLS interface.

Considerations

- Backups created by Oracle RMAN can only be restored using the SAP restore utilities.
- 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 you are restoring a sparse file, you can improve the performance by setting the sparse option. See “[Sparse files](#)” on page 203.
- 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](#)” on page 202.
- Restore preview is not supported.

Restoring using the Data Protector GUI

1. In the Context List, click **Restore**.
2. In the Scoping Pane, expand **SAP R/3**, expand the client 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**.

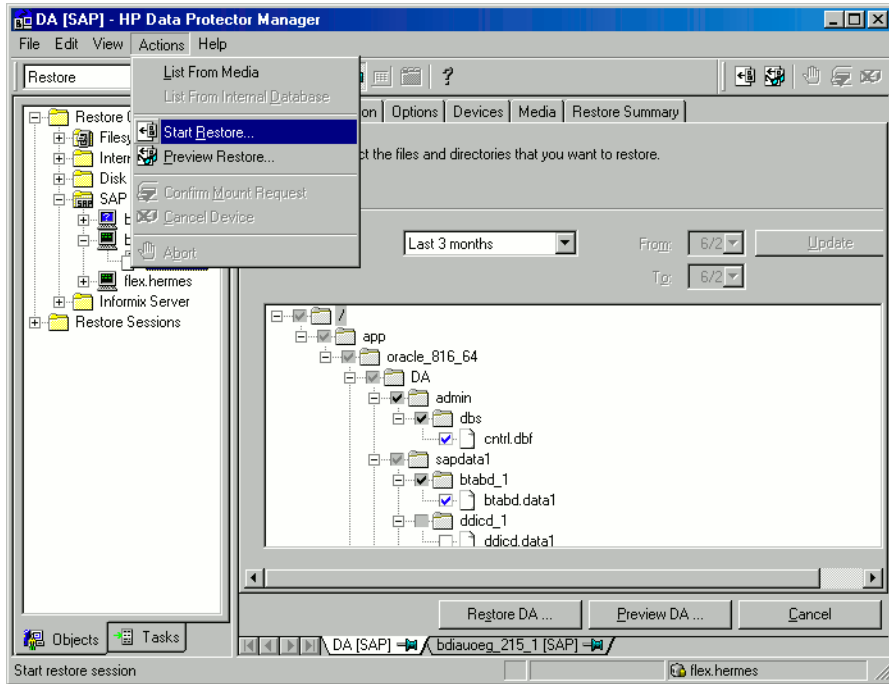


Figure 35 Selecting objects for restore

4. In the **Destination** tab, select the client to restore to (**Target client**). See [Figure 36](#) on page 199.

For details on options, press **F1**.

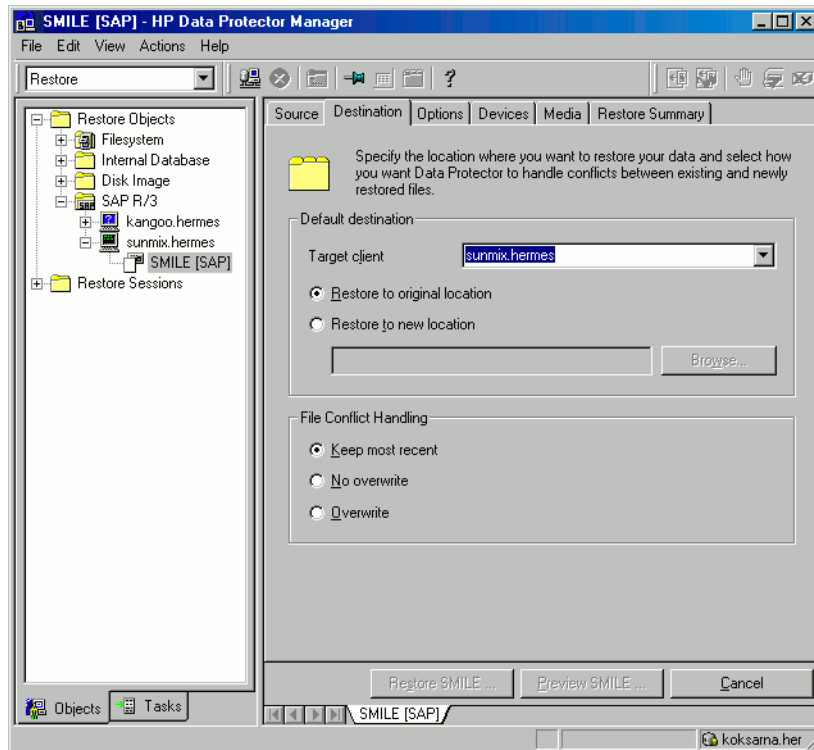


Figure 36 Selecting the target client

5. In the **Options** page, set the restore options. For information, press **F1**.
6. In the **Devices** page, select devices to use for the restore.
7. Click **Restore**.
8. In the **Start Restore Session** dialog box, click **Next**.
9. Specify **Report level** and **Network load**.
10. Click **Finish** to start the restore.

The message `Session completed successfully` is displayed at the end of a successful session.

Restoring using the Data Protector CLI

From the directory:

Windows: `Data_Protector_home\bin`

HP-UX, Solaris, and Linux: `/opt/omni/bin/`

Other UNIX: `/usr/omni/bin/`

run:

```
omnir -sap Client:Set -session SessionID -tree FileName
```

where `FileName` is the pathname of the SAP R/3 file to be restored.

Windows only: Specify the pathname in the UNIX format (using slashes to separate the drive letter, directories, and the filename. The drive letter must be preceded by a slash).

Example (Windows)

To restore the SAP R/3 file `btabd_1.dat` to the original location `C:\oracle\ABA\sapdata1\btabd_1` on the Windows system `computer1.company.com` from the backup session `2006/01/23-1`, run:

```
omnir -sap computer1.company.com:ABA.0 -session 2006/01/23-1  
-tree /C:/oracle/ABA/sapdata1/btabd_1/btabd_1.dat
```

Example (UNIX)

To restore the SAP R/3 file `btabd_1.dat` to the original location `/app/oracle/ABA/sapdata1/btabd_1` on the UNIX system `computer2.company.com` from the backup session `2006/01/23-1`, run:

```
omnir -sap computer2.company.com:ABA.0 -session 2006/01/23-1  
-tree /app/oracle/ABA/sapdata1/btabd_1/btabd_1.dat
```

TIP:

To get a list of backed up SAP R/3 objects, run:

```
omnidb -sap
```

To get details on a specific object, including the SessionID, run:

```
omnidb -sap object_name
```

Restoring using the SAP commands

You can start a restore of the SAP R/3 database using the SAP BRRESTORE command. The command uses the Data Protector backint interface to restore files backed up with Data Protector.

1. Log in to the SAP R/3 client as the Oracle OS user.
2. Ensure that you have enough disk space. BRRESTORE needs additional disk space to restore the control file and archived redo log files.
3. Specify the Oracle database to be restored by setting the OB2APPNAME environment variable:

UNIX: `export OB2APPNAME=ORACLE_SID`

Windows: `set OB2APPNAME=ORACLE_SID`

 **NOTE:**

If you have more than one database corresponding to the same ORACLE_SID name, also specify the client:

UNIX: `export OB2HOSTNAME=client_name`

Windows: `set OB2HOSTNAME=client_name`

4. If you plan do restores in the RMAN mode, ensure that the SBT_LIBRARY parameter in the `initSAP_instance.sap` file points to the correct platform-specific Data Protector MML. For details on the Data Protector MML location, see [Step 3](#) on page 85.
5. Run the SAP restore command.

Restoring using another device

You can restore using a device other than that used for backup.

Using the Data Protector GUI

On how to specify another device for restore using the Data Protector GUI, see the online Help index: “restore, selecting devices for”.

Using the Data Protector CLI or SAP R/3 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 [“Restore sessions fail due to invalid characters in filenames”](#) on page ?.

If you are restoring files using the Data Protector CLI and the names of backed up objects contain characters that cannot be displayed using the current language group (Windows) or code page (UNIX):

1. Set the environment variable `OB2_CLI_UTF8` to 1.
2. **Windows only:** Set the encoding used by the terminal to UTF-8.

Otherwise, the output of some commands is not displayed correctly (for example, backup objects returned by `omnidb`) and cannot be used as input for other commands (for example `omnir`).

Sparse files

You can improve performance of a sparse file restore by setting the `sparse` option. Set the option in any of the following ways:

- Using the Data Protector GUI: Select the **Restore sparse files** option in the **Options** page.
- Using the Data Protector CLI: Add the `-sparse` option when running the `omnir` command.
- Using the SAP commands: Before running the `BRRESTORE` command, set the Data Protector `OB2SPARSE` variable:

Windows: `set OB2SPARSE=sparse`

UNIX: `export OB2SPARSE=sparse`

Disaster recovery

For general information, see the *HP Data Protector disaster recovery guide*.

Restoring the control file

The control file contains all the information about the database structure. If the control file is lost, restore the control file before you restore any other part of the database:

1. Restore the control file using the standard Data Protector restore procedure.

The control files (`ctrlORACLE_SID.dbf`) are restored to the directory defined by the `SAPBACKUP` variable. If the variable is not set, the control files are restored to:

HP-UX, Solaris, and Linux: `/var/opt/omni/tmp`

Other UNIX: `/usr/opt/omni/tmp`

Windows: `Oracle_home\tmp`

2. Run:

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

where `TMP_FILENAME` is the folder to which the control file was restored.

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*.

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 the support matrices at <http://www.hp.com/support/manuals> for an up-to-date list of supported versions, platforms, and other information.

General troubleshooting

Data Protector reports “12:8422” error when using Data Protector Oracle integration after an upgrade of Oracle8i to Oracle9i

Problem

After Oracle8i is upgraded to Oracle9i, the following error is returned during the configuration of Oracle instance or during the backup:

```
*RETVAL*8422
```

Action

Rename the Oracle8i `svrmgr1` binary to something else so that Data Protector will not find it. The Oracle upgrade process from Oracle8i to Oracle9i does not remove the Oracle8i `svrmgr1` binary, rather it changes its permissions. Once the `svrmgr1` binary is renamed, Data Protector will use Oracle9i `sqlplus`, as it should, to complete the operations correctly.

Configuration fails due to a database operation failure

Problem

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”](#) on page 168.

Troubleshooting on Windows systems

Prerequisites concerning the Oracle side of the integration

The following steps should be performed to verify that Oracle is installed as required for the integration to work. These steps do not include verifying Data Protector components.

1. Verify that you can access the Oracle Target Database and that it is opened, as follows:

Set *ORACLE_HOME* and *ORACLE_SID* variables.

Start the Server Manager (Oracle8/8i) or SQL Plus (Oracle9i) from the *ORACLE_HOME* directory:

`bin\svrmgrl (Oracle8/8i) or`

`bin\sqlplus (Oracle9i)`

At the SVRMGR (Oracle8/8i) or SQL (Oracle9i) prompt, type:

```
connect user/passwd@service
```

```
select * from dba_tablespaces;
```

```
exit
```

If this fails, open the Oracle Target Database.

2. Verify that the TNS listener is correctly configured for the Oracle Target Database. This is required for properly establishing network connections:

Start the listener from the `ORACLE_HOME` directory:

```
bin\lsnrctl80 status service (for Oracle8) or
```

```
bin\lsnrctl status service (for Oracle8i/9i)
```

```
quit
```

If it fails, start up the TNS listener process and refer to the Oracle documentation for instructions on how to create a TNS configuration file (`LISTENER.ORA`).

The listener process can be started from the Windows desktop. In the **Control Panel**, go to **Administrative Tools, Services**.

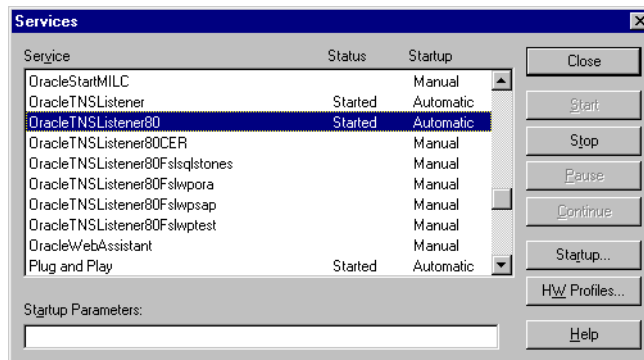


Figure 37 Checking the status of the Oracle listener

- The status of the respective listener service in the **Services** window should be **Started**, otherwise you must start it manually.
- Start the Server Manager (Oracle8/8i) or SQL Plus (Oracle9i) from the `ORACLE_HOME` directory:

```
bin\svrmgrl (Oracle8/8i) or
```

```
bin\sqlplus (Oracle9i)
```

At the `SVRMGR` (Oracle8/8i) or `SQL` (Oracle9i) prompt, type:

```
connect Target_Database_Login
```

```
exit
```

If it fails, refer to the Oracle documentation for instructions on how to create a TNS configuration file (`TNSNAMES.ORA`).

3. If you are running backups in RMAN mode, verify that the Oracle Target Database is configured to allow remote connections with system privileges:

Set `ORACLE_HOME` as described in [Step 1](#) on page 218 and start the Server Manager from the `ORACLE_HOME` directory:

```
bin\svrmgrl
```

At the `wSVRMGR` prompt, type

```
connect Target_Database_Login as SYSDBA;
```

```
exit
```

Repeat the procedure using `SYSOPER` instead of `SYSDBA`. Set the `ORACLE_HOME` directory

If you are using the recovery catalog:

```
bin\rman target Target_Database_Login rcvcat  
Recovery_Catalog_Login
```

If you are not using the recovery catalog:

```
bin\rman target Target_Database_Login nocatalog
```

If this fails, refer to the Oracle documentation for instructions on how to set up the password file and any relevant parameters in the `initORACLE_SID.ora` file.

Prerequisites on the SAP side of the integration

The following verification steps must be performed in order to verify that SAP is installed as required for the integration to work. These steps do not include Data Protector components.

1. Verify backup directly to disk as follows:

```
brbackup -d disk -u user/password
```

If this fails, check the error messages and resolve possible problems before you continue.

2. Verify restore directly to disk as follows:

```
brrestore -d disk -u user/password
```

If this fails, check the error messages and resolve possible problems before you continue.

3. If you are running backups in RMAN mode, verify backup and restore directly to disk using Recovery Manager channel type disk as follows:

- a.** You must define the parameter `init` in the initialization file `initORACLE_SID.ora`. Run the following commands:

```
brrestore -d pipe -u user/password -t online -m all
brrestore -d disk -u user/password
```

- b.** If this fails, refer to the SAP Online Help to learn how to execute backup and restore directly to disk using the SAP backup utility.

Check the error message and resolve these problems before you continue.

4. Verify that the SAP backup tools correctly start backint (which is provided by Data Protector):

Move the original `backint` and create a test script `namedbackint.bat` in the directory where the SAP backup utility resides, with the following entries:

```
echo "Test backint called as follows:"
echo "%0%1%2%3%4%5%6%7%8%9"
exit
```

Then start the following commands:

```
brbackup -t offline -d util_file -u user/password -c
```

If you receive `backint` arguments, this means that SAP is properly configured for backup using `backint`; otherwise you have to reconfigure SAP.

See “[Configuring SAP R/3 databases](#)” on page 173.

Configuration problems



IMPORTANT:

The procedure described in the previous sections must be performed before you start checking the Data Protector configuration.

1. Verify that the Data Protector software has been installed properly.

Refer to the *HP Data Protector installation and licensing guide* for details.

2. Perform a filesystem backup of the SAP Database Server.

Perform a filesystem backup of the SAP Database Server system so that you can eliminate any potential communication problems between the SAP Database Server and the Data Protector Cell Manager system.

Do not start troubleshooting an online database backup unless you have successfully completed a filesystem backup of the SAP Database Server system.

See the online Help index “standard backup procedure” for details about how to do a filesystem backup.

3. If the SAP backup utilities are installed in a shared directory, then the `inet` startup parameter must be specified as described in [Step 4](#) on page 211, or the Windows permissions must be set correctly.

Run the following command (if you use the default directory):

```
dir \\client_name\sapmnt\ORACLE_SID\SYS\exe\run\brbackup
```

or

```
dir \\client_name\SAPEXE\brbackup
```

If this fails, set the `inet` startup parameters, or set the correct permissions to access a Windows network directory.

4. If you use the command line to start the Data Protector commands, verify the inet startup parameters.

Check the `Data Protector Inet` service startup parameters on the SAP Database Server system. Proceed as follows:

- a.** In the **Control Panel**, go to **Administrative Tools, Services**.
- b.** Select **Data Protector Inet**.

In the **Services** window, select **Data Protector Inet, Startup**.

The service must run under a specified user account. Make sure that the same user is also added to the Data Protector `admin` user group.



Figure 38 Checking the Inet start-up parameters

5. Examine the environment variables.

If you need to export some variables before starting the Oracle Server Manager, TNS listener, or other Oracle utility, these variables must be defined in the `Environment` section of the Data Protector SAP configuration file on the Cell Manager. See “[Data Protector SAP R/3 configuration file](#)” on page 160.

6. Examine system errors.

System errors are reported in the `Data_Protector_home\log\debug.log` file on the SAP Server.

Configuration fails due to a script failure

Problem

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

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”](#) on page 167.

Backup problems

At this stage, you should have performed all the verification steps described in the previous sections. If backup still fails, proceed as follows:

1. Check your SAP Server configuration:

To check the configuration, start the following command on the SAP Server system:

```
Data_Protector_home\bin\util_sap.exe -CHKCONF ORACLE_SID
```

The *RETVAL*0 indicates successful configuration.

2. Verify Data Protector internal data transfer using the `testbar2` utility.

Before you run the `testbar2` utility, verify that the Cell Manager name is correctly defined on the SAP Database Server. Check the `Data_Protector_home\Config\client\cell_server` file, which contains the name of the Cell Manager system. Then run the following command:

```
Data_Protector_home\bin\testbar2 -type:SAP  
-appname:ORACLE_SID -bar:backup_specification_name  
-perform:backup
```

Examine the errors reported by the `testbar2` utility by clicking the **Details** button in the Data Protector **Monitor** context.

If the messages indicate problems concerning the Data Protector side of the integration, create an SAP backup specification to back up to a nul or file device. If the backup succeeds, the problem may be related to the backup devices. Refer to the *HP Data Protector troubleshooting guide* for instructions on troubleshooting devices. If the test fails again, call support.

3. Verify the backup using backint

```
export OB2BARLIST=barlist_name
```

```
export OB2APPNAME=ORACLE_SID
```

```
Data_Protector_home\bin\backint.exe -f backup -t file -u  
ORACLE_SID -i input_file
```

where *input_file* is a file with a list of full pathnames for backup.

Backint anticipates a list of files in the following
format:*pathName_1pathName_2pathName_3*

Backup fails at the beginning with the message “Internal heap ERROR 17112”

Problem

When using SAP 4.6D kernel on HP-UX 11.11, backup fails immediately after it was started due to a BRBACKUP core dump. A line similar to the following can be found at the beginning of the message:

```
Internal heap ERROR 17112 addr=0x800003ffff7f3660
```

Action

1. Login to the SAP server as the user who is owner of the backup specification.

2. Run the command

```
env | grep NLS_LANG
```

The output is similar to the following:

```
NLS_LANG=AMERICAN_AMERICA.US7ASCII
```

3. Add the NLS_LANG variable to the backup specification. For more details, see [“Setting, retrieving, listing and deleting Data Protector SAPR/3 configuration file parameters using the CLI”](#) on page 163.

4. Restart the backup.

Backup fails with “Connect to database instance failed”

Problem

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
```

Action

Before you start a backup, ensure that the database instance is in the open or shutdown mode.

Restore problems

At this stage, you should have performed all the verification steps described in the previous sections. After this, proceed as follows:

1. Verify that a backup object exists on the backup media and in the IDB:

This can be done by executing the command

```
Data_Protector_home\bin\omnidb -SAP "object_name" -session  
"Session_ID" -media
```

on the SAP Database Server system.

The output of the command lists detailed information about the specified backup object, session IDs of the backup sessions containing this object, and a list of the media used.

For detailed syntax of the `omnidb` command, run:

```
Data_Protector_home\bin\omnidb -help
```

You can also do this using the SAP tools:

Use `backint`, so that SAP tools also use this command to query:

```
Data_Protector_home\bin\backint.exe -f inquiry -u  
ORACLE_SID -i input_file
```

where the specified `input_file` is queried.

If this fails, check if the backup session was performed successfully and if the query was started under the appropriate user account.

`Backint` anticipates a list of files of the following format:

```
backup_ID_1 pathName_1 [targetDirectory_1]  
backup_ID_2 pathName_2 [targetDirectory_2]  
backup_ID_3 pathName_3 [targetDirectory_3]
```

To retrieve the `backup_ID` numbers, enter the following command:

```
echo #NULL #NULL | backint -f inquiry -u ORACLE_SID
```

or, alternatively, you can just specify `#NULL` as `backup_ID_1` in the `input_file`. In this case, the latest backup session for the file is used for the restore.

2. Verify the restore using the Data Protector User Interface

This test is possible if the objects have been backed up by `backint`.

If this fails, check if the backup session was performed successfully and if the query was started under the appropriate user account.

3. Simulate a Restore Session

Once you know the information about the object to be restored, you can simulate a restore using the Data Protector `testbar2` utility.

Before you run `testbar2`, verify that the Cell Manager name is correctly defined on the SAP Database Server.

Check the `Data_Protector_home\Config\client\cell_server`, which contains the name of the Cell Manager system.

Then, test the Data Protector internal data transfer using the `testbar2` utility:

```
Data_Protector_home\bin\testbar2 -type:SAP
-appname:ORACLE_SID
-perform:restore
-object:object_name
-version:object_version
-bar:backup_specification_name
```

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.

4. Verify the restore using backint

Run the following command:

```
Data_Protector_home\bin\backint.exe -f restore -u
ORACLE_SID -i input_file
```

where the contents of the `input_file` will be restored.

If this fails, check if the session was performed successfully and if the restore was started under the appropriate user account.

Backint anticipates a list of files in the following format:
`backup_ID_1 pathName_1 [targetDirectory_1] backup_ID_2 pathName_2`
`[targetDirectory_2] backup_ID_3 pathName_3`
`[targetDirectory_3]`

To retrieve the `backup_ID` numbers, enter the following command:

```
echo "#NULL #NULL" | backint -f inquiry -u ORACLE_SID
```

Restore sessions fail due to invalid characters in filenames

Problem

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, restore fails 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.

Troubleshooting on UNIX systems

Using Oracle after removing the Data Protector Oracle integration

This section is relevant only if Oracle8i RMAN has been used to back up the SAP R/3 datafiles and you have uninstalled the Data Protector Oracle integration on an Oracle server.

After uninstalling the Data Protector Oracle integration on an Oracle server, the Oracle server software is still linked to the Data Protector Database Library. You have to rebuild the Oracle binary to remove this link. If this is not done, the Oracle server cannot be started after the integration has been removed.

See [“Using Oracle after removing the Data Protector Oracle integration”](#) on page 134 for more information on how to make the Oracle server functional again.

Prerequisites concerning the Oracle side of the integration

The following steps should be performed to verify that Oracle is installed as required for the integration to work. These steps do not include verifying Data Protector components.

1. Verify that you can access the Oracle Target Database and that it is opened, as follows:

Export *ORACLE_HOME* and *ORACLE_SID* as follows:

- if you are using an SH - like shell enter the following commands:

```
ORACLE_HOME="ORACLE_HOME"  
export ORACLE_HOME  
ORACLE_SID = "ORACLE_SID"  
export ORACLE_SID
```

- if you are using a CSH - like shell enter the following commands:

```
setenv ORACLE_HOME "ORACLE_HOME"  
setenv ORACLE_SID "ORACLE_SID"
```

Start the Server Manager (Oracle8/9i) or SQL Plus (Oracle9i) from the *ORACLE_HOME* directory:

bin\svrmgrl (Oracle8/8i) or

bin\sqlplus (Oracle9i)

At the SVRMGR (Oracle8/8i) or SQL (Oracle9i) prompt, type:

```
connect user/passwd@service  
select * from dba_tablespaces;  
exit
```

If it fails, open the Oracle Target Database.

2. Verify that the TNS listener is correctly configured for the Oracle Target Database. This is required for properly establishing network connections:

Export `ORACLE_HOME` as described in [Step 1](#) on page 218 and start the listener from the `ORACLE_HOME` directory:

```
bin/lsnrctl start service  
exit
```

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

Export `ORACLE_HOME` as described in [Step 1](#) on page 218 and start the Server Manager (Oracle8/8i) or SQL Plus (Oracle9i) from the `ORACLE_HOME` directory:

```
bin/svrmgrl (Oracle8/8i)  
bin/svrmgrl (Oracle9i)
```

At the `SVRMGR` (Oracle8/8i) or `SQL` (Oracle9i) prompt, type:

```
connect Target_Database_Login  
exit
```

If it fails, refer to the Oracle documentation for instructions on how to create a TNS configuration file (`TNSNAMES.ORA`).

3. If you run backups in RMAN mode, verify that the Oracle Target Database is configured to allow remote connections with system privileges:

Export `ORACLE_HOME` as described in [Step 1](#) on page 218 and start the Server Manager (Oracle8/8i) or SQL Plus (Oracle9i) from the `ORACLE_HOME` directory:

```
bin/svrmgrl (Oracle8/8i)
```

```
bin/svrmgrl (Oracle9i)
```

At the SVRMGR (Oracle8/8i) or SQL (Oracle9i) prompt, type:

```
connect Target_Database_Login as SYSDBA;
```

```
exit
```

Repeat the procedure using `SYSOPER` instead of `SYSDBA`. Set the `ORACLE_HOME` directory

If you use the Recovery Catalog:

```
bin/rman target Target_Database_Login rcvcat  
Recovery_Catalog_Login
```

If you do not use the Recovery Catalog:

```
bin/rman target Target_Database_Login nocatalog
```

If this fails, refer to the Oracle documentation for instructions on how to set up the password file and any relevant parameters in the `initORACLE_SID.ora` file.

4. If you run backups in the RMAN mode, verify backup and restore directly to disk using the Recovery Manager channel type disk.

If you use the Recovery Catalog:

Export `ORACLE_HOME` as described in [Step 1](#) on page 218 and start Recovery Manager:

```
bin/rman target Target_Database_Login rcvcat
Recovery_Catalog_Login cmd_file=rman_script
```

If you do not use the Recovery Catalog:

Export `ORACLE_HOME` as described in [Step 1](#) on page 218 and start Recovery Manager:

```
bin/rman target Target_Database_Login nocatalog
cmd_file=rman_script
```

An example of the `rman_script` is listed 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' release
channel 'dev0';
}
```

If one of the above procedures fails, refer to the Oracle documentation to learn how to execute backup and restore directly to disk using the Recovery Manager.

Prerequisites on the SAP side of the integration

The following verification steps must be performed in order to verify that SAP is installed as required for the integration to work. These steps do not include Data Protector components.

1. Verify backup directly to disk as follows:

```
brbackup -d disk -u user/password
```

If this fails, check the error messages and resolve possible problems before you continue.

2. Verify restore directly to disk as follows:

```
brrestore -d disk -u user/password
```

If this fails, check the error messages and resolve possible problems before you continue.

3. If you are running backups in RMAN mode, verify backup and restore directly to disk using Recovery Manager channel type disk as follows:

- a. Re-link the Oracle Server with the Database Library provided by SAP (`libobk.sl`).

Oracle8i: Use the same procedure as described in “[Linking Oracle Server with the Data Protector MML](#)” on page 36.

Oracle9i/10g: For each RMAN channel, set the `SBT_LIBRARY` parameter to point to the `libobk.sl` file.

 **IMPORTANT:**

Before you can use Data Protector again in the RMAN mode, you have to re-link the Oracle again with the Data Protector Database Library.

- b. You have to define the parameter `init` in the initialization file `initORACLE_SID.ora`.

Run the following commands:

```
brrestore -d pipe -u user/password -t online -m all
```

```
brrestore -d disk -u user/password
```

If this fails, refer to the SAP Online Help to learn how to execute backup and restore directly to disk using the SAP backup utility. Check the error message and resolve this issues before you continue.

4. Verify that the SAP backup tools correctly start backint (which is provided by Data Protector):

Move the original backint and create a test script named `backint` in the directory where the SAP backup utility resides, with the following entries:

```
#!/usr/bin/sh
echo "Test backint called as follows:"
echo "$0 $*"
echo "exiting 3 for a failure"
exit 3
```

Then start the following commands as the Oracle database user described in [“Configuring user accounts”](#) on page 168:

```
brbackup -t offline -d util_file -u user/password -c
```

If you receive backint arguments, this means that SAP is properly configured for backup using backint; otherwise you have to reconfigure SAP.

See [“Configuring SAP R/3 databases”](#) on page 173.

Configuration problems



IMPORTANT:

The procedure described in the previous sections must be performed before you start checking the Data Protector configuration.

1. Verify that the Data Protector software has been installed properly.

Refer to the *HP Data Protector installation and licensing guide* for details.

2. On clients with Oracle8i, verify that the Data Protector Database Library is linked with the Oracle executable:

Use the following command to check if the `libob2oracle8.so` on Solaris and `libob2oracle8.sl` (`libob2oracle8_64bit.sl`) on HP-UX is linked with the Oracle executable.

Export the `ORACLE_HOME` and the `ORACLE_SID` as described in [Step 1](#) on page 218.

HP-UX platform:

```
/usr/bin/chatr ORACLE_HOME/bin/oracle
```

Solaris platform:

```
/usr/bin/ldd -s ORACLE_HOME/bin/oracle
```

The output has to state that the respective Data Protector library is required by Oracle executable.

The following is an extract of the command output on HP-UX:

```
bin/oracle:
shared executable
shared library dynamic path search:
SHLIB_PATH  enabled second
embedded path disabled first Not Defined
shared library list:
static
/opt/omni/lib/libob2oracle8.sl(libob2oracle8_64bit.sl)
dynamic /usr/lib/librt.2
dynamic /usr/lib/libnss_dns.1
dynamic /usr/lib/libdld.2
```

The line starting with `SHLIB_PATH` should be returned as in the example above. If this line is different, then enable the Data Protector Database Library dynamic path as follows:

```
/usr/bin/chatr +s enable ORACLE_HOME/bin/oracle
```


3. Perform a filesystem backup of the SAP R/3 Database Server:

Perform a filesystem backup of the SAP Database Server system so that you can eliminate any potential communication problems between the SAP Database Server and the Data Protector Cell Manager system.

Do not start troubleshooting an online database backup unless you have successfully completed a filesystem backup of the SAP Database Server system.

See the online Help index “standard backup procedure” for details about how to do a filesystem backup.

4. Examine the environment variables:

If you need to export some variables before starting the Oracle Server Manager, TNS listener, or other Oracle utility, these variables must be defined in the `Environment` section of the Data Protector SAP configuration file on the Cell Manager. See “[Data Protector SAP R/3 configuration file](#)” on page 160.

5. Verify the permissions of the currently used user account:

Your user account has to enable you to perform backup or restore using Data Protector. Use the `testbar2` utility to check the permissions:

```
/opt/omni/bin/utilns/testbar2 -perform:checkuser
```

If the user account holds all required permissions, you will receive only NORMAL messages displayed on the screen.

See also “[Configuring user accounts](#)” on page 168.

6. Examine system errors:

System errors are reported in the `/var/opt/omni/log/debug.log` (HP-UX, Solaris, and Linux systems) or `/usr/omni/log/debug.log` (other UNIX systems) file on the SAP Server.

Configuration fails due to a script failure

Problem

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

Resolve the problem by reviewing the user account configuration. For details, see “[Configuring user accounts](#)” on page 168.

Backup problems

At this stage, you should have performed all the verification steps described in the previous sections. If backup still fails, proceed as follows:

1. Check your SAP Server configuration:

To check the configuration, start the following command on the SAP Server system:

```
/opt/omni/lbin/util_sap.exe -CHKCONF ORACLE_SID (HP-UX, Solaris, and Linux systems) or
```

```
/usr/omni/bin/util_sap.exe -CHKCONF ORACLE_SID (other UNIX systems)
```

In case of an error, the error number is displayed in the form *RETVAL**Error_number*.

To get the error description, start the command:

```
/opt/omni/lbin/omnigetmsg 12 Error_number (HP-UX, Solaris, and Linux systems) or
```

```
/usr/omni/bin/omnigetmsg 12 Error_number (other UNIX systems)
```

The *RETVAL*0 indicates successful configuration.

2. Verify Data Protector internal data transfer using the `testbar2` utility.

Before you run the `testbar2` utility, verify that the Cell Manager name is correctly defined on the SAP Database Server. Check the `/etc/opt/omni/client/cell_server` (HP-UX, Solaris, and Linux systems) or `/usr/omni/config/cell/cell_server` (other UNIX systems) file, which contains the name of the Cell Manager system. Then run the following command:

```
/opt/omni/bin/utilns/testbar2 -type:SAP -appname:ORACLE_SID  
-bar:backup_specification_name -perform:backup (HP-UX, Solaris,  
and Linux systems)
```

```
/usr/omni/bin/utilns/testbar2 -type:SAP -appname:ORACLE_SID  
-bar:backup_specification_name -perform:backup (other UNIX  
systems)
```

Examine the errors reported by the `testbar2` utility by clicking the **Details** button in the Data Protector **Monitor** context.

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

- a. Check that the owner of the backup specification is the Oracle OS user described in “[Configuring user accounts](#)” on page 168
- b. Check that the respective Data Protector user group has the `See private objects` user right enabled.
- c. Create an SAP backup specification to back up to a null or file device. If the backup succeeds, the problem may be related to the backup devices.

Refer to the *HP Data Protector troubleshooting guide* for instructions on troubleshooting devices.

If the test fails again, call support.

3. Verify the backup using backint

```
export OB2BARLIST=barlist_name
export OB2APPNAME=ORACLE_SID
/opt/omni/lbin/backint -f backup -t file -u ORACLE_SID -i
input_file (HP-UX, Solaris, and Linux systems)
/usr/omni/bin/backint -f backup -t file -u ORACLE_SID -i
input_file (other UNIX systems)
```

where *input_file* is a file with a list of full pathnames for backup.

Backint expects the list of files in the following format:*pathName_1 pathName_2 pathName_3*

Backup fails at the beginning with the message “Internal heap ERROR 17112”

Problem

When using SAP 4.6D kernel on HP-UX 11.11, backup fails immediately after it was started due to a BRBACKUP core dump. A line similar to the following can be found at the beginning of the message:

```
Internal heap ERROR 17112 addr=0x800003ffff7f3660
```

Action

1. Login to the SAP server as the user who is owner of the backup specification.
2. Run the command:

```
env | grep NLS_LANG
```

The output is similar to the following:

```
NLS_LANG=AMERICAN_AMERICA.US7ASCII
```

3. Add the NLS_LANG variable to the backup specification. For more details, see [“Setting, retrieving, listing and deleting Data Protector SAP R/3 configuration file parameters using the CLI”](#) on page 163.
4. Restart the backup.

Util_File_Online SAP backup fails with “semop() error”

Problem

When the *util_file_online* option is used with BRBACKUP (for example, if you select the *Brbackup_Util_File_Online* template), the tablespaces are switched

into/from backup mode individually. As there can be only one process communicating with BRBACKUP, several sapback processes are using a semaphore to synchronize their interaction with BRBACKUP.

The number of sapback processes is calculated as the sum of concurrencies of all devices used for backup. With a large number of sapback processes, the maximum number of processes that can have undo operations pending on any given IPC semaphore on the system may be exceeded. In such case, several sapback agents will fail with the following error:

```
[28] No space left on device.
```

Action

Perform any of the following actions to resolve the problem:

- Reduce the number of backup devices or their concurrency.
- Increase the value of the `semnmu` kernel parameter. After you increase the value, rebuild the kernel and reboot the system.

Backup fails with “Connect to database instance failed”

Problem

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
```

Action

Before you start a backup, ensure that the database instance is in the `open` or `shutdown` mode.

Restore problems

At this stage, you should have performed all the verification steps described in the previous sections. After this, proceed as follows:

1. Verify a user for the restore:

Verify that user specified for the restore session is the user of backup session and that he/she belongs to the Data Protector `operator` or `admin` group.

See “[Configuring user accounts](#)” on page 168.

2. Verify that a backup object exists on the backup media and in the IDB:

This can be done by executing the command

```
/opt/omni/bin/omnidb -SAP "object_name" -session  
"Session_ID" -media (HP-UX, Solaris, and Linux systems) or  
  
/usr/omni/bin/omnidb -SAP "object_name" -session  
"Session_ID" -media (other UNIX systems)
```

on the SAP Database Server system.

The output of the command lists detailed information about the specified backup object, session IDs of the backup sessions containing this object, and a list of the media used.

For detailed syntax of the `omnidb` command, run:

```
/opt/omni/bin/omnidb -help (HP-UX, Solaris, and Linux systems)  
  
/usr/omni/bin/omnidb -help (other UNIX systems)
```

You can also do this using the SAP tools:

Use `backint`, so that SAP tools will also use this command to query:

```
/opt/omni/lbin/backint -f inquiry -u ORACLE_SID -i  
input_file (HP-UX, Solaris, and Linux systems)  
  
/usr/omni/bin/backint -f inquiry -u ORACLE_SID -i  
input_file (other UNIX systems)
```

where the specified `input_file` is queried.

If this fails, check if the backup session was performed successfully and if the query was started under the appropriate user account.

`Backint` anticipates a list of files of the following format:

```
backup_ID_1 pathName_1 [targetDirectory_1]  
backup_ID_2 pathName_2 [targetDirectory_2]  
backup_ID_3 pathName_3 [targetDirectory_3]
```

To retrieve the `backup_ID` numbers, enter the following command:

```
echo "#NULL #NULL" | backint -f inquiry -u ORACLE_SID
```

or, alternatively, you can just specify `#NULL` as `backup_ID_1` in the `input_file`. In this case, the latest backup session for the file is used for the restore.

3. Verify the restore using the Data Protector user interface

This test is possible if the objects have been backed up by `backint`.

If this fails, check if the backup session was performed successfully and if the query was started under the appropriate user account.

4. Simulate a restore session

Once you know the information about the object to be restored, you can simulate a restore using the Data Protector `testbar2` utility.

Before you run `testbar2`, verify that the Cell Manager name is correctly defined on the SAP Database Server.

Check the `/etc/opt/omni/client/cell_server` (HP-UX, Solaris, and Linux systems) or `/usr/omni/config/cell/cell_server` (other UNIX systems) file, which contains the name of the Cell Manager system.

Then, test the Data Protector internal data transfer using the `testbar2` utility:

```
opt/omni/bin/utilns/testbar2 -type:SAP
-appname:ORACLE_SID
-perform:restore
-object:object_name
-version:object_version
-bar:backup_specification_name
```

(HP-UX, Solaris, and Linux systems) or

```
/opt/omni/bin/utilns/testbar2 -type:SAP
-appname:ORACLE_SID
-perform:restore
-object:object_name
-version:object_version
-bar:backup_specification_name (other UNIX systems)
```

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.

5. Verify the restore using backint

Run the following command:

- On HP-UX, Solaris, and Linux: `/opt/omni/lbin/backint -f restore -u ORACLE_SID -i input_file`
- On other UNIX: `/usr/omni/bin/backint -f restore -u ORACLE_SID -i input_file`

where the contents of the `input_file` will be restored.

If this fails, check if the session was performed successfully and if the restore was started under the appropriate user account.

Backint anticipates a list of files in the following format:
`backup_ID_1 pathName_1 [targetDirectory_1] backup_ID_2 pathName_2 [targetDirectory_2] backup_ID_3 pathName_3 [targetDirectory_3]`

To retrieve the `backup_ID` numbers, enter the following command:

```
echo #NULL #NULL | backint -f inquiry -u ORACLE_SID
```

Restore of SAP R/3 tablespaces located on raw partitions fails

Problem

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 Integrating SAP DB/MaxDB and Data Protector

Introduction

This chapter explains how to configure and use the Data Protector SAP DB/MaxDB integration (**SAP DB integration**). It describes the concepts and methods you need to understand to back up and restore SAP DB/MaxDB database objects (**SAP DB objects**).

Data Protector integrates with SAP DB/MaxDB Server to offer online backup of an SAP DB/MaxDB Server instance (**SAP DB instance**). You can back up the following SAP DB objects using the Data Protector SAP DB integration:

- SAP DB data
- SAP DB configuration
- SAP DB archive logs

During backup, the database is online and actively used. It can be in the `Admin` or `Online` mode.

Data Protector offers interactive and scheduled backups of the following types:

Table 14 Backup types

Full	SAP DB complete backup. Backs up all selected objects.
Diff	SAP DB incremental backup. Backs up changes made to the database since the last full backup. ¹
Trans	SAP DB log backup. Backs up archived logs ¹ .

¹What is actually backed up depends on which objects you select. For details, see [Table 15](#) on page 246.

You can restore SAP DB objects:

- To the original location

- To another SAP DB client
- To another SAP DB instance

As part of the restore session, you can also recover the database to a specific point in time or to the last archive log.

You can also back up and restore SAP DB objects using SAP DB utilities.

This chapter provides information specific to the Data Protector SAP DB integration. For general Data Protector procedures and options, see online Help.

Integration concepts

Data Protector integrates with the SAP DB Server through the SAP DB integration component using the SAP DB database management server and the `backint` interface.

[Figure 39](#) shows the architecture of the Data Protector SAP DB integration.

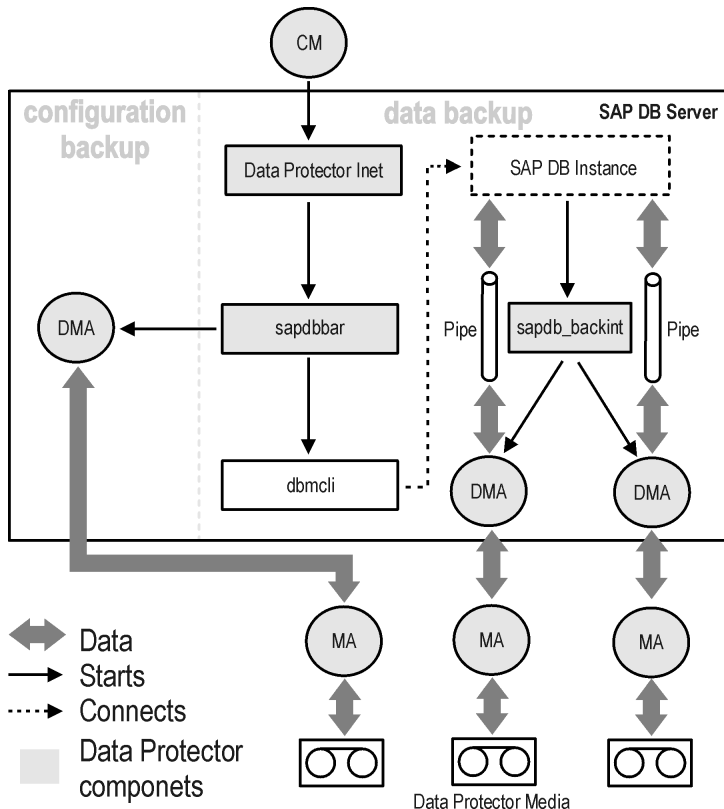


Figure 39 SAP DB integration architecture

The Data Protector integration software consists of the following components:

- The `sapdbbar` module, installed on the SAP DB Server system, which controls activities between the SAP DB Server and Data Protector backup and restore processes.
- The `sapdb_backint` component, installed on the SAP DB Server system, is a binary interface between Data Protector and backup and restore functionality of the SAP DB.
- The DMA (Data Mover Agent) component, installed on the SAP DB Server system, is the actual data transferring module, called by the `sapdb_backint`.
- The `util_sapdb` utility, which is used by Data Protector to configure an SAP DB instance to use with Data Protector and check the instance configuration.

SAP DB data and archive logs are backed up or restored in streams, whereas the SAP DB configuration is backed up or restored as ordinary files. After the backup

has finished, the archive logs can either be deleted or kept on the SAP DB Server, depending on the selected options.

The integration also takes advantage of the concept of **SAP DB media** and **media groups**, thus providing parallel backup and restore of SAP DB objects. Several SAP DB media are grouped in an SAP DB media group, which is then backed up or restored in streams. This is referred to as SAP DB **parallelism**. See [Table 16](#) for more information on the Data Protector `Parallelism` option.

 **NOTE:**

When running a backup using SAP DB utilities, SAP DB media and pipes must be configured manually.

Backup flow

When a backup session is started, the Cell Manager starts the `sapdbbar` with selected backup parameters from the backup specification. The `sapdbbar` module then starts an SAP DB session using the SAP DB `dbmcli`. The `sapdbbar` module issues `dbmcli` commands that configure SAP DB backup media (parallelism), configure `sapdb_backint` and then start the backup using SAP DB `dbmcli`. SAP DB then starts the configured `sapdb_backint` component. For every SAP DB medium (pipe) `sapdb_backint` starts a DMA, which transfers the data from SAP DB media (pipes) to Data Protector media. This procedure is the same for full, differential, and transactional backup. Additionally, if the configuration (including media specification and the backup history) is selected for backup, it is backed up directly by the `sapdbbar` module and DMA. The list of configuration files to be backed up is retrieved through `dbmcli`.

Restore flow

When a restore session is started, the Cell Manager starts the `sapdbbar` module, which starts SAP DB `dbmcli`. The `sapdbbar` module issues commands to SAP DB `dbmcli` to configure `sapdb_backint` and SAP DB backup media (parallelism). SAP DB then starts the configured `sapdb_backint`, which starts streaming data to media (pipes) that SAP DB created. For every SAP DB medium (pipe) the `sapdb_backint` starts a DMA, which transfers the data from Data Protector media to SAP DB media (pipes). If SAP DB configuration is being restored, it is the `sapdbbar` module and DMA that perform the restore.

Configuring the integration

You need to configure SAP DB users and every SAP DB instance you intend to back up from or restore to.

Prerequisites

- Ensure that you have correctly installed and configured the SAP DB system.
 - See *HP Data Protector product announcements, software notes, and references* or <http://www.hp.com/support/manuals> for supported versions, platforms, devices, and other information
 - See SAP DB documentation for information on installing, configuring, and using SAP DB Server.

To enable transactional backups (log backups), you need to activate the SAP DB Automatic Log Backup.

- Ensure that you have correctly installed Data Protector. See the *HP Data Protector installation and licensing guide* on how to install Data Protector in various architectures.

Every SAP DB system you intend to back up from or restore to must have the Data Protector `SAP DB Integration` component installed.

Limitations

The SAP DB transactional backup of a database instance is only supported with SAP DB 7.04.03 and newer SAP DB versions.

The following are not supported:

- Instance names in UNICODE format
- Pre- and post-exec options on the level of the backup specification
- Preview for SAP DB restore sessions
- Integrated offline backup of SAP DB objects

Before you begin

- Configure devices and media for use with Data Protector.
- To test whether the SAP DB system and the Cell Manager communicate properly, configure and run a Data Protector filesystem backup and restore on the SAP DB system.

Cluster-aware clients

Configure SAP DB instances only on one cluster node, since the configuration files reside on the Cell Manager.

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`

Configuring SAP DB users

Create or identify an **SAP DB database user** with at least the following SAP DB permissions:

- Saving backups (Backup)
- Restoring backups (Recovery)
- Installation management (InstallMgm)
- Parameter access (ParamCheckWrite)

The last two permissions are required for the Data Protector configuration.

UNIX only: Add the OS user under whose account SAP DB is running (**SAP DB OS user**) and user `root` to the Data Protector `admin` or `operator` group. For more information, see the online Help index: “adding users”. For example, by default, the SAP DB OS user is the user `sapdb` in the group `sapsys`.

Configuring SAP DB instances

You need to provide Data Protector with the following configuration parameters for the SAP DB instance:

- Username of the SAP DB database user.
- Password of the SAP DB database user.
- Optionally, the SAP DB independent program path parameter

To configure an SAP DB instance, use the Data Protector GUI or CLI.

Data Protector then creates the SAP DB instance configuration file on the Cell Manager and verifies the connection to the instance.

**TIP:**

Once the configuration file is created, you can set, retrieve, and list the configuration file parameters using the Data Protector `util_cmd` command. For details, see the `util_cmd` man page.

To configure an SAP DB instance, use the Data Protector GUI or CLI.

Before you begin

- Ensure that the SAP DB instance is online.

Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **SAP DB Server**, and click **Add Backup**.
3. In the **Create New Backup** dialog box, select the **Blank SAPDB Backup** template. Click **OK**.

4. In **Client**, select the SAP DB Server system. In a cluster environment, select the virtual server.

In **Application database**, type the SAP DB instance name.

UNIX only: Select/type the username and group name of the SAP DB OS user.

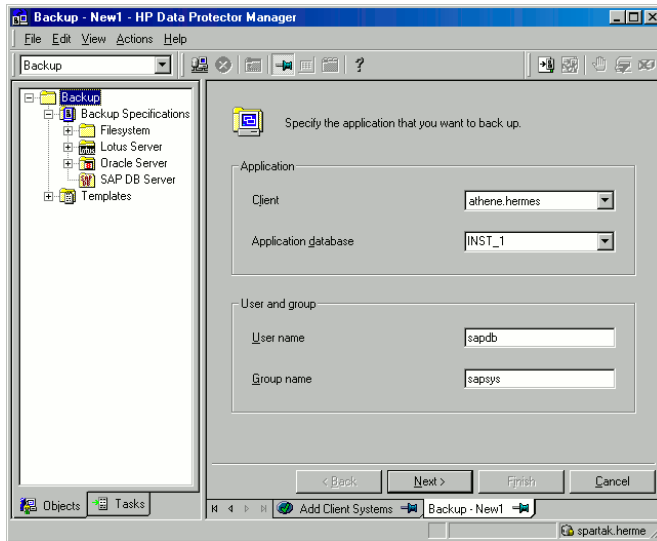


Figure 40 Specifying an SAP DB instance

Click **Next**.

5. In the **Configure SAP DB** dialog box, specify the **SAP DB independent program path** parameter. This parameter is the independent program path directory specified during the installation of the SAP DB application. To automatically detect the directory, leave the **Auto-detect** option selected.

Under **Connection**, type the username and password of the SAP DB database user as described in [Configuring SAP DB users](#).

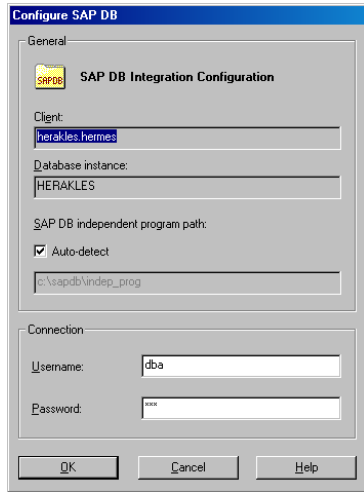


Figure 41 SAP DB configuration

Click **OK**.

6. The SAP DB instance is configured. Exit the GUI or proceed with creating the backup specification at [Step 3](#).

Using the Data Protector CLI

UNIX only: Log in to the SAP DB Server system as the SAP DB OS user.

From the directory:

Windows: `Data_Protector_home\bin`

HP-UX: `/opt/omni/lbin`

Other UNIX: `/usr/omni/bin/`

run:

```
util_sapdb.exe \[-homedir SAPDB_independent_program_directory]  
\-config Instance Name username password
```

Parameter description

SAPDB_independent_program_directory

The SAP DB independent program path parameter. This parameter is the independent program path directory specified during the installation of the SAP DB application on the SAP DB Server.

This parameter is optional. If it is not specified, the directory is detected automatically.

Instance_Name

The name of the SAP DB instance to be configured.

username

The username of the SAP DB database user created or identified as described in [Configuring SAP DB users](#).

password

The password of the SAP DB database user created or identified as described in [Configuring SAP DB users](#)



NOTE:

The username and the SAP DB independent program path parameter must not contain the single quote character (').

Example

To configure the instance `sapdb_inst` by specifying the database user `sapdb_user` with the password `sapdb_pass`, and the SAP DB independent program path `/opt/sapdb/indep_prog` (UNIX) or `c:\program files\sapdb\indep_prog` (Windows), run:

Windows:

```
util_sapdb.exe -homedir "SAPDB_independent_program_directory"  
-config sapdb_inst sapdb_user sapdb_pass
```

UNIX:

```
util_sapdb.exe -homedir  
SAPDB_independent_program_directory/indep_prog -config  
sapdb_inst sapdb_user sapdb_pass
```



TIP:

To change the configuration parameters, run the same command using new values.

Handling errors

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

UNIX only: To obtain an error description, from the directory:

HP-UX: `/opt/omni/lbin`

Other UNIX: `/usr/omni/bin/`

run:

`omnigetmsg 12 Error_number`

Checking the configuration

Check the configuration of an SAP DB instance after you have created at least one backup specification for the SAP DB instance. 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 DB Server**. Click the backup specification to display the SAP DB instance to be checked.
3. Right-click the SAP DB instance and click **Check configuration**.

Using the Data Protector CLI

UNIX only: Log in to the SAP DB Server system as the SAP DB OS user.

From the directory:

Windows: `Data_Protector_home\bin`

HP-UX: `/opt/omni/lbin`

Other UNIX: `/usr/omni/bin/`

run:

```
util_sapdb.exe -chkconf Instance_Name
```

where *Instance_Name* is the name of the SAP DB instance.

A successful configuration check displays the message *RETVAL*0.

Backup

The integration provides online database backups of different types. What is backed up depends on which objects and backup type you select. See [Table 15](#).

Table 15 What is backed up

		SAP DB Backup Mode		
		Full	Diff	Trans
GUI Selections	Data	data	diff on data	archive logs
	Configuration	configuration	configuration	configuration
	Instance	data + configuration	diff on data + configuration	archive logs + configuration

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 DB Server**, and click **Add Backup**.
3. In the **Create New Backup** dialog box, select the **Blank SAPDB Backup** template. Click **OK**.

4. In **Client**, select the SAP DB Server system. In a cluster environment, select the virtual server.

In **Application database**, type the SAP DB instance name.

UNIX only: Type the username and group name of the SAP DB OS user. This user will be the backup owner.

Click **Next**.

5. If the SAP DB instance is not configured yet for use with Data Protector, the **Configure SAP DB** dialog box is displayed. Configure it as described in [Configuring SAP DB instances](#).
6. Select the SAP DB objects you want to back up.

 **IMPORTANT:**

To back up SAP DB archive logs, select the **Data** item. The archive log backup is then triggered by selecting the **Trans** backup type when scheduling the backup or running the backup interactively.

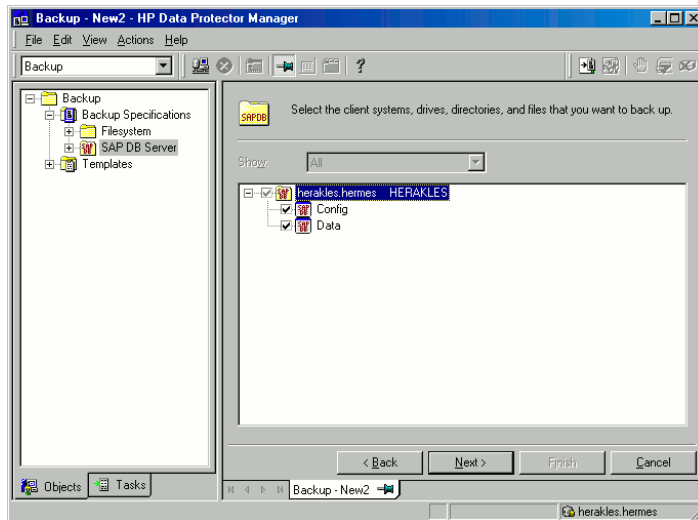


Figure 42 Selecting SAP DB objects

7. Select devices to use for the backup.

To specify device options, right-click the device and click **Properties**. Specify the device **Concurrency**, media pool, and preallocation policy.

Click **Next**.

8. Set backup options. For information on application specific options (Figure 43), see Table 16.

Click **Next**.

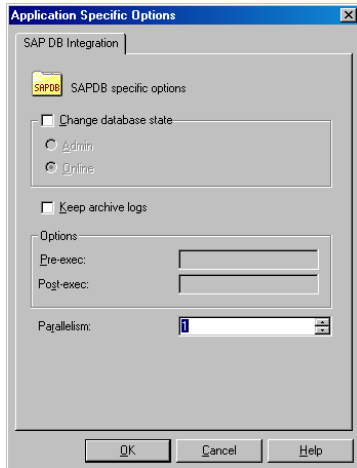


Figure 43 Application specific options

9. Optionally, schedule the backup. See [Scheduling backup specifications](#).

Click **Next**.

10. Save the backup specification, specifying a name and a backup specification group.

 **TIP:**

Save the backup specifications in the group **SAP DB Integration**.

 **TIP:**

Preview your backup specification before using it for real. See [Previewing backup sessions](#).

Table 16 SAP DB backup options

Option	Description
Change database state	Specifies the SAP DB database mode during backup (Admin or Online). If this option is OFF, the database remains in the current mode.
Keep archive logs	Specifies whether to keep (ON) or delete (OFF) archive logs on the SAP DB Server after the backup has finished.
Parallelism	<p>Specifies the number of SAP DB media created on the SAP DB Server and consequently the number of SAP DB backup data streams.</p> <p>The value must be equal to or lower than:</p> <ul style="list-style-type: none">• The SAP DB <code>MAXBACKUPDEVS</code> parameter.• The sum of concurrency values of all backup devices selected in the backup specification. <p>For more information on the Data Protector <code>Concurrency</code> option, see the online Help index: "concurrency".</p> <p>Default value: 1</p> <p>Maximum value: 32</p> <p>Recommended value: the number of SAP DB data volumes to be backed up</p>

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 back up SAP DB objects 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**. See [Figure 44](#).

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.

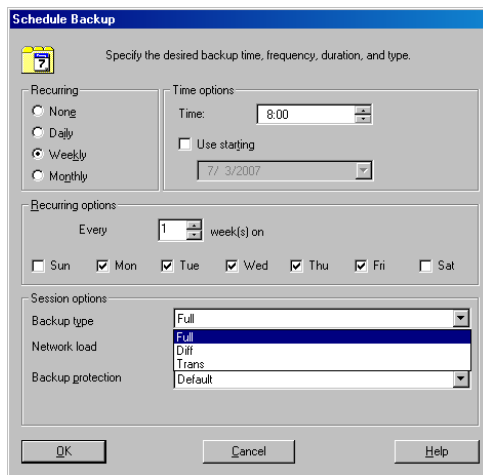


Figure 44 Scheduling the backup specification

Previewing backup sessions

Preview the backup session using the Data Protector GUI or CLI to test it.

This interactive test does not back up any data. However, as a result of this test, the following file is created on the SAP DB Server system:

Windows:

```
Data_Protector_home\tmp\Backup_Specification_Name_TEST_FILE
```

UNIX:

```
/var/opt/omni/tmp/Backup_Specification_Name_TEST_FILE
```

Delete the file after the test.

Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **SAP DB Server**. Right-click the backup specification you want to preview and click **Preview Backup**.
3. Specify the **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 and Solaris: `/opt/omni/bin/`

Other UNIX: `/usr/omni/bin/`

run:

```
omnib -sapdb_list backup_specification_name -test_bar
```

What happens during the preview?

1. The `sapdbbar` program is started, which then starts the Data Protector `testbar2` command.
2. Data Protector tests the Data Protector part of the configuration. The following is tested:
 - Communication between the SAP DB instance and Data Protector
 - 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 the SAP DB objects selected in the backup specification in any of the following ways:

- Use the Data Protector GUI.
- Use the Data Protector CLI.
- Use the SAP DB utilities.

Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, and then **SAP DB Integration**. Right-click the backup specification you want to start and click Start Backup.
3. Select the **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

From the directory:

Windows: `Data_Protector_home\bin`

HP-UX and Solaris: `/opt/omni/bin/`

Other UNIX: `/usr/omni/bin/`

run:

```
omnib -sapdb_list ListName [-barmode sapdbmode] [list_options]  
[-preview]
```

ListName is the name of the backup specification.

sapdbmode specifies the backup type. You can select full, diff, or trans.

For *List_options*, see the omnib man page.

Example

To start a full backup using an existing SAP DB backup specification called `TEST`, and to set data protection to 10 weeks, run:

```
omnib -sapdb_list TEST -barmode full -protect weeks 10
```

Using SAP DB utilities

For a description of the variables listed below, see [Parameter description](#).

1. Create the `bsi_env` file on the SAP DB Server system.

UNIX only: Give the SAP DB OS user read permission for this file.

The file must contain the following lines:

Windows:

```
BACKINT Data_Protector_home\bin\sapdb_backint.exe
INPUT Data_Protector_home\tmp\inst_name.bsi_in
OUTPUT Data_Protector_home\tmp\inst_name.bsi_out
ERROROUTPUT Data_Protector_home\tmp\inst_name.bsi_err
PARAMETERFILE name_of_backup_spec
TIMEOUT_SUCCESS 60
TIMEOUT_FAILURE 30
```

HP-UX:

```
BACKINT /opt/omni/bin/sapdb_backint
INPUT /var/opt/omni/tmp/inst_name.bsi_in
OUTPUT /var/opt/omni/tmp/inst_name.bsi_out
ERROROUTPUT /var/opt/omni/tmp/inst_name.bsi_err
PARAMETERFILE name_of_backup_spec
TIMEOUT_SUCCESS 60
TIMEOUT_FAILURE 30
```

Other UNIX:

```
BACKINT /usr/omni/bin/sapdb_backint
INPUT /var/opt/omni/tmp/inst_name.bsi_in
OUTPUT /var/opt/omni/tmp/inst_name.bsi_out
ERROROUTPUT /var/opt/omni/tmp/inst_name.bsi_err
PARAMETERFILE name_of_backup_spec
TIMEOUT_SUCCESS 60
TIMEOUT_FAILURE 30
```

2. Log in to the SAP DB database manager as the SAP DB database user by running:

```
dbmcli -d inst_name -u username,password
```

3. In the SAP DB database manager, register the location of the `bsi_env` file created in [Step 1](#) of this procedure by running:

Windows:

```
dbm_configset -raw BSI_ENV location\inst_name.bsi_env
```

UNIX:

```
dbm_configset -raw BSI_ENV location/inst_name.bsi_env
```

4. Create SAP DB media, grouping them under the same name (*media_group_name*). The number of created media should equal the parallelism you plan to use for backup. To create a medium *medium_name*, run:

```
medium_put media_group_name/medium_name pipe_name
medium_type backup_type
```

where *backup_type* can be one of the following:

- DATA for full backup
- PAGES for differential backup
- LOG for log backup

 **IMPORTANT:**

When creating SAP DB media for the purpose of a Data Protector backup and restore, the media group name must begin with the “BACK” string.

Example

The commands below create two media and two pipes (parallelism = 2) in the media group BACKDP-Data[2].

Windows:

```
medium_put BACKDP-Data[2]/1 \
\\.\Pipe\inst_name.BACKDP_Data[2].1 PIPE DATA

medium_put BACKDP-Data[2]/2 \
\\.\Pipe\inst_name.BACKDP_Data[2].2 PIPE DATA
```

UNIX:

```
medium_put BACKDP-Data[2]/1 \
/var/opt/omni/tmp/inst_name.BACKDP_Data[2].1 PIPE DATA

medium_put BACKDP-Data[2]/2 \
/var/opt/omni/tmp/inst_name.BACKDP_Data[2].2 PIPE DATA
```

5. Start the SAP DB utility session by running:

```
util_connect
```

6. Start the backup. The following exemplary command starts the full backup for the media created in [Step 4](#) on page 255 of this procedure:

```
backup_start BACKDP-Data[2] DATA
```

7. The progress of the session is displayed in the Data Protector **Monitor** context. For more information, see [Monitoring sessions](#).

Parameter description

<i>inst_name</i>	Name of the instance to be backed up.
<i>name_of_backup_spec</i>	Name of the Data Protector backup specification to be used for backup.
<i>username,password</i>	Connection string for the SAP DB database user.
<i>location</i>	Location of the <code>bsi_env</code> file.
<i>media_group_name</i>	Name of the SAP DB media group.
<i>medium_name</i>	Name of the SAP DB medium.
<i>pipe_name</i>	Name of the SAP DB pipe.
<i>medium_type</i>	Type of the SAP DB medium.

Restore

Restore SAP DB objects in any of the following ways:

- Use the Data Protector GUI. See [Restoring using the Data Protector GUI](#).
- Use the Data Protector CLI. See [Restoring using the Data Protector CLI](#).
- Use the SAP DB utilities. See [Restoring using SAP DB utilities](#).

Restore and recovery overview

This section provides an overview of restore and recovery process with regard to Data Protector restore and recovery options selection. For a detailed description of these options, see [SAP DB restore options](#).

At the beginning of a restore session, Data Protector switches the SAP DB database to the `Admin` mode. If the database cannot be switched to the `Admin` mode, an error is issued in the Data Protector monitor.

Depending on the type of restore and on the selected restore and recovery options, the SAP DB database can be switched to the following modes after the restore:

- If the Data Protector `Recovery` option is selected, the database is switched to the `Online` mode after the restore.
- If the Data Protector `Recovery` option is *not* selected and archive logs have not been restored (if restore from a full or diff backup session is performed), the database remains in the `Admin` mode after the restore.
- If the Data Protector `Recovery` option is *not* selected and archive logs have been restored, the database is, if the restored archive logs allow it, switched to the `Online` mode. If the database, however, cannot be switched to the `Online` mode (because the restored archive logs do not allow it), it remains in the `Admin` mode.



IMPORTANT:

There are several scenarios, depending on the backup option `Keep archive logs` and the recovery option `Use existing archive logs`, in which a gap of transactions between the sequence of redo logs on the SAP DB Server and the restored volumes can occur. When performing recovery (when the database is switched to the `Online` mode), SAP DB always checks whether such a gap exists, regardless of the point in time selected for recovery. If such a gap exists, the recovery is not performed and the database remains in the `Admin` mode, unless the existing redo logs are manually deleted before starting the restore.

If a full or diff backup session is restored, only the data (no archive logs) from the selected backup session is restored. The data on the SAP DB Server is overwritten.

If a trans backup session is restored, only the archive logs (no data) from the selected backup session are restored.

During the restore, the redo logs that existed on the SAP DB Server before the restore are not deleted during the restore.

When restoring, the existing redo logs on the SAP DB Server can be, depending on the Data Protector `Use existing archive logs` option selection (it can be selected only if the `Recovery` option is selected), handled as follows:

- If the `Use existing archive logs` option is selected, the existing archive logs on the SAP DB Server are applied to the redo logs.
When a transactional backup session is selected for restore, or when it is a part of the needed restore chain, and the `Use existing archive logs` option is selected at the same time, the archive logs from Data Protector media are

applied to redo logs. Thereafter, the archive logs on the SAP DB Server are applied to redo logs.

- If the `Use existing archive logs` option is not selected, the backed up archive logs on backup media are applied to the redo logs (if trans backup session is restored), or the redo logs are left intact together with the existing archive logs on the SAP DB Server (if full or diff backup session is restored).

 **NOTE:**

The `Use existing archive logs` option is disabled in case of SAP DB migration, thus allowing only for the restore of redo logs from the backed up archive logs on backup media (if trans backup session is restored).

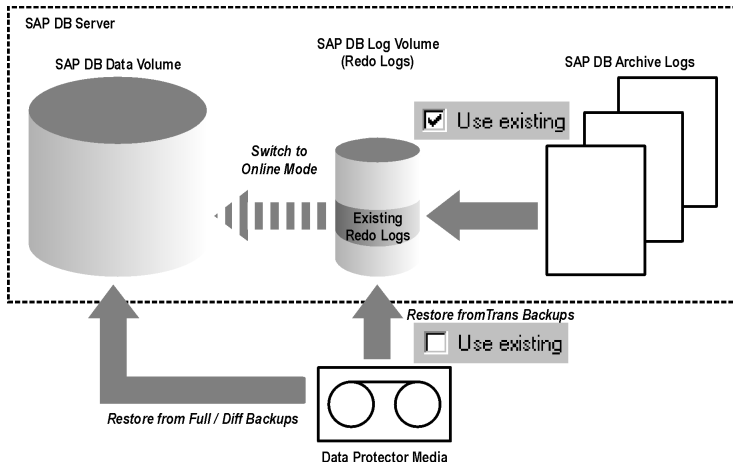


Figure 45 SAP DB restore process

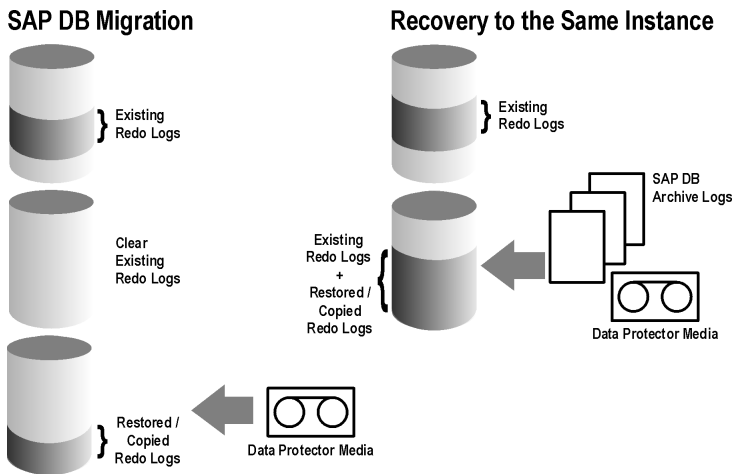


Figure 46 SAP DB archive logs restore process—redo logs details

If you select a differential or a transactional backup session to be restored, you can set the integration to:

- Perform a full database restore. In this case, the integration automatically determines the chain of needed full, differential or transactional backup sessions when performing the restore. After the restore has finished, the database is, if the `Recovery` option is selected, switched to the `Online` mode.
- Restore only the selected differential or the selected transactional backup session. If the database is consistent after such a restore and if the `Recovery` option is

selected, it is switched to the `Online` mode. Otherwise, the database is left in the `Admin` mode.

Restoring only the selected trans or diff backup session is useful if the database remains offline or in the `Admin` mode after a restore from full backup session, which is then followed by a restore from diff or trans backup session.



NOTE:

During the restore or migration, the archive logs on the SAP DB Server are never deleted.

Before you begin

If you intend to restore to another SAP DB instance:

- Install the Data Protector SAP DB/Max DB integration on the SAP DB Server system to which you want to restore.
- Add the SAP DB client to the Data Protector cell.
- Configure SAP DB/Max DB users as described in [Configuring SAP DB users](#).
- Configure the instance to which you want to perform the restore. See [Configuring SAP DB instances](#).



NOTE:

If you are using the Data Protector GUI, you can configure the instance during the restore process.

During the restore to another SAP DB instance, the existing data is overwritten and the existing redo logs are deleted.

Restoring using the Data Protector GUI

1. In the **Context List**, click **Restore**.
2. In the Scoping Pane, expand **SAP DB Server**, expand the client from which the data to be restored was backed up, and then click the SAP DB/Max DB instance you want to restore.

3. In the **Source** page, select objects for restore.

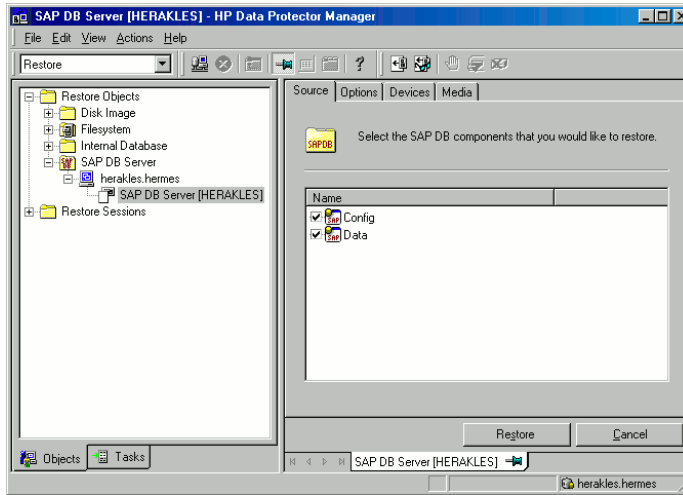


Figure 47 Selecting objects for restore

To restore SAP DB objects from a specific backup session, right-click the **Data** item, click **Properties**, and specify **Backup version** in the **Properties for Data** dialog box.

Selecting a `Trans` or `Diff` backup session enables you to:

- Perform a full restore of the database (`Full restore of database`). In this case, the integration automatically determines the chain of needed full, differential, or transactional backup sessions.
- Restore only the selected backup session (`Restore only this backup`). Restoring only the selected `Trans` or `Diff` backup session is useful if the database remains offline or in the `Admin` mode after a restore from a full backup session.

To restore SAP DB archive logs, select the `Data` item and a `Trans` backup session to restore from.

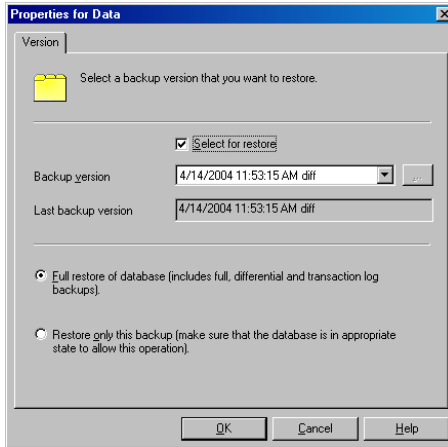


Figure 48 Properties for data



IMPORTANT:

The **Configuration** item is restored from the same backup session as selected for the **Data** item, regardless of what you select for the **Configuration** item.

4. In the **Options** page, set the restore and recovery options. For information, see [SAP DB restore options](#).
5. In the **Devices** page, select devices to use for the restore.

The **Automatic device selection** option is selected by default, but it is recommended to select the **Original device selection** option.



IMPORTANT:

If you decide to select the **Automatic device selection** option, ensure that the number of available devices is equal to or greater than the number of devices that were used for backup.

6. In the **Media** page, view the media needed for the restore and verify its availability.
7. Click **Restore**.
8. In the **Start Restore Session** dialog box, click **Next**.

9. Specify the **Report level** and **Network load**.

Click **Finish** to start the restore.

The message `Session completed successfully` is displayed at the end of a successful session.

Restoring using the Data Protector CLI

From the directory:

Windows: `Data_Protector_home\bin`

HP-UX: `/opt/omni/bin/`

Other UNIX: `/usr/omni/bin/`

run:

```
omnir -sapdb -barhost ClientName -instance InstanceName
[-destination ClientName]
[-newinstance DestinationInstanceName]
[-session SessionID]
[-recover [-endlogs | -time: YYYY-MM-DD.hh.mm.ss]
[-from_disk]]
[-nochain]
```

The `-barhost` option sets the name of the SAP DB Server that was backed up.

The `-instance` option sets the name of the SAP DB instance that was backed up.

The `-session` option selects the backup session to be restored. If this option is not specified, the last backup session is restored, regardless of the `-endlogs` or the `-time` option selection.

When restoring objects that have copies do not use the copy session ID, but the object's backup ID, which equals the object's backup session ID.

The `-nochain` option instructs the integration to restore only the selected or last backup session; the integration does not restore the whole restore chain of full, differential, and transactional backups.

For descriptions of all other options, see [SAP DB restore options](#). Refer also to the `omnir` man page.

Example

To restore an instance named "inst1" (together with configuration), backed up on an SAP DB Server named "srv1.company.com" from the last backup session and then perform a recovery until the end of logs, enter the following command:

On HP-UX:

```
/opt/omni/bin/omnir -sapdb -barhost srv1.company.com -instance  
inst1 -recover -endlogs
```

On other UNIX:

```
/usr/omni/bin/omnir -sapdb -barhost srv1.company.com -instance  
inst1 -recover -endlogs
```

On Windows:

```
Data_Protector_home\bin\omnir -sapdb -barhost srv1.company.com  
-instance inst1 -recover -endlogs
```

On how to find information needed for restore, see [Finding information needed for restore](#) (TBD).

Restoring using SAP DB utilities

Using this integration, it is also possible to run an integrated Data Protector restore of an SAP DB Server from SAP DB utilities.

To perform a restore to an existing SAP DB Server instance, see [SAP DB restore and recovery](#).

To migrate an SAP DB instance, see [SAP DB migration](#).

Finding information needed for restore

To find the information needed for a restore, follow the steps below:

Execute the following commands:

- `omnidb -sapdb`
to get a list of SAP DB objects.
- `omnidb -sapdb object_name`
to get details on a specific object, including the SessionID.

SAP DB restore and recovery

Follow the procedure on the next few pages to restore and recover a database using SAP DB utilities from existing Data Protector SAP DB backup session(s). In the procedure, the following conventions are used:

inst_name is the name of the instance to be restored

username,password is the connection string for the SAP DB database user created or identified as described in [Configuring SAP DB users](#)

location is the location of the *bsi_env* file

media_group_name is the name of the SAP DB media group

medium_name is the name of the SAP DB medium

pipe_name is the name of the SAP DB pipe

medium_type is the type of the SAP DB medium

SessionID is the Data Protector session ID of the session to be restored

1. Skip this step if the `bsi_env` file is already present and configured on the SAP DB Server.

On the SAP DB Server create the `bsi_env` file in a directory of your choice. It must contain the following lines:

Windows:

```
BACKINT Data_Protector_home\bin\sapdb_backint.exe
INPUT Data_Protector_home\tmp\inst_name.bsi_in
OUTPUT Data_Protector_home\tmp\inst_name.bsi_out
ERROROUTPUT Data_Protector_home\tmp\inst_name.bsi_err
TIMEOUT_SUCCESS 60
TIMEOUT_FAILURE 30
```

HP-UX:

```
BACKINT /opt/omni/bin/sapdb_backint
INPUT /var/opt/omni/tmp/inst_name.bsi_in
OUTPUT /var/opt/omni/tmp/inst_name.bsi_out
ERROROUTPUT /var/opt/omni/tmp/inst_name.bsi_err
TIMEOUT_SUCCESS 60
TIMEOUT_FAILURE 30
```

Other UNIX:

```
BACKINT /usr/omni/bin/sapdb_backint
INPUT /var/opt/omni/tmp/inst_name.bsi_in
OUTPUT /var/opt/omni/tmp/inst_name.bsi_out
ERROROUTPUT /var/opt/omni/tmp/inst_name.bsi_err
TIMEOUT_SUCCESS 60
TIMEOUT_FAILURE 30
```

2. Login to the SAP DB database manager as the SAP DB database user created or identified as described in [Configuring SAP DB users](#). On the SAP DB Server, execute the following command to login:

```
dbmcli -d inst_name -u username,password
```

3. In the SAP DB database manager, switch the database to the Admin mode by executing the following command:

```
db_admin
```

4. Skip this step if the location of the `bsi_env` file is already registered on the SAP DB Server.

Register the location of the `bsi_env` file as follows:

Windows:

```
dbm_configset -raw BSI_ENV location\inst_name.bsi_env
```

UNIX:

```
dbm_configset -raw BSI_ENV location/inst_name.bsi_env
```

5. Skip this step if the SAP DB media and pipes to be used with Data Protector are already existing on the SAP DB Server.

Note that to restore a Data Protector SAP DB backup session, the number of SAP DB media and pipes required equals the parallelism value used during the backup session.

Create SAP DB media in an SAP DB media group. Execute the following command for every medium to be created:

```
medium_put media_group_name/medium_name pipe_name
media_type backup_type
```

Where *backup_type* can be one of the following:

- DATA for full backup
- PAGES for differential (diff) backup
- LOG for transactional (trans) backup

 **IMPORTANT:**

When creating SAP DB media and pipes for the purpose of a Data Protector backup and restore, the media group name must begin with the “BACK” string. The commands below create two media and two pipes (parallelism = 2) in a media group:

Windows:

```
medium_put BACKDP-Data[2]/1 \
\\.\Pipe\inst_name.BACKDP_Data[2].1 PIPE DATA
```

```
medium_put BACKDP-Data[2]/2 \
\\.\Pipe\inst_name.BACKDP_Data[2].2 PIPE DATA
```

UNIX:

```
medium_put BACKDP-Data[2]/1 \
/var/opt/omni/tmp/inst_name.BACKDP_Data[2].1 PIPE DATA
```

```
medium_put BACKDP-Data[2]/2 \
/var/opt/omni/tmp/inst_name.BACKDP_Data[2].2 PIPE DATA
```

6. Start the SAP DB utility session by executing the following command:

```
util_connect
```

7. Start the restore from a Data Protector backup session by executing the following command:

```
recover_start media_group_name backup_type EBID "inst_name
SessionID:1 pipe_name1,inst_name SessionID:2 pipe_name2[,
...]"
```

Windows:

```
recover_start BACKDP-Data[2] DATA EBID "inst_name
SessionID:1 \\.\Pipe\inst_name.BACKDP-Data[2].1,TEST
SessionID:2 \\.\Pipe\inst_name.BACKDP-Data[2].2"
```

UNIX:

```
recover_start BACKDP-Data[2] DATA EBID "inst_name
SessionID:1
/var/opt/omni/tmp/inst_name.BACKDP-Data[2].1,inst_name
SessionID:2 /var/opt/omni/tmp/inst_name.BACKDP-Data[2].2"
```

Repeat this step for every session in the required chain of backup sessions.

8. When the restore has finished, the database can be recovered either until the last redo log or until the specified point in time.
 - a. To recover the database until the last redo log, execute the following command in the SAP DB database manager:

```
db_online
```

- b. To recover the database until the specified point in time, execute the following command in the SAP DB database manager:

```
db_warm -f -u yyyymmdd hhmmss
```

Where *yyyymmdd* and *hhmmss* parameters set the time for the last redo log to be applied.

SAP DB migration

When performing an SAP DB migration, some additional tasks must first be done in order to prepare the SAP DB Server or instance. These tasks are described in [Before you begin](#).

Follow the procedure in the section [SAP DB restore and recovery](#) to migrate the SAP DB database using SAP DB utilities from existing Data Protector SAP DB backup session(s). When following the mentioned procedure, *before* executing the `recover_start` command, delete the existing redo logs on the SAP DB Server by executing the following command in the SAP DB database manager:

```
util_execute clear log
```

SAP DB restore options

Figure 49 shows the SAP DB GUI restore and recovery options.

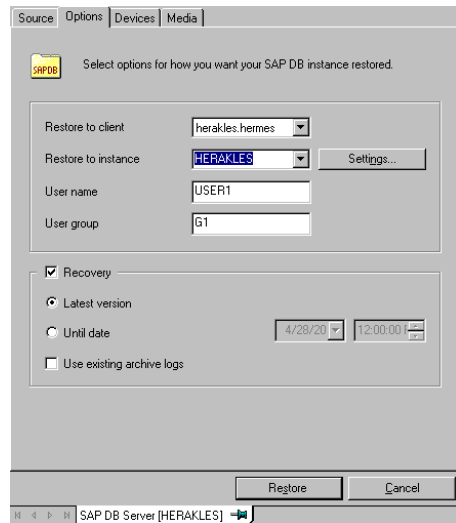


Figure 49 SAP DB restore and recovery options

The following are SAP DB specific backup options:

Migration Options

To restore selected SAP DB object to the same SAP DB Server and instance, leave the migration options as they are. Use the migration options only in case of SAP DB migration (when restoring to some other SAP DB Server or to some other instance than those that were backed up).

The following are descriptions of the migration options. First the GUI option is given, followed by a slash (/), CLI equivalent, and then description.

Restore to client / `-destination ClientName`

When using the GUI, in the drop-down list, select an SAP DB Server to which you want to restore the database.

When using the CLI, specify the `-destination` option and the name of the SAP DB Server as the `ClientName` argument.

The selected SAP DB Server must be a part of the Data Protector cell and must have the Data Protector SAP DB Integration software component installed.

Restore to instance / `-newinstance DestinationInstanceName`

When using the GUI, you can either:

- Select an instance in the **Restore to instance** drop-down list. The drop-down list shows only the instances that are already configured for use with this integration. See [Configuring SAP DB instances](#) for information on how to configure an SAP DB Server for use with this integration.
- Enter the name of an existing instance, not yet configured for use with this integration. In this case, click on the **Settings** button to configure the specified instance.

When using the CLI, the instance specified as the `DestinationInstanceName` argument to the `-newinstance` option must already be configured for use with this integration. See [Configuring SAP DB instances](#) for information on how to configure an SAP DB Server for use with this integration.

User name and User group / N/A

On UNIX, you can change the user name and the group name for the OS system user, under whose account the SAP DB application is running on the SAP DB Server (for example, the `sapdb` user in the `sapsys` group). By default, the user that started the Data Protector GUI is set for this option.

When using the CLI, it is not possible to change the user name and the group name. The same user as used during the backup session is used.

Settings / N/A

Click this button if the instance you are restoring to is not yet configured for use with this integration. See [Configuring SAP DB instances](#) for information on parameters that must be entered.

When using the CLI, this option is not available. To configure the instance, use the `util_sapdb.exe` utility as described in [Configuring SAP DB instances](#).

Recovery Options

Use the recovery options to recover the database by applying the redo logs until the latest version or until the specified date and time.

 **IMPORTANT:**

There are several scenarios, depending on the backup option `Keep archive logs` and the recovery option `Use existing archive logs`, in which a gap of transactions between the sequence of redo logs on the SAP DB Server and the restored volumes can occur. When performing recovery (when the database is switched to the `Online` mode), SAP DB always checks whether such a gap exists, regardless of the point in time selected for recovery. If such a gap exists, the recovery is not performed and the database remains in the `Admin` mode, unless the existing redo logs are manually deleted before starting the restore.

The following are descriptions of the recovery options. First the GUI option is given, followed by a slash (/), CLI equivalent, and then description.

Recovery / `-recover`

When this option is selected, the database is recovered after the restore (it is switched to `Online` mode) by applying the redo logs until the latest version (if the **Latest version** option is selected) or until the specified date and time (if the **Until date** option is selected).

 **IMPORTANT:**

When using this option, make sure that the backup session selected in the Properties for Data dialog box (when using GUI) or by the `-session` option (when using CLI) will restore enough data for the integration to apply the redo logs until the latest version or until the specified date and time. For information on how to access the Properties for Data dialog box, see [Step 3](#) on page 261. For information on the `-session` option, refer to [“Restoring using the Data Protector CLI”](#) on page 263.

When this option is not selected, all other recovery options are disabled and the following happens after the restore:

- If archive logs are not restored (if restore from a full backup session is performed), the database remains in the `Admin` mode after the restore.
- If archive logs are restored, the database is, if the restored archive logs allow it, switched to the `Online` mode. If the database, however, cannot be switched to the `Online` mode (because the restored archive logs do not allow it), it remains in the `Admin` mode.

Latest version / `-endlogs`

Select this option to recover the database until the last log.

When using the CLI, this is the default option.

Until date / `-time: YYYY-MM-DD.hh.mm.ss`

When using the GUI, select this option to recover the database until the point you select in the **Until date** drop-down menu.

When using the CLI, specify the `-time:` option if you want to recover the database until the point specified by the `YYYY-MM-DD.hh.mm.ss` argument.

NOTE:

The selected time is the system time on the system running the Data Protector GUI or CLI. If the system to be recovered is not in the same time zone as the system running the Data Protector GUI or CLI, the point of recovery is adjusted to the local time setting on the system to be restored.

Use existing archive logs / `-from_disk`

Select this option to copy the existing archive logs on the SAP DB Server to SAP DB Server redo logs.

If this option is not selected, the backed up archive logs on backup media are applied to the redo logs (if trans backup session is restored), or the redo logs are left intact together with the existing archive logs on the SAP DB Server (if full or diff backup session is restored).

When a transactional backup session is selected for restore or when it is a part of the needed restore chain, and the **Use existing archive logs** option is selected at the same time, the archive logs from Data Protector media are applied to the redo logs. Thereafter, the archive logs on the SAP DB Server are applied to redo logs.

NOTE:

The **Use existing archive logs** option is disabled in case of SAP DB migration, thus allowing only for the restore of redo logs from the backed up archive logs on backup media (if trans backup session is restored).

Restoring using another device

You can restore using a device other than that used for backup.

On how to specify another device for restore using the Data Protector GUI, see the online Help index: “restore, selecting devices for”.

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”.

Troubleshooting

This section lists problems you might encounter when using the Data Protector SAP DB integration.

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. See 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.

Problems

Problem

Data Protector reports the following error during backup or restore:

```
[Critical] From: OB2BAR_SAPDBBAR@machine.company.com "INSTANCE" Time: 02/06/04 18
-24920,ERR_BACKUPOP: backup operation was unsuccessful The database was unable to
(-2025, Invalid number of backup devices).
```

Action

Increase the value of the SAP DB `MAXBACKUPDEVS` parameter to a value that is greater than or equal to the value of the Data Protector `Parallelism` option, or reduce the value of the Data Protector `Parallelism` option.

Problem

An SAP DB instance cannot be started after restore.

Action

Using the SAP DB `db_restartinfo` command, check if the instance can be restarted.

- If the instance cannot be restarted, most probably the existing log volumes do not contain enough data to restart the instance from data volumes. The required differential or transactional backups might not have been restored.
- If the instance can be restarted, check the SAP DB instance kernel error file for errors.

If there was insufficient space for SAP DB logs at some point of time, logs might have been corrupted: delete the logs (using the `dbmcli util_execute clear log` command) or contact SAP DB or Data Protector support.

Problem

A restore from an object copy hangs.

Action

Before restarting the restore:

- Increase the number of Disk Agent buffers for the device used for the restore.
- If all objects of the backup are recorded in the IDB, perform the following steps:

1. In the Internal Database context of the Data Protector GUI, search for all objects belonging to the same backup. The objects are identified by the same backup ID.
2. Copy each object in a separate object copy session to a separate device, for example a file library. For each object, use a separate medium with the non-appendable media policy.
3. Set the highest media location priority for the newly created copies.

Problem

Data Protector reports the following error:

```
Error: SAPDB responded with:  
Error! Connection failed to node (local) for database CLUSTER:  
connection refused: x_server not running.
```

Action

Start the SAP DB x_server. Refer to the SAP DB documentation for information on how to do that.

Problem

Data Protector reports the following error:

```
Error: SAPDB responded with:  
-24988,ERR_SQL: sql error  
1,database not running
```

Action

Start the SAP DB instance. Refer to the SAP DB documentation for information on how to do that.

Problem

Data Protector reports the following error:

```
Error: SAPDB responded with:  
-24988,ERR_SQL: sql error1,utility session is already in use
```

Action

Some other user is connected to the SAP DB instance and is performing administrative tasks (utility session). Such SAP DB tasks are of the "Utility" type and can be displayed using the `dbmcli show task` command. Finish these tasks.

Problem

Data Protector reports the following error:

```
Error: SAPDB responded with:  
-24950,ERR_USRFAIL: user authorization failed
```

Action

Reconfigure the SAP DB instance as described in the section [Configuring SAP DB instances](#).

Problem

Data Protector reports the following error during backup or restore:

```
Error: SAPDB responded with:  
-24920,ERR_BACKUPOP: backup operation was unsuccessful  
The backup tool was killed with -1 as sum of exit codes. The database request ended
```

Action

Set the `TimeoutSuccess` environment variable on the Cell Manager by running the following command:

```
util_cmd -putopt SAPDB SAPDB_instance TimeoutSuccess 1000  
-sublist Environment
```

For more information, see the `util_cmd` man page.

You can also set the `TimeoutSuccess` environment variable using the Data Protector GUI. Select the backup specification in the Scoping Pane, then right-click the SAP DB instance object in the Results Pane under the **Source** tab and select the **Set Environment Variables** from the pop-up menu.

SAP DB cluster-related troubleshooting

In a cluster environment, the environment variable `OB2BARHOSTNAME` must be defined as the virtual hostname before performing some procedures run from the command line (on the client). The `OB2BARHOSTNAME` variable is set as follows:

UNIX

```
export OB2BARHOSTNAME=virtual_hostname
```

Windows

```
set OB2BARHOSTNAME=virtual_hostname
```

Glossary

access rights	See user rights .
ACSL	<i>(StorageTek specific term)</i> The Automated Cartridge System Library Server (ACSL) software that manages the Automated Cartridge System (ACS).
Active Directory	<i>(Windows specific term)</i> 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.
AML	<i>(EMASS/GRAU specific term)</i> Automated Mixed-Media library.
application agent	A component needed on a client to back up or restore online database integrations. See also Disk Agent .
application system	<i>(ZDB specific term)</i> A system the application or database runs on. The application or database data is located on source volumes. See also backup system and source volume .
archived redo log	<i>(Oracle specific term)</i> 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">• 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.

- NOARCHIVELOG - The filled online redo log files are not archived.

See also [online redo log](#).

archive logging	<i>(Lotus Domino Server specific term)</i> Lotus Domino Server database mode where transaction log files are overwritten only after they have been backed up.
ASR Set	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 <code>Data_Protector_home\Config\Server\dr\asr</code> on a Windows Cell Manager or in <code>/etc/opt/omni/server/dr/asr/</code> on a UNIX Cell Manager) as well as on the backup medium. The ASR archive file is extracted to three diskettes for 32-bit Windows systems or four diskettes for 64-bit Windows systems after a disaster occurs. You need these diskettes to perform ASR.
Audit Logs	Data files to which auditing information is stored.
Audit Report	User-readable output of auditing information created from data stored in audit log files.
Auditing Information	Data about every backup session that was performed over an extended, user-defined period for the whole Data Protector cell.
autochanger	See library .
autoloader	See library .
Automatic Storage Management	<i>(Oracle specific term)</i> Automatic Storage Management is an Oracle 10g integrated filesystem and volume manager that manages Oracle database files. It eliminates complexity associated with managing data and disk and provides striping and mirroring capabilities to optimize performance.
automigration	<i>(VLS specific term)</i> 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.

See also [Virtual Library System \(VLS\)](#) and [virtual tape](#).

BACKINT	<p>(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.</p>
backup API	<p>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.</p>
backup chain	<p>See restore chain.</p>
backup device	<p>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.</p>
backup generation	<p>One backup generation includes one full backup and all incremental backups until the next full backup.</p>
backup ID	<p>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.</p>
backup object	<p>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). A backup object is defined by:</p> <ul style="list-style-type: none">• Client name: Hostname of the Data Protector client where the backup object resides.• 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.• 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).

- Type: Backup object type. For filesystem objects — filesystem type (for example, WinFS). For integration objects — “Bar”.

backup owner	Each backup object in the IDB has an owner. The default owner of a backup is the user who starts the backup session.
backup session	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 (full or incremental). The result of a backup session is a set of media, which was written to, also called the backup or media set. See also backup specification , incremental backup , and full backup .
backup set	A complete set of integration objects associated with a backup.
backup set	<i>(Oracle specific term)</i> 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.
backup specification	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.
backup system	<i>(ZDB specific term)</i> A system connected to target volumes of one or multiple application systems. The backup system is typically connected to a backup device to perform the backup of the data in a replica. See also application system , target volume , and replica .
backup types	See incremental backup , differential backup , transaction backup , full backup , and delta backup .

backup to IAP	A Data Protector based backup to the HP Integrated Archiving Platform (IAP) appliance. It takes advantage of the IAP capability to eliminate redundancies in the stored data at a block (or chunk) level, by creating a unique content address for each data chunk. Only changed chunks are transmitted over the network and added to the store.
backup view	Data Protector provides different views for backup specifications: By Type - according to the type of data available for backups/templates. Default view. By Group - according to the group to which backup specifications/templates belong. By Name - according to the name of backup specifications/templates. 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.
BC	<i>(EMC Symmetrix specific term)</i> Business Continuance are processes that allow customers to access and manage instant copies of EMC Symmetrix standard devices. See also BCV .
BC	<i>(HP StorageWorks Disk Array XP specific term)</i> The Business Copy XP allows to maintain internal copies of HP StorageWorks Disk Array XP LDEVs for purposes such as data backup or data duplication. The copies (secondary volumes or S-VOLs) can be separated from the primary volumes (P-VOLs) and connected to a different system for various purposes, such as backup and development. For backup purposes, P-VOLs should be connected to the application system, and one of the S-VOL mirror sets should be connected to the backup system. See also HP StorageWorks Disk Array XP LDEV, CA, Main Control Unit, application system, and backup system .
BC EVA	<i>(HP StorageWorks EVA specific term)</i> Business Copy EVA is a local replication software solution enabling you to create point-in-time copies (replicas) of the source volumes using the snapshot and clone capabilities of the EVA firmware. See also replica , source volume , snapshot , and CA+BC EVA .
BC Process	<i>(EMC Symmetrix specific term)</i> 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.
See also [BCV](#).

BC VA *(HP StorageWorks Virtual Array specific term)* Business Copy VA allows you to maintain internal copies of HP StorageWorks Virtual Array LUNs for data backup or data duplication within the same virtual array. The copies (child or Business Copy LUNs) can be used for various purposes, such as backup, data analysis or development. When used for backup purposes, the original (parent) LUNs are connected to the application system and the Business Copy (child) LUNs are connected to the backup system. See also [HP StorageWorks Virtual Array LUN](#), [application system](#), and [backup system](#).

BCV *(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. See also [BC](#) and [BC Process](#).

Boolean operators 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.

boot volume/disk/partition 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.

BRARCHIVE *(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. See also [BRBACKUP](#), and [BRRESTORE](#).

BRBACKUP *(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.

See also [BRARCHIVE](#), and [BRRESTORE](#).

BRRESTORE

(SAP R/3 specific term) An SAP R/3 tool that can be used to restore files of the following type:

- Database data files, control files, and online redo log files saved with BRBACKUP
- Redo log files archived with BRARCHIVE
- Non-database files saved with BRBACKUP

You can specify files, tablespaces, complete backups, log sequence numbers of redo log files, or the session ID of the backup.

See also [BRBACKUP](#), and [BRARCHIVE](#).

BSM

The Data Protector Backup Session Manager controls the backup session. This process always runs on the Cell Manager system.

CA

(HP StorageWorks Disk Array XP specific term) Continuous Access XP allows you to create and maintain remote copies of HP StorageWorks Disk Array XP LDEVs for purposes such as data duplication, backup, and disaster recovery. CA operations involve the main (primary) disk arrays and the remote (secondary) disk arrays. The main disk arrays contain the CA primary volumes (P-VOLs), which contain the original data and are connected to the application system. The remote disk arrays contain the CA secondary volumes (S-VOLs) connected to the backup system.

See also [BC](#) *(HP StorageWorks Disk Array XP specific term)*, [Main Control Unit](#) and [HP StorageWorks Disk Array XP LDEV](#).

CA+BC EVA

(HP StorageWorks EVA specific term) The combination of Continuous Access (CA) EVA and Business Copy (BC) EVA enables you to create and maintain copies (replicas) of the source volumes on a remote EVA, and then use these copies as the source for local replication on this remote array.

See also [BC EVA](#), [replica](#), and [source volume](#).

CAP

(StorageTek specific term) Cartridge Access Port is a port built into the door panel of a library. The purpose is to enter or eject media.

catalog protection	Defines how long information about backed up data (such as file names and file versions) is kept in the IDB. See also data protection .
CDB	The Catalog Database is a part of the IDB that contains information about backups, object copies, restores, media management sessions,, and backed up data. 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. See also MMDB .
CDF file	<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.
cell	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. Central control is available to administer the backup and restore policies and tasks.
Cell Manager	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.
centralized licensing	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. See also MoM .
Centralized Media Management Database (CMMDB)	See CMMDB .

Change Journal	<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.
Change Log Provider	<i>(Windows specific term)</i> A module that can be queried to determine which objects on a filesystem have been created, modified, or deleted.
channel	<p><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:</p> <ul style="list-style-type: none"> • type 'disk' • type 'sbt_tape' <p>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.</p>
chunking	<p><i>(IAP specific term)</i> The process of dividing data into blocks (chunks), where each chunk gets a unique content address. This address is then used to determine whether a particular chunk is already backed up to the IAP appliance. If the duplicate data is identified (two addresses are identical, that is the address is the same as for another data chunk already stored into IAP), it is not backed up. This way, the data redundancy is eliminated and the optimal data storage is achieved.</p> <p>See also backup to IAP.</p>
circular logging	<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.
client backup	A backup of all filesystems mounted on a client. Filesystems mounted on the client after the backup specification was created are not automatically detected.
client backup with disk discovery	A backup of all filesystems mounted on a client. When the backup starts, Data Protector discovers the disks on the clients. Client backup with disk discovery simplifies backup configuration

and improves backup coverage of systems that often mount or dismount disks.

client or **client system** Any system configured with any Data Protector functionality and configured in a cell.

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).

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](#).

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 MoM Manager. A reliable network connection between the MoM cell and the other Data Protector cells is highly recommended. See also [MoM](#).

COM+ Registration Database	<i>(Windows specific term)</i> The COM+ Registration Database and the Windows Registry store COM+ application attributes, class attributes, and computer-level attributes. This guarantees consistency among these attributes and provides common operation on top of these attributes.
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.
Command View (CV) EVA	<i>(HP StorageWorks EVA specific term)</i> The user interface that enables you to configure, manage, and monitor your HP StorageWorks EVA storage system. It is used to perform various storage management tasks, for example, creating virtual disk families, managing storage system hardware, and creating snapclones and snapshots of virtual disks. The Command View EVA software runs on the HP OpenView Storage Management Appliance, and is accessed by a Web browser. See also HP StorageWorks EVA SMI-S Agent and HP StorageWorks SMI-S EVA provider .
Command View VLS	<i>(VLS specific term)</i> A web browser-based GUI that is used to configure, manage, and monitor the VLS through a LAN. See also Virtual Library System (VLS) .
concurrency	See Disk Agent concurrency .
control file	<i>(Oracle and SAP R/3 specific term)</i> An Oracle data file that contains entries specifying the physical structure of the database. It provides database consistency information used for recovery.
copy set	<i>(HP StorageWorks EVA specific term)</i> A pair that consists of the source volumes on a local EVA and their replica on a remote EVA. See also source volume , replica , and CA+BC 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 <code>root</code> .

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.
data file	<i>(Oracle and SAP R/3 specific term)</i> 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. <i>See also catalog protection.</i>
data stream	Sequence of data transferred over the communication channel.
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.
Data Replication (DR) group	<i>(HP StorageWorks EVA specific term)</i> A logical grouping of EVA virtual disks. It can contain up to eight copy sets provided they have common characteristics and share a common CA EVA log. <i>See also copy set.</i>
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.
Dobject	<i>(Informix Server specific term)</i> 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 <code>dcbf</code> directory and is located in the <code>Data_Protector_home\db40</code> directory on a Windows Cell Manager and in the <code>/var/opt/omni/server/db40</code> directory on a UNIX Cell

Manager. You can create more DC directories and use a custom location. Up to 10 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 file system 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	<i>(EMC Symmetrix specific term)</i> 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.
device streaming	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.
DHCP server	A system running the Dynamic Host Configuration Protocol (DHCP) providing dynamic IP address assignment and network configuration for DHCP clients.

differential backup	An incremental backup that backs up changes made since the last full backup. To perform this type of backup, specify the <code>Incr1</code> backup type. See also incremental backup .
differential backup	<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. See also backup types .
differential database backup	A differential database backup records only those data changes made to the database after the last full database backup.
direct backup	A SAN-based backup solution in which data movement directly from disk to tape (or to other secondary storage) is facilitated by the SCSI Extended Copy (Xcopy) command. Direct backup lessens the backup I/O load on systems in a SAN environment. The data movement is facilitated directly from disk to tape (or to other secondary storage) by the SCSI Extended Copy (XCOPY) command. The command is provided by any element of the infrastructure including bridges, switches, tape libraries, and disk subsystems. See also XCOPY engine .
directory junction	<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.
disaster recovery	A process to restore a client's main system disk to a state close to the time when a (full) backup was performed.
Disk Agent	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.
Disk Agent concurrency	The number of Disk Agents that are allowed to send data to one Media Agent concurrently.
disk discovery	The detection of disks during client backup with disk discovery. During this backup, Data Protector discovers (detects) the disks that are present on the client — even though they might not have

been present on the system when the backup was configured — and backs them up. This is particularly useful in dynamic environments, where configurations change rapidly. After the disks are expanded, each inherits all options from its master client object. Even if pre- and post-exec commands are specified once, they are started many times, once per each object.

disk group	<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.
disk image (rawdisk) backup	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.
disk quota	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.
disk staging	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).
distributed file media format	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. See also virtual full backup .
Distributed File System (DFS)	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.
DMZ	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.

DNS server	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.
domain controller	A server in a network that is responsible for user security and verifying passwords within a group of other servers.
DR image	Data required for temporary disaster recovery operating system (DR OS) installation and configuration.
DR OS	A disaster recovery operating system is an operating system environment in which disaster recovery runs. It provides Data Protector a basic runtime environment (disk, network, tape, and filesystem access). The OS has to be installed and configured before the Data Protector disaster recovery can be performed. DR OS not only hosts the Data Protector disaster recovery process but is also a part of the restored system because it replaces its own configuration data with the original configuration data.
drive	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.
drive index	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.
dynamic client	See client backup with disk discovery .
EMC Symmetrix Agent (SYMA) (EMC Symmetrix specific term)	See Symmetrix Agent (SYMA) .
emergency boot file	<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>SERVERNUM</code> configuration parameter. Each line of the emergency boot file corresponds to one backup object.

enhanced incremental backup	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.
Enterprise Backup Environment	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. <i>See also MoM.</i>
Event Log (Data Protector Event Log)	A central repository of all Data Protector related notifications. By default, all notifications are sent to the Event Log. The Event Log is accessible only to Data Protector users in the <code>Admin</code> group and to Data Protector users who are granted the <code>Reporting and notifications</code> user rights. You can view or delete all events in the Event Log.
Event Logs	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.
Exchange Replication Service	<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. <i>See also Cluster Continuous Replication and Local Continuous Replication.</i>
exchanger	Also referred to as SCSI Exchanger. <i>See also library.</i>
exporting media	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. <i>See also importing media.</i>

Extensible Storage Engine (ESE)	<i>(Microsoft Exchange Server specific term)</i> A database technology used as a storage system for information exchange in Microsoft Exchange Server.
failover	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.
failover	<i>(HP StorageWorks EVA specific term)</i> An operation that reverses the roles of source and destination in CA+BC EVA configurations. See also CA+BC EVA .
FC bridge	See Fibre Channel bridge .
Fibre Channel	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.
Fibre Channel bridge	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.
file depot	A file containing the data from a backup to a file library device.
file jukebox device	A device residing on disk consisting of multiple slots used to store file media.
file library device	A device which resides on a disk emulating a library with several media, hence containing multiple files, referred to as file depots.
File Replication Service (FRS)	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.

file tree walk	<i>(Windows specific term)</i> The process of traversing a filesystem to determine which objects have been created, modified, or deleted.
file version	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.
filesystem	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.
first-level mirror	<i>(HP StorageWorks Disk Array XP specific term)</i> HP StorageWorks Disk Array XP allows up to three mirror copies of a primary volume and each of these copies can have additional two copies. The three mirror copies are called first-level mirrors. See also primary volume and MU number .
flash recovery area	<i>(Oracle specific term)</i> Flash recovery area is an Oracle 10g managed directory, filesystem, or Automatic Storage Management disk group that serves as a centralized storage area for files related to backup and recovery (recovery files). See also recovery files .
fnames.dat	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.
formatting	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.
free pool	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.
full backup	A backup in which all selected objects are backed up, whether or not they have been recently modified. See also backup types .

full database backup	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.
full mailbox backup	A full mailbox backup is a backup of the entire mailbox content.
full ZDB	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. See also incremental ZDB .
global options file	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 in the <code>/etc/opt/omni/server/options</code> directory on HP-UX and Solaris systems and in the <code>Data_Protector_home\Config\Server\Options</code> directory on Windows systems.
group	<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.
GUI	A cross-platform (HP-UX, Solaris, Linux, and Windows) graphical user interface, provided by Data Protector for easy access to all configuration, administration, and operation tasks. Besides the original Data Protector GUI, Data Protector also provides a Java-based graphical user interface with the same look and feel. As Java can run on numerous platforms, the Data Protector Java GUI is supported on a larger number of platforms than the original Data Protector GUI.
hard recovery	<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.
heartbeat	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.

Hierarchical Storage Management (HSM)	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.
Holidays file	A file that contains information about holidays. You can set different holidays by editing the Holidays file: <code>/etc/opt/omni/server/Holidays</code> on the UNIX Cell Manager and <code>Data_Protector_home\Config\Server\holidays</code> on the Windows Cell Manager.
host backup	See client backup with disk discovery .
hosting system	A working Data Protector client used for Disk Delivery Disaster Recovery with a Data Protector Disk Agent installed.
HP ITO	See OM .
HP OpC	See OM .
HP OpenView SMART Plug-In (SPI)	A fully integrated, out-of-the-box solution which "plugs into" HP Operations Manager software, extending the managed domain. Through the Data Protector integration, which is implemented as an HP OpenView SMART Plug-In, a user can have an arbitrary number of Data Protector Cell Managers monitored as an extension to HP Operations Manager software (OM).
HP OM	See OM
HP StorageWorks Disk Array XP LDEV	A logical partition of a physical disk within an HP StorageWorks Disk Array XP. LDEVs are entities that can be replicated in the Continuous Access XP (CA) and Business Copy XP (BC) configurations, or can be used as standalone entities. See also BC , CA (<i>HP StorageWorks Disk Array XP specific term</i>), and replica .
HP StorageWorks EVA SMI-S Agent	A Data Protector software module that executes all tasks required for the HP StorageWorks Enterprise Virtual Array integration. With the EVA SMI-S Agent, the control over the array is established through HP StorageWorks SMI-S EVA provider, which directs communication between incoming requests and CV EVA.

See also [Command View \(CV\) EVA](#) and [HP StorageWorks SMI-S EVA provider](#).

HP StorageWorks SMI-S EVA provider

An interface used for controlling HP StorageWorks Enterprise Virtual Array. SMI-S EVA provider runs as a separate service on the HP OpenView Storage Management Appliance system and acts as a gateway between incoming requests and Command View EVA. With the Data Protector HP StorageWorks EVA integration, SMI-S EVA provider accepts standardized requests from the EVA SMI-S Agent, communicates with Command View EVA for information or method invocation, and returns standardized responses.

See also [HP StorageWorks EVA SMI-S Agent](#) and [Command View \(CV\) EVA](#).

HP StorageWorks Virtual Array LUN

A logical partition of a physical disk within an HP StorageWorks Virtual Array. LUNs are entities that can be replicated in the HP StorageWorks Business Copy VA configuration, or can be used as standalone entities.

See also [BC VA](#) and [replica](#).

HP VPO

See [OM](#).

ICDA

(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.

IDB

The Data Protector Internal Database is an embedded database located on the Cell Manager that keeps information regarding which data is backed up, on which media it is backed up, how backup and restore sessions are run, and which devices and libraries are configured.

IDB recovery file

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.

importing media	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. <i>See also exporting media.</i>
incremental backup	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. <i>See also backup types.</i>
incremental backup	<i>(Microsoft Exchange Server specific term)</i> 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. <i>See also backup types.</i>
incremental mailbox backup	An incremental mailbox backup backs up all the changes made to the mailbox after the last backup of any type.
incremental1 mailbox backup	An incremental1 mailbox backup backs up all the changes made to the mailbox after the last full backup.
incremental (re)-establish	<i>(EMC Symmetrix specific term)</i> 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.
incremental restore	<i>(EMC Symmetrix specific term)</i> 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.

- incremental ZDB** 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.
See also [full ZDB](#).
- Inet** 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.
- Information Store** (*Microsoft Exchange Server specific term*) 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.
See also [Key Management Service](#) and [Site Replication Service](#).
- Informix Server** (*Informix Server specific term*) Refers to Informix Dynamic Server.
- initializing** See [formatting](#).
- Installation Server** A computer system that holds a repository of the Data Protector software 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.
- instant recovery** (*ZDB specific term*) 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.
See also [replica](#), [zero downtime backup \(ZDB\)](#), [ZDB to disk](#), and [ZDB to disk+tape](#).

integration object	A backup object of a Data Protector integration, such as Oracle or SAP DB.
Internet Information Services (IIS)	<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).
IP address	An Internet Protocol address is a numeric address of a system used to uniquely identify the system on the network. The IP address consists of four groups of numbers separated by periods (full stops).
ISQL	<i>(Sybase specific term)</i> A Sybase utility used to perform system administration tasks on Sybase SQL Server.
ITO	See OM .
Java GUI Client	The Java GUI Client is a component of the Java GUI that contains only user interface related functionalities and requires connection to the Java GUI Server to function.
Java GUI Server	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.
jukebox	See library .
jukebox device	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”.
Key Management Service	<i>(Microsoft Exchange Server specific term)</i> The Microsoft Exchange Server service that provides encryption functionality for enhanced security. See also Information Store and Site Replication Service .
keychain	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.

LBO *(EMC Symmetrix specific term)* 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.

library 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.

lights-out operation or **unattended operation** 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.

LISTENER.ORA *(Oracle specific term)* An Oracle configuration file that describes one or more Transparent Network Substrate (TNS) listeners on a server.

load balancing 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.

local and remote recovery 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.

Local Continuous Replication	<p>(<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. 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. 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. See also Cluster Continuous Replication and Exchange Replication Service.</p>
lock name	<p>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.</p>
log_full shell script	<p>(<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 <code>ALARMPROGRAM</code> 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 <code>ALARMPROGRAM</code> configuration parameter to <code>INFORMIXDIR/etc/no_log.sh</code>.</p>
logging level	<p>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</p>

influence the IDB growth, backup speed, and the convenience of browsing data for restore.

logical-log files

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.

login ID

(Microsoft SQL Server specific term) 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.

login information to the Oracle Target Database

(Oracle and SAP R/3 specific term) The format of the login information is *user_name/password@service*, where:

- *user_name* 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.
- *password* must be the same as the password specified in the Oracle password file (*orapwd*), which is used for authentication of users performing database administration.
- *service* is the name used to identify an SQL*Net server process for the target database.

login information to the Recovery Catalog Database

(Oracle specific term) The format of the login information to the Recovery (Oracle) Catalog Database is *user_name/password@service*, 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, *service* is the name of the service to the Recovery Catalog Database, not the Oracle target database. Note that the Oracle user specified here must be the owner of the Oracle Recovery Catalog.

Lotus C API

(Lotus Domino Server specific term) An interface for the exchange of backup and recovery information between Lotus Domino Server and a backup solution, like Data Protector.

LVM

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.

Magic Packet	See Wake ONLAN .
mailbox	<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.
mailbox store	<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.
Main Control Unit (MCU)	<i>(HP StorageWorks Disk Array XP specific term)</i> An HP StorageWorks XP disk array that contains the primary volumes for the CA and BC configurations and acts as a master device. See also BC (HP StorageWorks Disk Array XP specific term), CA (HP StorageWorks Disk Array XP specific term), and HP StorageWorks Disk Array XP LDEV .
Manager-of-Managers (MoM)	See MoM .
MAPI	<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.
MCU	See Main Control Unit (MCU) .
Media Agent	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 session, a Media Agent locates data on the backup medium and sends it to the Disk Agent. The Disk Agent then writes the data to the disk. A Media Agent also manages the robotics control of a library.
media allocation policy	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.
media condition	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.
media condition factors	The user-assigned age threshold and overwrite threshold used to determine the state of a medium.
medium ID	A unique identifier assigned to a medium by Data Protector.
media label	A user-defined identifier used to describe a medium.
media location	A user-defined physical location of a medium, such as "building 4" or "off-site storage".
media management session	A session performing some action on a medium, such as initializing, scanning the content, verifying data on a medium, or copying a medium.
media pool	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.
media set	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.
media type	The physical type of media, such as DDS or DLT.
media usage policy	The media usage policy controls how new backups are added to the already used media. It can be <code>Appendable</code> , <code>Non-Appendable</code> , or <code>Appendable for incrementals only</code> .
merging	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. <i>See also overwrite.</i>

Microsoft Exchange Server	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.
Microsoft Management Console (MMC)	<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.
Microsoft SQL Server	A database management system designed to meet the requirements of distributed "client-server" computing.
Microsoft Volume Shadow Copy Service (VSS)	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. See also shadow copy , shadow copy provider , replica , and writer .
mirror (EMC Symmetrix and HP StorageWorks Disk Array XP specific term)	See target volume .
mirror rotation (HP StorageWorks Disk Array XP specific term)	See replica set rotation .
MMD	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.
MMDB	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.

See also [CMMDB](#), [CDB](#).

- MoM** 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.
- mount request** 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.
- mount point** 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.
- MSM** The Data Protector Media Session Manager, which runs on the Cell Manager and controls media sessions, such as copying media.
- MU number** (*HP StorageWorks Disk Array XP specific term*) Mirror Unit number. An integer number (0, 1 or 2), used to indicate a first-level mirror.
See also [first-level mirror](#).
- multi-drive server** A license that allows you to run an unlimited number of Media Agents on a single system. This license, which is bound to the IP address of the Cell Manager, is no longer available.
- obdrindex.dat** See [IDB recovery file](#).
- OBDR capable device** 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.
- object** See [backup object](#).
- object consolidation** 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.

object consolidation session	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.
object copy	A copy of a specific object version that is created during an object copy session or a backup session with object mirroring.
object copy session	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.
object copying	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.
object ID	<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.
object mirror	A copy of a backup object created using object mirroring. Object mirrors are often referred to as object copies.
object mirroring	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.
offline backup	<p>A backup during which an application database cannot be used by the application.</p> <ul style="list-style-type: none">• For simple backup methods (non ZDB), the database is generally put into a quiescent state that allows use by the backup system, but not the application, for the whole backup period (several minutes or hours). For instance, for backup to tape, until streaming of data to the tape is finished.• For ZDB methods, the database is also put into the quiescent state, but for the period of the data replication process only (several seconds). Normal database operation can then be resumed for the rest of the backup process. <p>See also zero downtime backup (ZDB) and online backup.</p>

offline recovery	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.
offline redo log	See archived redo log .
On-Bar	<p><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:</p> <ul style="list-style-type: none"> • the <code>onbar</code> command • Data Protector as the backup solution • the XBSA interface • ON-Bar catalog tables, which are used to back up dbobjects and track instances of dbobjects through multiple backups.
ONCONFIG	<p><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 <code>onconfig</code> file in the directory <code>INFORMIXDIR\etc</code> (on Windows) or <code>INFORMIXDIR/etc/</code> (on UNIX).</p>
online backup	<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 that the backup application requires access to the original data objects. During this period, the database is fully operational, but there may be a small performance impact and log files may grow very quickly.</p> <ul style="list-style-type: none"> • For simple backup methods (non ZDB), backup mode is required for the whole backup period (several minutes or hours). For instance, for backup to tape, until streaming of data to tape is finished. • For ZDB methods, backup mode is required for the short period of the data replication process only (several seconds). Normal database operation can then be resumed for the rest of the backup process. <p>In some cases, transaction logs may also have to be backed up to allow a consistent database to be restored. See also zero downtime backup (ZDB), and offline backup.</p>

online redo log	<i>(Oracle specific term)</i> 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. See also archived redo log .
OpC	See OM .
OpenSSH	A set of network connectivity tools used to access remote machines securely, by using a variety of authentication and encryption methods. It needs to be installed and configured on the Installation Server and the client if you perform remote installation using secure shell.
Oracle Data Guard	<i>(Oracle specific term)</i> 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.
Oracle instance	<i>(Oracle specific term)</i> Each installation of an Oracle database on one or more systems. One computer system can have several instances of a database running.
ORACLE_SID	<i>(Oracle specific term)</i> 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.
original system	The system configuration backed up by Data Protector before a computer disaster hits the system.
overwrite	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. See also merging .

OM	<p>HP Operations Manager software for UNIX provides powerful capabilities for operations management of a large number of systems and applications on in a network. Data Protector provides an integration into this management product. This integration is implemented as a SMART Plug-In for OM management servers on HP-UX, Solaris, and Linux. Earlier versions of OM were called IT/Operation, Operations Center and Vantage Point Operations.</p> <p>See also merging.</p>
ownership	<p>The ownership of a backup determines who can restore from the backup. The user who starts an interactive backup is the session owner. If a user starts an existing backup specification without modifying it, the session is not considered interactive. In that case, if the backup owner has been defined in the backup specification, they remain the session owner. Otherwise, the session owner becomes the user who started the backup in question. For the scheduled backups, by default, the session owner is for the UNIX Cell Manager: <code>root.sys@Cell Manager</code>, and for the Windows Cell Manager, the user that was specified during the installation of the Cell Manager. It is possible to modify the ownership, so that the specific user becomes the session owner.</p>
P1S file	<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 backup medium and on Cell Manager into <code>Data_Protector_home\Config\Server\dr\pls</code> directory on a Windows Cell Manager or in <code>/etc/opt/omni/server/dr/pls</code> directory on a UNIX Cell Manager with the filename <code>recovery.pls</code>.</p>
package	<p><i>(MC/ServiceGuard and Veritas Cluster specific term)</i> A collection of resources (for example volume groups, application services, IP names and addresses) that are needed to run a specific cluster-aware application.</p>
pair status	<p><i>(HP StorageWorks Disk Array XP specific term)</i> A mirrored pair of disks can have various status values depending on the action performed on it. The three most important status values are:</p>

- COPY - The mirrored pair is currently re-synchronizing. Data is transferred from one disk to the other. The disks do not contain the same data.
- PAIR - The mirrored pair is completely synchronized and both disks (the primary volume and the mirrored volume) contain identical data.
- SUSPENDED - The link between the mirrored disks is suspended. That means that both disks are accessed and updated independently. However, the mirror relationship is still maintained and the pair can be re-synchronized without transferring the complete disk.

parallel restore	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.
parallelism	The concept of reading multiple data streams from an online database.
physical device	A physical unit that contains either a drive or a more complex unit such as a library.
post-exec	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. See also pre-exec .
pre- and post-exec commands	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.
prealloc list	A subset of media in a media pool that specifies the order in which media are used for backup.

pre-exec	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. See also post-exec .
primary volume (P-VOL)	<i>(HP StorageWorks Disk Array XP specific term)</i> Standard HP StorageWorks Disk Array XP LDEVs that act as a primary volume for the CA and BC configurations. The P-VOL is located in the MCU. See also secondary volume (S-VOL) and Main Control Unit (MCU) .
protection	See data protection and also catalog protection .
public folder store	<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 <code>.edb</code> file and a streaming native internet content <code>.stm</code> file.
public/private backed up data	When configuring a backup, you can select whether the backed up data will be: <ul style="list-style-type: none"> • public, that is visible (and accessible for restore) to all Data Protector users • private, that is, visible (and accessible for restore) only to the owner of the backup and administrators
RAID	Redundant Array of Inexpensive Disks.
RAID Manager Library	<i>(HP StorageWorks Disk Array XP specific term)</i> The RAID Manager Library is used internally by Data Protector on Solaris systems to allow access to HP StorageWorks Disk Array XP configuration, status, and performance data and to key HP StorageWorks Disk Array XP features through the use of function calls translated into a sequence of low level SCSI commands.
RAID Manager XP	<i>(HP StorageWorks Disk Array XP specific term)</i> The RAID Manager XP application provides an extensive list of commands to report and control the status of the CA and BC applications. The commands communicate through a RAID Manager instance with the HP StorageWorks Disk Array XP Disk Control Unit. This

instance translates the commands into a sequence of low level SCSI commands.

rawdisk backup	See disk image backup .
RCU	See Remote Control Unit (RCU) .
RDBMS	Relational Database Management System.
RDF1/RDF2	<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.
RDS	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.
Recovery Catalog	<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">• The physical schema of the Oracle target database• Data file and archived log backup sets• Data file copies• Archived Redo Logs• Stored scripts
Recovery Catalog Database	<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.
recovery files	<i>(Oracle specific term)</i> Recovery files are Oracle 10g 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. See also flash recovery area .
RecoveryInfo	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.

Recovery Manager (RMAN)	<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.
recycle	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.
redo log	<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.
Remote Control Unit (RCU)	<i>(HP StorageWorks Disk Array XP specific term)</i> The Remote Control Unit (RCU) acts as a slave of an MCU in a CA configuration. In bidirectional configurations, the RCU can act as an MCU.
Removable Storage Management Database	<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.
reparse point	<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.
replica	<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.

See also [snapshot](#), [snapshot creation](#), [split mirror](#), and [split mirror creation](#).

- replica set** *(ZDB specific term)* A group of replicas, all created using the same backup specification.
See also [replica](#) and [replica set rotation](#).
- replica set rotation** *(ZDB specific term)* 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.
See also [replica](#) and [replica set](#).
- restore chain** 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.
- restore session** A process that copies data from backup media to a client.
- resync mode** *(HP StorageWorks Disk Array XP VSS provider specific term)* One of two XP VSS hardware provider operation modes. When the XP 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.
See also [VSS compliant mode](#), [source volume](#), [primary volume \(P-VOL\)](#), [replica](#), [secondary volume \(S-VOL\)](#), [MU number](#), and [replica set rotation](#).
- RMAN (Oracle specific term)** See [Recovery Manager](#).
- RSM** The Data Protector Restore Session Manager controls the restore session. This process always runs on the Cell Manager system.

RSM	<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.
scan	A function that 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).
scanning	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.
Scheduler	A function that controls when and how often automatic backups occur. By setting up a schedule, you automate the start of backups.
secondary volume (S-VOL)	<i>(HP StorageWorks Disk Array XP specific term)</i> secondary volumes, or S-VOLs, are XP LDEVs that act as a secondary CA or BC mirror of another LDEV (a P-VOL). In the case of CA, S-VOLs can be used as failover devices in a MetroCluster configuration. The S-VOLs are assigned separate SCSI addresses, different from the addresses used by the P-VOLs. See also primary volume (P-VOL) and Main Control Unit (MCU)
session	See backup session , media management session , and restore session .
session ID	An identifier of a backup, restore, object copy, object consolidation, or media management session, consisting of the date when the session ran and a unique number.
session key	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 <code>omnimnt</code> , <code>omnistat</code> , and <code>omniabort</code> commands.

shadow copy	<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. See also Microsoft Volume Shadow Copy Service and replica .
shadow copy provider	<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). See also shadow copy .
shadow copy set	<i>(Microsoft VSS specific term)</i> A collection of shadow copies created at the same point in time. See also shadow copy and replica set .
shared disks	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.
SIBF	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.
single instancing	<i>(IAP specific term)</i> The process of recognizing redundancy of data, at both a whole object and a chunk level. It computes a strong hash for each data chunk and uses it as a unique content address needed to determine whether attempts to store duplicates are being made. See also backup to IAP .
Site Replication Service	<i>(Microsoft Exchange Server specific term)</i> The Microsoft Exchange Server 2000/2003 service that permits compatibility with Microsoft Exchange Server 5.5 by emulating the Exchange Server 5.5 directory service. See also Information Store and Key Management Service .
slot	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.

- SMB** See [split mirror backup](#).
- smart copy** (*VLS specific term*) 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. See also [Virtual Library System \(VLS\)](#).
- smart copy pool** (*VLS specific term*) A pool that defines which destination library slots are available as smart copy targets for a specified source virtual library. See also [Virtual Library System \(VLS\)](#) and [smart copy](#).
- SMBF** The Session Messages Binary Files (SMBF) part of the IDB stores session messages generated during backup, restore, object copy, object consolidation, and media management sessions. One binary file is created per session. The files are grouped by year and month.
- snapshot** (*HP StorageWorks VA and HP StorageWorks EVA specific term*) A form of replica produced using snapshot creation techniques. A range of snapshot types is available, with different characteristics, depending on the arrays/techniques used. Such replicas are dynamic and may be either virtual copies, still reliant upon the contents of the source volumes, or independent exact duplicates (clones), depending on the snapshot type and the time since creation. See also [replica](#) and [snapshot creation](#).
- snapshot backup**
(*HP StorageWorks VA and HP StorageWorks EVA specific term*) See [ZDB to tape](#), [ZDB to disk](#), and [ZDB to disk+tape](#).
- snapshot creation** (*HP StorageWorks VA and HP StorageWorks EVA specific term*) A replica creation technique, in which copies of source volumes are created using storage virtualization techniques. The replicas are considered to be created at one particular point in time, without pre-configuration, and are immediately available for

use. However background copying processes normally continue after creation.

See also [snapshot](#).

- source (R1) device** *(EMC Symmetrix specific term)* 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.
See also [target \(R2\) device](#).
- source volume** *(ZDB specific term)* A storage volume containing data to be replicated.
- sparse file** 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.
- split mirror** *(EMC Symmetrix and HP StorageWorks Disk Array XP specific term)* A replica created using split mirror techniques. Such a replica provides an independent, exact duplicate, or clone of the contents of the source volumes.
See also [replica](#) and [split mirror creation](#).
- split mirror backup** *(EMC Symmetrix specific term)* See [ZDB to tape](#).
- split mirror backup** *(HP StorageWorks Disk Array XP specific term)* See [ZDB to tape](#), [ZDB to disk](#), and [ZDB to disk+tape](#).
- split mirror creation** *(EMC Symmetrix and HP StorageWorks Disk Array XP specific term)* 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.
See also [split mirror](#).

- split mirror restore** *(EMC Symmetrix and HP StorageWorks Disk Array XP specific term)* A process in which data backed up in a ZDB-to-tape or a ZDB-to-disk+tape session is restored from tape media to a split mirror replica, which is then synchronized to the source volumes. Individual backup objects or complete sessions can be restored using this method.
See also [ZDB to tape](#), [ZDB to disk+tape](#), and [replica](#).
- sqlhosts file** *(Informix Server specific term)* 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.
- SRD file** The Data Protector System Recovery Data (SRD) file contains system information required for installing and configuring the operating system in case of a disaster. The SRD file is an ASCII file, generated when a CONFIGURATION backup is performed on a Windows client and stored on the Cell Manager.
- SRDF** *(EMC Symmetrix specific term)* 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.
- SSE Agent** *(HP StorageWorks Disk Array XP specific term)* A Data Protector software module that executes all tasks required for a split mirror backup integration. It communicates with the HP StorageWorks Disk Array XP storing system using the RAID Manager XP utility (HP-UX and Windows systems) or RAID Manager Library (Solaris systems).
- sst.conf file** The file `/usr/kernel/drv/sst.conf` 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.
- st.conf file** The file `/kernel/drv/st.conf` 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.

stackers	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.
standalone file device	A file device is a file in a specified directory to which you back up data.
Storage Group	<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.
StorageTek ACS library	<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.
storage volume	<i>(ZDB specific term)</i> A storage volume represents an object that may be presented to an operating system or some other entity (for example, a virtualization mechanism) upon which volume management systems, file systems, 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.
switchover	See failover .
Sybase Backup Server API	<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.
Sybase SQL Server	<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.
Symmetrix Agent (SYMA)	<i>(EMC Symmetrix specific term)</i> The Data Protector software module that prepares the EMC Symmetrix environment for backup and restore operations.

synthetic backup	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.
synthetic full backup	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.
System Backup to Tape	<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.
system databases	<i>(Sybase specific term)</i> The four system databases on a newly installed Sybase SQL Server are the: <ul style="list-style-type: none"> • master database (master) • temporary database (tempdb) • system procedure database (sybssystemprocs) • model database (model).
System State	<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.
system volume/disk/partition	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.
SysVol	<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.

tablespace	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.
tapeless backup (ZDB specific term)	See ZDB to disk .
target database	<i>(Oracle specific term)</i> In RMAN, the target database is the database that you are backing up or restoring.
target (R2) device	<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. See also source (R1) device .
target system	<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 crashed system and a target system is that a target system has all faulty hardware replaced.
target volume	<i>(ZDB specific term)</i> A storage volume to which data is replicated.
Terminal Services	<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.
thread	<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.
TimeFinder	<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).
TLU	Tape Library Unit.

TNSNAMES.ORA	<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.
transaction	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.
transaction backup	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.
transaction backup	<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.
transaction log backup	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.
transaction log files	Files that record transactions of the database modifications, and provide fault tolerance in case of a database disaster.
transaction logs	<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.
transaction log table	<i>(Sybase specific term)</i> A system table in which all changes to the database are automatically recorded.
transportable snapshot	<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. See also Microsoft Volume Shadow Copy Service (VSS) .
TSANDS.CFG file	<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.

UIProxy	The Java GUI Server (UIProxy 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.
unattended operation	See lights-out operation .
user account (Data Protector user account)	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.
user disk quotas	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.
user group	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.
user profile	<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.
user rights	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.
vaulting media	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.
verify	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.
Virtual Controller Software (VCS)	<i>(HP StorageWorks EVA specific term)</i> The firmware that manages all aspects of storage system operation, including communication with Command View EVA through the HSV controllers. See also Command View (CV) EVA .
Virtual Device Interface	<i>(Microsoft SQL Server specific term)</i> This is a SQL Server programming interface that allows fast backup and restore of large databases.
virtual disk	<i>(HP StorageWorks EVA specific term)</i> A unit of storage allocated from an HP StorageWorks Enterprise Virtual Array storage pool. Virtual disks are the entities that are replicated using the HP StorageWorks Enterprise Virtual Array snapshot functionality. See also source volume and target volume .
virtual full backup	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.
Virtual Library System (VLS)	A disk-based data storage device hosting one or more virtual tape libraries (VTLs).
virtual server	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.
virtual tape	<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. See also Virtual Library System (VLS) and Virtual Tape Library .

Virtual Tape Library (VTL)	<i>(VLS specific term)</i> An emulated tape library that provides the functionality of traditional tape-based storage. See also Virtual Library System (VLS) .
volser	<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.
volume group	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.
volume mount point	<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.
Volume Shadow Copy Service	See Microsoft Volume Shadow Copy Service .
VPO	See OM .
VSS	See Microsoft Volume Shadow Copy Service .
VSS compliant mode	<i>(HP StorageWorks Disk Array XP VSS provider specific term)</i> One of two XP VSS hardware provider operation modes. When the XP 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 of 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.

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	<i>(Microsoft VSS specific term)</i> 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.
XBSA interface	<i>(Informix Server specific term)</i> ON-Bar and Data Protector communicate with each other through the X/Open Backup Services Application Programmer's Interface (XBSA).
XCopy engine	<i>(direct backup specific term)</i> A SCSI-3 copy command that allows you to copy data from a storage device having a SCSI source address to a backup device having a SCSI destination address, thus enabling direct backup. The data flows from a source device (either block or streaming, that is, disk or tape) to the destination device (either block or streaming) through XCopy. This releases

the controlling server of reading the data from the storage device into memory and then writing the information to the destination device.

See also [direct backup](#).

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 for ZDB, instant recovery, and split mirror restore.

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 on split mirror arrays, 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. On split mirror arrays, split mirror restore can also be used.

See also [zero downtime backup \(ZDB\)](#), [ZDB to disk](#), [instant recovery](#), [ZDB to disk+tape](#), 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

- architecture
 - SAP DB integration, 237
 - SAP R/3 integration, 155
- audience, 15

B

- backing up Oracle
 - backup templates, 57
 - offline, 75
 - online, 76
- backing up SAP DB
 - backup flow, 238
 - concepts, scheme, 237
 - parallelism, concepts, 238
 - previewing backups, 250
 - starting backups, 252
- backing up SAP R/3
 - backup options, 188
 - previewing backups, 190
- backing up Oracle, 75 - 82 - 90
 - backup options, 66
 - backup specifications, creating, 58
 - backup types, 27
 - examples, using RMAN, 87
 - recovery catalog, 78
 - scheduling backups, 78
 - starting backups, 80 - 82 - 90
 - starting backups, using CLI, 81
 - starting backups, using GUI, 80
 - starting backups, using RMAN, 82
- backing up SAP DB, 246 - 256
 - architecture, 237
 - backup options, 249
 - backup modes, 246
 - backup specification, modifying, 249
 - backup specifications, creating, 246
 - backup types, 235
 - differential backups, 235
 - full backups, 235
 - online backups, 235
 - parallelism, 249
 - scheduling backups, 249
 - scheduling backups, example, 250
 - transaction log backups, 235
- backing up SAP R/3, 181 - 196
 - backup flow, 158
 - backup specifications, creating, 183
 - backup templates, 184
 - backup types, 181
 - backup modes, 154
 - backup specification, modifying, 189
 - full backups, 153, 181
 - incremental backups, 153, 181
 - SAP backup utilities, 154
 - scheduling backups, 189
 - scheduling backups, example, 189
 - starting backups, 191

- backing up SAP R/3
 - architecture, 155
 - backup types, 153
 - manual balancing, 188, 196
 - SAP compliant ZDB, 193
 - SAP R/3 parameter file, 183
 - using RMAN, 195
 - using RMAN, 182
- backint mode
 - SAP R/3 integration, 154
- backup options
 - Oracle integration, 66
- backup specifications, scheduling
 - Oracle integration, 78
 - SAP DB integration, 249
- backup flow
 - SAP R/3 integration, 158
 - SAP DB integration, 238
- backup flow, Oracle integration, 30 - 31
- backup modes
 - SAP DB integration, 246
- backup modes, SAP R/3 integration, 154
- backup options
 - SAP DB integration, 249
 - SAP R/3 integration, 188
- backup specifications, creating
 - SAP DB integration, 246
- backup specifications, creating
 - Oracle integration, 58
 - SAP R/3 integration, 183
- backup specifications, modifying
 - SAP DB integration, 249
 - SAP R/3 integration, 189
- backup specifications, ownership
 - Oracle integration, 44
- backup templates
 - Oracle integration, 57
 - SAP R/3 integration, 184

- backup types
 - Oracle integration, 27
 - SAP DB integration, 235
 - SAP R/3 integration, 153, 181
- BRARCHIVE
 - SAP R/3 integration, 154
- BRBACKUP
 - SAP R/3 integration, 154
- BRRESTORE, 201
 - SAP R/3 integration, 154

C

- checking configuration
 - Oracle integration, 55
 - SAP DB integration, 245
 - SAP R/3 integration, 179
- concepts
 - Oracle integration, 28
 - SAP DB integration, 236
 - SAP R/3 integration, 154 - 160
- configuration file
 - SAP R/3 integration, 160
- configuring Oracle, 33 - 56
 - checking configuration, 55
 - example, CLI, 54
 - prerequisites, 35
- configuring SAP DB, 239 - 246
 - checking configuration, 245
- configuring SAP R/3, 160 - 181
 - checking configuration, 179
 - configuration file, 160
- configuring SAP R/3
 - authentication modes, 173
- control files, Oracle integration
 - restore, 97
- conventions
 - document, 22
- creating backup specifications
 - Oracle integration, 58
 - SAP DB integration, 246
 - SAP R/3 integration, 183

D

- Data Guard, Oracle integration
 - configuration, example, 54
 - limitations, 35
 - primary databases, restore, 105
 - standby databases, restore, 105
- database recovery
 - Oracle integration, options, 110
- differential backups
 - SAP DB integration, 235
- disaster recovery
 - Oracle integration, 94
 - Oracle integration, 132
 - SAP R/3 integration, 203
- document
 - conventions, 22
 - related documentation, 15
- documentation
 - HP website, 15
 - providing feedback, 25

E

- examples
 - SAP R/3 integration, starting interactive backups, 192
- examples, Oracle integration
 - backing up using RMAN, 87
 - restoring using RMAN, 114
- examples, SAP DB integration
 - scheduling backups, 250
 - starting interactive backups, 253
- examples, scheduling backups
 - SAP R/3 integration, 189

F

- finding users
 - Oracle integration, 44
- full backups
 - SAP DB integration, 235
 - SAP R/3 integration, 153, 181

H

- help
 - obtaining, 24
- HP
 - technical support, 24

I

- incremental backups
 - Oracle integration, 79
 - SAP R/3 integration, 153, 181
- Informix backup
 - backup specifications, creating, 183
- interactive backups
 - SAP R/3 integration, 191
- interactive backups
 - Oracle integration, 80
 - SAP DB integration, 252
- introduction
 - Oracle integration, 27
 - SAP DB integration, 235
 - SAP R/3 integration, 153

L

- limitations
 - SAP DB integration, 239

M

- manual balancing
 - SAP R/3 integration, 188, 196
- MC/ServiceGuard
 - clusters, configuration, 45
 - linking Oracle with the MML, 36
- Media Management Library
 - See MML
- migration, restore
 - SAP DB integration, 260
- MML (Data Protector Media Management Library)
 - linking with Oracle, UNIX, 36

- MML (Data Protector Media Management Library)
 - linking with Oracle, OpenVMS, 38
- modifying backup specifications
 - SAP DB integration, 249
 - SAP R/3 integration, 189
- monitoring sessions
 - Oracle integration, 133
 - SAP DB integration, 274
 - SAP R/3 integration, 204

O

- OB2_MIRROR_COMP, SAP R/3 integration, 193
- OB2RMANSAVE, Oracle integration, 152
- online backups
 - SAP DB integration, 235
- Oracle backup
 - scheduling backups, 78
- Oracle configuration
 - example, CLI, 54
- Oracle integration
 - concepts, 28
 - introduction, 27
 - monitoring sessions, 133
- Oracle restore
 - database items, 91
 - recovery catalog, 130
 - using another device, 131
 - using GUI, 93
- Oracle troubleshooting, 137 - 152
- Oracle backup, 75 - 82 - 90
 - backup concepts, scheme, 32
 - backup specifications, creating, 58
 - backup templates, 57
 - backup types, 27
 - starting backups, using CLI, 81
 - starting backups, using GUI, 80
 - starting backups, using RMAN, 82
 - starting backups, 80 - 82 - 90

- Oracle configuration
 - checking configuration, 55
 - prerequisites, 35
- Oracle integration
 - backup, 75 - 82 - 90
 - configuration, 33 - 56
 - disaster recovery, 132
 - removing the integration, 134
 - restore, 91 - 133
 - troubleshooting, 137 - 152
 - viewing sessions, 134
- Oracle restore, 91 - 133
 - restorable items, 91
 - control files, 97
 - database objects, 99
 - disaster recovery, 132
 - editing RMAN scripts, 152
 - examples, using RMAN, 114
 - preparing databases for restore, 115
 - primary databases, Data Guard, 105
 - recovery catalog, 95
 - restore options, 110
 - restore flow, 31
 - restore methods, 91
 - restore types, 28
 - standby databases, Data Guard, 105
 - tablespaces and datafiles, 104
 - using RMAN, 114
- Oracle RMAN metadata, 136
- Oracle RMAN script, 68
- overview, restore
 - SAP DB integration, 256
- ownership, backup specifications
 - Oracle integration, 44

P

- parallelism
 - SAP DB integration, 249
- parallelism, concepts
 - SAP DB integration, 238, 238

- previewing backups
 - SAP DB integration, [250](#)
 - SAP R/3 integration, [190](#)
- primary databases, Oracle integration
 - restore, [105](#)

R

- RAC, configuring Oracle Servers
 - on HP-UX, [35](#)
 - on other UNIX systems, [35](#)
- recovery
 - Oracle integration, options, [110](#)
- recovery catalog, Oracle integration
 - backup, [78](#)
 - restore, [95](#)
- related documentation, [15](#)
- removing the Oracle integration, [134](#)
 - from HP-UX, [135](#)
 - from Solaris and other UNIX systems, [135](#)
- restore flow
 - SAP R/3 integration, [160](#)
 - SAP DB integration, [238](#)
- restore methods
 - SAP R/3 integration, [196](#)
- restore options
 - SAP DB integration, [270](#)
- restore types
 - Oracle integration, [28](#)
- restoring Oracle
 - database objects, [99](#)
 - tablespaces and datafiles, [104](#)
 - using RMAN, [114](#)
- restoring SAP DB, [256 - 274](#)
 - parallelism, concepts, [238](#)
 - restore flow, [238](#)
 - using another device, [274](#)
 - using CLI, [263](#)
 - using GUI, [260](#)
- restoring SAP R/3
 - using another device, [201](#)
- restoring SAP R/3
 - disaster recovery, [203](#)
 - using SAP BRTOOLS, [201](#)
- restoring Oracle, [91 - 133](#)
 - control files, [97](#)
 - disaster recovery, [132](#)
 - editing RMAN scripts, [152](#)
 - methods, [91](#)
 - primary databases, Data Guard, [105](#)
 - recovery catalog, [95, 130](#)
 - restore flow, [31](#)
 - standby databases, Data Guard, [105](#)
 - using another device, [131](#)
 - using GUI, [93](#)
- restoring SAP R/3
 - using GUI, [197](#)
- restoring SAP DB migration, [260](#)
 - overview, [256](#)
 - restore options, [270](#)
 - using SAP DB utilities, [264](#)
- restoring SAP R/3, [196 - 204](#)
 - restore flow, [160](#)
 - restore methods, [196](#)
 - SAP restore utilities, [154](#)
 - using CLI, [200](#)
- restoring SAP R/3 architecture, [155](#)
 - using BRRESTORE, [201](#)
- RMAN mode
 - SAP R/3 integration, [154](#)
- RMAN, backup
 - SAP R/3 integration, [182, 195](#)
- RMAN, Oracle integration, [82](#)
 - backup, [87](#)
 - restore, [114](#)
 - scripts, examples, [87](#)
- running backups
 - See starting backups

S

- SAP DB backup
 - backup options, 249
 - scheduling backups, 249
- SAP DB configuration, 239 - 246
- SAP DB integration
 - introduction, 235
 - limitations, 239
 - monitoring sessions, 274
- SAP DB restore
 - restore options, 270
- SAP R/3 backup
 - backup flow, 158
 - backup templates, 184
 - backup types, 181
 - manual balancing, 188, 196
 - SAP compliant ZDB, 193
 - scheduling backups, 189
 - starting backups, 191
 - using RMAN, 195
- SAP R/3 configuration
 - authentication modes, 173
- SAP R/3 integration
 - disaster recovery, 203
 - introduction, 153
 - monitoring sessions, 204
- SAP R/3 restore
 - restore flow, 160
 - using BRRESTORE, 201
- SAP R/3 troubleshooting, 204 - 233
- SAP compliant ZDB
 - SAP R/3 integration, 193
- SAP DB backup, 246 - 256
 - concepts, scheme, 237
 - architecture, 237
 - backup flow, 238
 - backup specification, modifying, 249
 - backup specifications, creating, 246
 - backup modes, 246
 - backup types, 235
 - differential backups, 235
 - full backups, 235
 - online backups, 235
 - parallelism, 249
 - parallelism, concepts, 238
 - previewing backups, 250
 - scheduling backups, example, 250
 - starting backups, 252
 - transaction log backups, 235
- SAP DB configuration
 - checking configuration, 245
- SAP DB integration
 - backup, 246 - 256
 - concepts, 236
 - configuration, 239 - 246
 - restore, 256 - 274
 - troubleshooting, 274 - 278
- SAP DB restore, 256 - 274
 - migration, 260
 - overview, 256
 - parallelism, concepts, 238
 - restore flow, 238
 - using another device, 274
 - using CLI, 263
 - using GUI, 260
 - using SAP DB utilities, 264
- SAP DB troubleshooting, 274 - 278
- SAP DB utilities
 - restore, 264
- SAP R/3 backup
 - backup types, 153
- SAP R/3 configuration, 160 - 181
- SAP R/3 restore
 - using CLI, 200

- SAP R/3 backup, [181 - 196](#)
 - architecture, [155](#)
 - backup specification, modifying, [189](#)
 - backup modes, [154](#)
 - backup options, [188](#)
 - full backups, [153, 181](#)
 - incremental backups, [153, 181](#)
 - previewing backups, [190](#)
 - SAP backup utilities, [154](#)
 - SAP R/3 parameter file, [183](#)
 - scheduling backups, example, [189](#)
 - using RMAN, [182](#)
 - SAP R/3 configuration
 - checking configuration, [179](#)
 - configuration file, [160](#)
 - SAP R/3 integration
 - backup, [181 - 196](#)
 - concepts, [154 - 160](#)
 - configuration, [160 - 181](#)
 - restore, [196 - 204](#)
 - troubleshooting, [204 - 233](#)
 - SAP R/3 restore, [196 - 204](#)
 - architecture, [155](#)
 - disaster recovery, [203](#)
 - restore methods, [196](#)
 - SAP restore utilities, [154](#)
 - using another device, [201](#)
 - using BRRESTORE, [201](#)
 - using GUI, [197](#)
 - SAP R/3 troubleshooting
 - on UNIX, [217 - 233](#)
 - on Windows, [205 - 217](#)
 - SBT_LIBRARY, Oracle integration, [36, 85, 117](#)
 - scheduling backups
 - Oracle integration, [78](#)
 - SAP DB integration, [249](#)
 - SAP R/3 integration, [189](#)
 - standby databases, Oracle integration
 - restore, [105](#)
 - starting backups
 - SAP DB integration, [252](#)
 - starting backups
 - SAP R/3 integration, [191](#)
 - starting backups, Oracle integration
 - using CLI, [81](#)
 - using GUI, [80](#)
 - using RMAN, [82](#)
 - starting backups, Oracle integration, [80 - 82 - 90](#)
 - Subscriber's Choice, HP, [25](#)
- ## T
- technical support
 - HP, [24](#)
 - technical support
 - service locator website, [25](#)
 - transaction log backups
 - SAP DB integration, [235](#)
 - troubleshooting Oracle, [137 - 152](#)
 - troubleshooting SAP DB, [274 - 278](#)
 - troubleshooting SAP R/3 , [204 - 233](#)
 - on UNIX, [217 - 233](#)
 - on Windows, [205 - 217](#)
- ## U
- users, finding
 - Oracle integration, [44](#)
 - users, configuring
 - Oracle integration, [44](#)
- ## V
- viewing sessions
 - Oracle integration, [134](#)
- ## W
- websites
 - HP Subscriber's Choice for Business, [25](#)
 - HP , [25](#)
 - product manuals, [15](#)

