HP OpenView Storage Data Protector Integration Guide

for

Oracle
SAP

Manual Edition: July 2006

Manufacturing Part Number: B6960-96008

Release A.06.00

Legal Notices

Contents

1. Integrating Oracle and Data Protector

Contents

Contents

Printing History

Manual 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	Manual Edition	Product
B6960-90109	October 2004	Data Protector Release A.05.50
B6960-96008	July 2006	Data Protector Release A.06.00

Conventions

The following typographical conventions are used in this manual.

Table 2

Convention	Meaning	Example
Italic	Book or manual titles, and manual page names	Refer to the *HP OpenView Storage Data Protector Integration Guide* for more information.
	Provides emphasis	You *must* follow these steps.
	Specifies a variable that you must supply when entering a command	At the prompt type: `rlogin` *`your_name`* where you supply your login name.
Bold	New terms	The Data Protector **Cell Manager** is the main ...
`Computer`	Text and items on the computer screen	The system replies: `Press Enter`
	Command names	Use the `grep` command ...
	File and directory names	`/usr/bin/X11`
	Process names	Check to see if `Data Protector Inet` is running.
	Window/dialog box names	In the `Backup Options` dialog box...
	Text that you must enter	At the prompt, type: `ls -l`
Keycap	Keyboard keys	Press **Return**.

Data Protector provides a cross-platform (Windows and UNIX) graphical user interface. Refer to the online Help for information about the Data Protector graphical user interface.

Figure 1 Data Protector Graphical User Interface

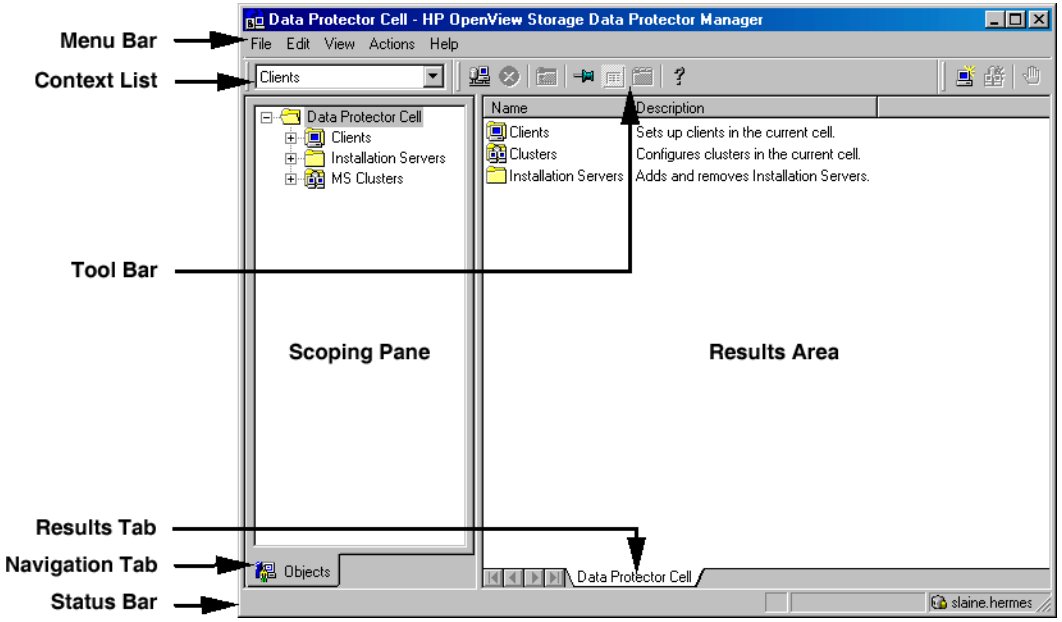

Contact Information

General Information

General information about Data Protector can be found at

http://www.hp.com/go/dataprotector

Technical Support

Technical support information can be found at the HP Electronic Support Centers at

http://www.itrc.hp.com

Information about the latest Data Protector patches can be found at

http://www.itrc.hp.com

HP does not support third-party hardware and software. Contact the respective vendor for support.

Documentation Feedback

Your comments on the documentation help us to understand and meet your needs. You can provide feedback at

storagedocs.feedback@hp.com

Training Information

For information on currently available HP OpenView training, see the HP OpenView World Wide Web site at

http://www.openview.hp.com/training/

Follow the links to obtain information about scheduled classes, training at customer sites, and class registration.

Data Protector Documentation

Data Protector documentation comes in the form of manuals and online Help.

Manuals

Data Protector manuals are available in printed format and in PDF format. Install the PDF files during the Data Protector setup procedure by selecting the `User Interface` component on Windows or the `OB2-DOCS` component on UNIX. Once installed, the manuals reside in the *<Data_Protector_home>*`\docs` directory on Windows and in the `/opt/omni/doc/C/` directory on UNIX. You can also find the manuals in PDF format at http://www.hp.com/support/manuals

HP OpenView Storage Data Protector Concepts Guide

This manual 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 OpenView Storage Data Protector Installation and Licensing Guide

This manual describes how to install the Data Protector software, taking into account the operating system and architecture of your environment. This manual also gives details on how to upgrade Data Protector, as well as how to obtain the proper licenses for your environment.

HP OpenView Storage Data Protector Troubleshooting Guide

This manual describes how to troubleshoot problems you may encounter when using Data Protector.

HP OpenView Storage Data Protector Disaster Recovery Guide

This manual describes how to plan, prepare for, test and perform a disaster recovery.

HP OpenView Storage Data Protector Integration Guide

This manual describes how to configure and use Data Protector to back up and restore various databases and applications. It is intended for backup administrators or operators. There are four versions of this manual:

- *HP OpenView Storage Data Protector Integration Guide for Microsoft Applications: SQL Server, Exchange Server, and Volume Shadow Copy Service*

  This manual describes the integrations of Data Protector with the following Microsoft applications: Microsoft Exchange Server 2000/2003, Microsoft SQL Server 7/2000/2005, and Volume Shadow Copy Service.

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

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

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

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

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

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

HP OpenView Storage Data Protector Integration Guide for HP OpenView

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

HP OpenView Storage Data Protector Integration Guide for HP OpenView Operations for UNIX

This manual describes how to monitor and manage the health and performance of the Data Protector environment with HP OpenView Operations (OVO), HP OpenView Service Navigator, and HP OpenView Performance (OVP) on UNIX.

HP OpenView Storage Data Protector Integration Guide for HP OpenView Operations for Windows

This manual describes how to monitor and manage the health and performance of the Data Protector environment with HP OpenView Operations (OVO), HP OpenView Service Navigator, and HP OpenView Performance (OVP) on Windows.

There are two versions of the manual:

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

HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide

This manual 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 OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide* and the *HP OpenView Storage Data Protector Zero Downtime Backup Integration Guide*.

HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide

This manual 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 OpenView Storage Data Protector Zero Downtime Backup Integration Guide

This manual 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 2000/2003, and Microsoft

SQL Server 2000 databases. The manual also describes how to configure and use Data Protector to perform backup and restore using the Microsoft Volume Shadow Copy Service.

HP OpenView Storage Data Protector MPE/iX System User Guide

This manual describes how to configure MPE/iX clients and how to back up and restore MPE/iX data.

HP OpenView Storage Data Protector Media Operations User's Guide

This manual 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 OpenView Storage Data Protector Product Announcements, Software Notes, and References

This manual gives a description of new features of HP OpenView Storage Data Protector A.06.00. 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:

- OVO UNIX integration
- OVO 7.1x/7.2x Windows integration
- OVO 7.5 Windows integration
- Media Operations

Online Help

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

Documentation Map

Abbreviations

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

Abbreviation	Manual
CLI	Command Line Interface Reference Guide
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-OV	Integration Guide—HP OpenView Service Information Portal/OpenView Reporter
IG-OVOU	Integration Guide—HP OpenView Operations, UNIX
IG-OVOW	Integration Guide—HP OpenView Operations 7.1x, 7.2x, Windows
IG-OVOW	Integration Guide—HP OpenView Operations 7.5, Windows
IG-Var	Integration Guide—Sybase, Network Node Manager, NDMP and VMware
Install	Installation and Licensing Guide
MO GS	Media Operations Getting Started Guide
MO RN	Media Operations Product Announcements, Software Notes, and References
MO UG	Media Operations User Guide
MPE/iX	MPE/iX System User Guide

Abbreviation	Manual
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.

								Integration Guides							ZDB			MO				
	Help	GS	Concepts	Install	Trouble	DR	PA	MS	O/S	IBM	Var	OV	OVOU	OVOW	Concpt	Admin	IG	GS	User	PA	MPE/iX	CLI
Backup	X	X	X					X	X	X	X				X	X	X				X	
CLI																						X
Concepts/Techniques	X		X					X	X	X	X	X	X	X	X	X	X				X	
Disaster Recovery	X		X			X																
Installation/Upgrade	X	X		X			X					X	X	X				X	X		X	
Instant Recovery	X		X												X	X	X					
Licensing	X			X			X												X			
Limitations	X				X		X	X	X	X	X			X			X			X		
New features	X						X													X		
Planning strategy	X		X									X			X							
Procedures/Tasks	X			X	X	X		X	X	X	X	X	X	X		X	X		X			
Recommendations			X				X								X					X		
Requirements				X			X	X	X	X	X			X				X	X	X		
Restore	X	X	X					X	X	X	X					X	X				X	
Support matrices							X															
Supported configurations															X							
Troubleshooting	X			X	X			X	X	X	X	X				X	X					

Integrations

Look in these manuals for details of the following integrations:

Integration	Guide
HP OpenView Operations (OVO)	IG-OVOU, IG-OVOW
HP OpenView Reporter (OVR)	IG-OV
HP OpenView Reporter Light	IG-OVOW
HP OpenView Service Information Portal (OVSIP)	IG-OV
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 Servers	IG-MS, ZDB IG
Microsoft Exchange Single Mailbox	IG-MS
Microsoft SQL Servers	IG-MS, ZDB IG
Microsoft Volume Shadow Copy Service (VSS)	IG-MS, ZDB IG
NDMP Server	IG-Var
Network Node Manager (NNM)	IG-Var
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

In This Book

This guide describes how to configure and use Data Protector with Oracle and SAP applications.

Audience

It 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 OpenView Storage Data Protector Concepts Guide,* which is recommended to fully understand the fundamentals and the model of Data Protector.

Organization

The manual is organized as follows:

Chapter 1	“Integrating Oracle and Data Protector” on page 1.
Chapter 2	“Integrating SAP R/3 and Data Protector” on page 143.
Chapter 3	“Integrating SAP DB/MaxDB and Data Protector” on page 249.
Glossary	Definition of terms used in this manual.

1 Integrating Oracle and Data Protector

In This Chapter

This chapter explains how to configure and use the Data Protector Oracle integration.

The chapter is organized into the following sections:

Introduction

Data Protector offers offline as well as online backup of the Oracle Server instances. To enable 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 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 Concept

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 1-1 on page 9, and the related terms are explained in the following table.

Figure 1-1 Data Protector Oracle Integration Concept

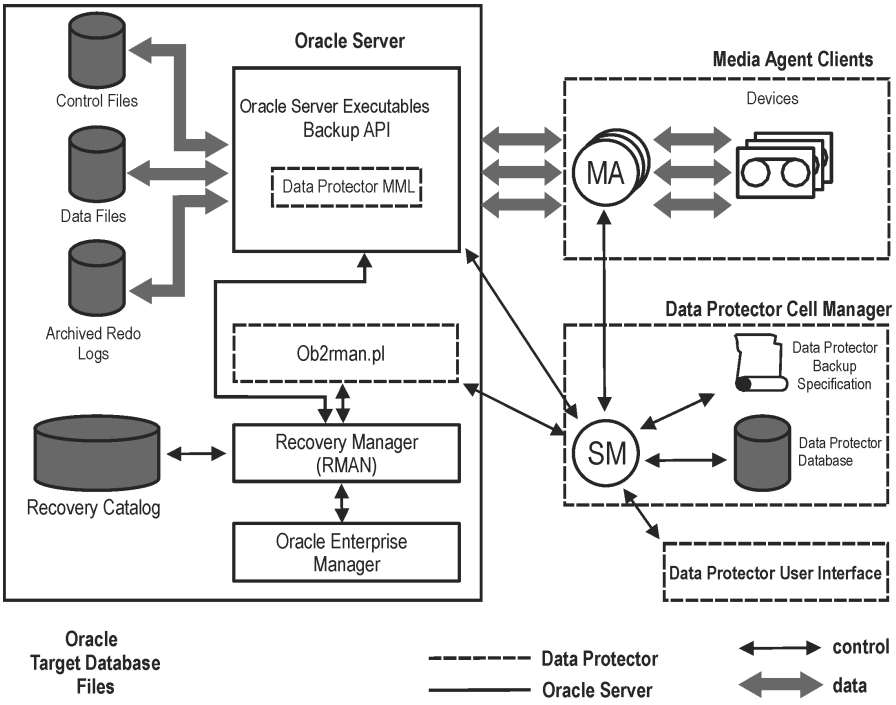

Oracle 10g database files can also be part of ASM configuration. They can reside in the flash recovery area.

Legend:

SM The Data Protector Session Manager, which can be the Data Protector Backup Session Manager during a backup session and the Data Protector Restore Session Manager during a restore session.

RMAN The Oracle Recovery Manager.

Data Protector MML The Data Protector Oracle integration Media Management Library, which is a set of routines that enables data transfer between the Oracle Server and Data Protector.

Backup API The Oracle-defined application programming interface.

IDB	The IDB where all the information about Data Protector sessions, including session messages, objects, data, used devices, and media is written.
MA	The Data Protector General Media Agent, which reads and writes data from and to media devices.

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 OpenView Storage 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 OpenView Storage 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 OpenView Storage 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 116 for details about how to check the prerequisites listed above.

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

- On Windows, 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 "Backup of Archive Logs on RAC Cannot be Performed" on page 132.

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

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 with the Data Protector Oracle Integration Media Management Library (MML) on UNIX

To use the Data Protector Oracle integration, you need to manually link the Oracle server software and MML on the Data Protector Oracle Server system.

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

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 UNIX and OpenVMS Systems” on page 112 for information on removing the integration link.

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

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

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

Other UNIX: `/usr/omni/lib`

The filename for MML depends on the platform:

Table 1-1 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 Oracle with MML 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

```
!rdbsm 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$$"mape
xtra$$"
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 24), 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 24).

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>
```

Figure 1-2 Finding the Oracle User

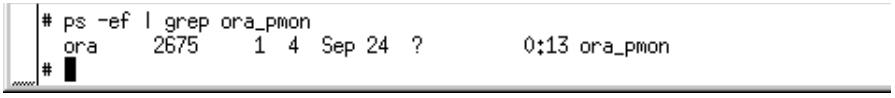

The example above states that the user `ora` has sufficient privileges within the Oracle database to back up and restore the database. Therefore, this user must be added to the corresponding Data Protector user group (`admin` or `operator`) and must also become the owner of the backup specification to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

IMPORTANT Additionally, the user `root` (UNIX only) 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 (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 authorize
      UAF>show onmiadmin
      UAF>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]ONMI_INFO`, for example:

   ```
   -key oracle8 -desc "Oracle Integration" -nlsset 159 -nlsId
   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 “Creating a Data Protector Oracle Backup Specification” on page 37 and at step 5 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 1-3 Configuring Oracle - General (Windows)

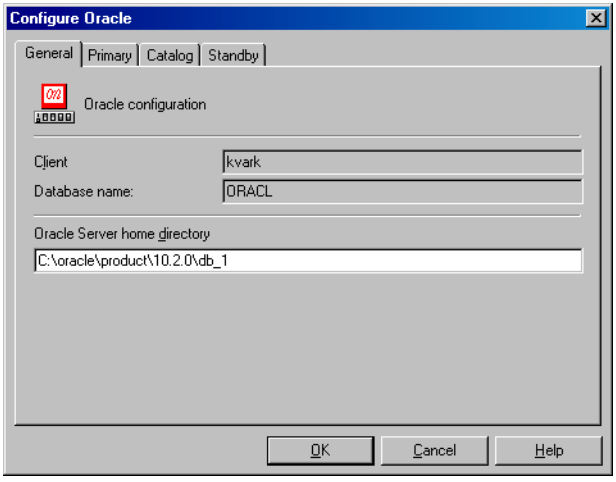

Figure 1-4 **Configuring Oracle - General (UNIX)**

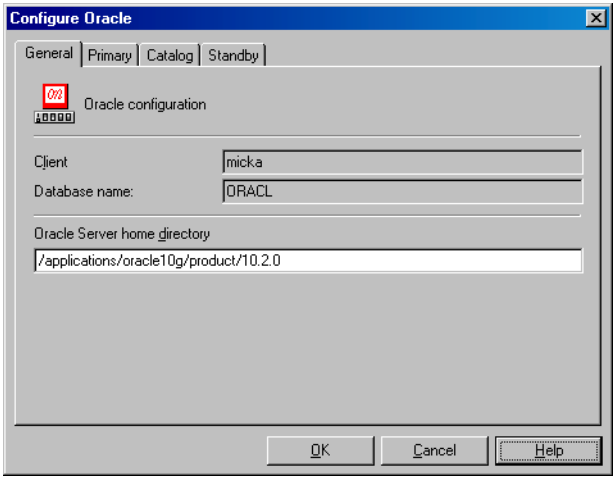

2. In the `Primary` page, specify the login information to the primary database.

   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.

   Note that each database instance for which the net service name is provided must be linked with the Data Protector MML. This means that the Data Protector Oracle integration software component must be installed on each system where the specific database instance is running.

   RAC: List all net services names for the primary database separated by a comma. Each net service name must resolve into a specific database instance.

NOTE You cannot specify a net service name that uses Oracle Net to distribute RMAN connections to more than one instance. In any RMAN connection made through a net service, each net service must specify only one instance.

Figure 1-5 Configuring Oracle - Primary

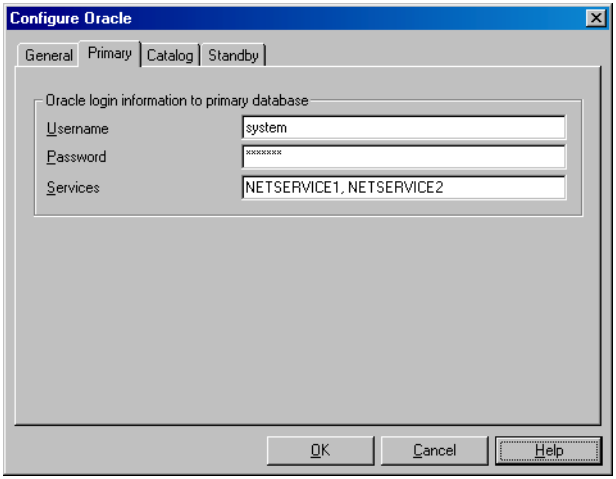

3. In the `Catalog` page, select `Use target database control file instead of recovery catalog` to use the primary database control file.

   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 1-6 **Configuring Oracle - Catalog**

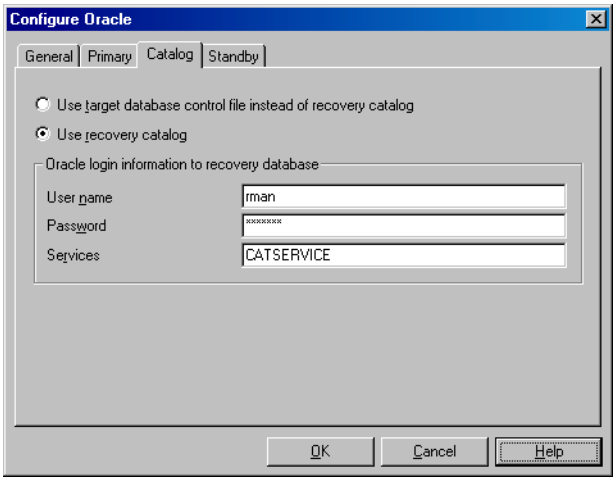

4. ***Oracle Data Guard***: If you intend to back up a standby database, configure also the standby database:

   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.

   Note that each database instance for which the net service name is provided must be linked with the Data Protector MML. This means that the Data Protector Oracle integration software component must be installed on each system where the specific database instance is running.

   RAC: List all net services names for the standby database separated by a comma. Each net service name must resolve into a specific database instance.

NOTE You cannot specify a net service name that uses Oracle Net to distribute RMAN connections to more than one instance. In any RMAN connection made through a net service, each net service must specify only one instance.

Figure 1-7 Configuring Oracle - Standby

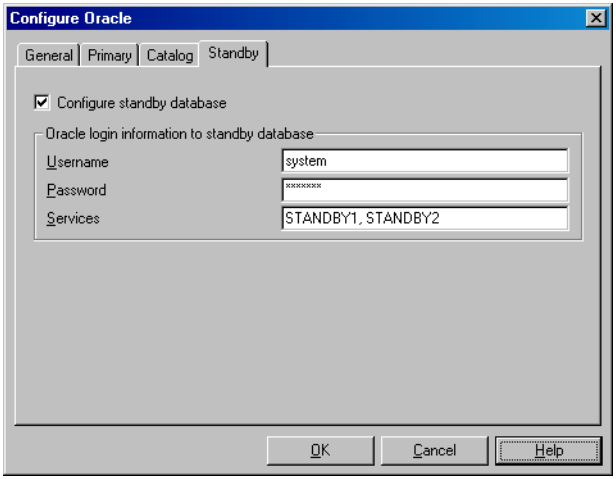

5. Click OK.

The Oracle database is configured. Exit the GUI or proceed with creating the backup specification at step 6 on page 1-40.

Using the Data Protector CLI

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

1. ***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 21.
2. On the Oracle Server system, from the directory:

   Windows: *<Data_Protector_home>*\bin

   HP-UX and Solaris: /opt/omni/lbin

   Other UNIX: /usr/omni/bin/

   OpenVMS: OMNI$ROOT:[BIN]

   run:

On 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>]
```

On 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 catalog 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: `netservice1`

primary net service name 2: `netservice2`

recovery catalog user name: `rman`

recovery catalog password: `manager`

recovery catalog net service name: `catservice`

standby user name: `system`

standby password: `manager`

standby net service name 1: `netservicesb1`

standby net service name 2: `netservicesb2`

Syntax

```
/opt/omni/lbin/util_oracle8.pl -config -dbname oracl \
-orahome /app10g/oracle10g/product/10.1.0 -prmuser system \
-prmpasswd manager -prmservice netservice1,netservice2 \
rcuser rman -rcpasswd manager -rcservice catservice \
-stbuser system -stbpasswd manager -stbservice \
netservicesb1,netservicesb2 -zdb_method BACKUP_SET -pfile \
/app10g/oracle10g/product/10.1.0/dbs/pfile.ora
```

If you need to export some variables before starting SQL*Plus, TNS 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 returns a successful result, 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 21.

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 and Solaris: /opt/omni/lbin

   Other UNIX: /usr/omni/bin/

   OpenVMS: OMNI$ROOT:[BIN]

   run:

On Windows:

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

On 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:

Windows: On the Cell Manager, see the file `<Data_Protector_home>\help\enu\Trouble.txt`

HP-UX and Solaris: Run:

```
/opt/omni/lbin/omnigetmsg 12 <error_number>
```

Other UNIX: Run:

```
/usr/omni/bin/omnigetmsg 12 <error_number>
```

OpenVMS: Run:

```
$@OMNI$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 21.

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 returns a successful result, you may still receive the error `12:8300` when trying to start a backup session. Such a backup session will not start. For more information, see “Troubleshooting” on page 116.

Configuring an Oracle 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 a Data Protector Oracle Backup Specification” on page 37.

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.

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 `fscmd` 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.

Creating a New Template

You can use backup templates to apply the 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 a Data Protector Oracle Backup Specification” on page 37 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 a Data Protector Oracle Backup Specification

To create an Oracle backup specification, proceed as follows:

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.

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

   The database name can be obtained as follows:

   ```
   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 21 for information on how to identify that user.

Figure 1-8 Specifying an Oracle Server System (Windows)

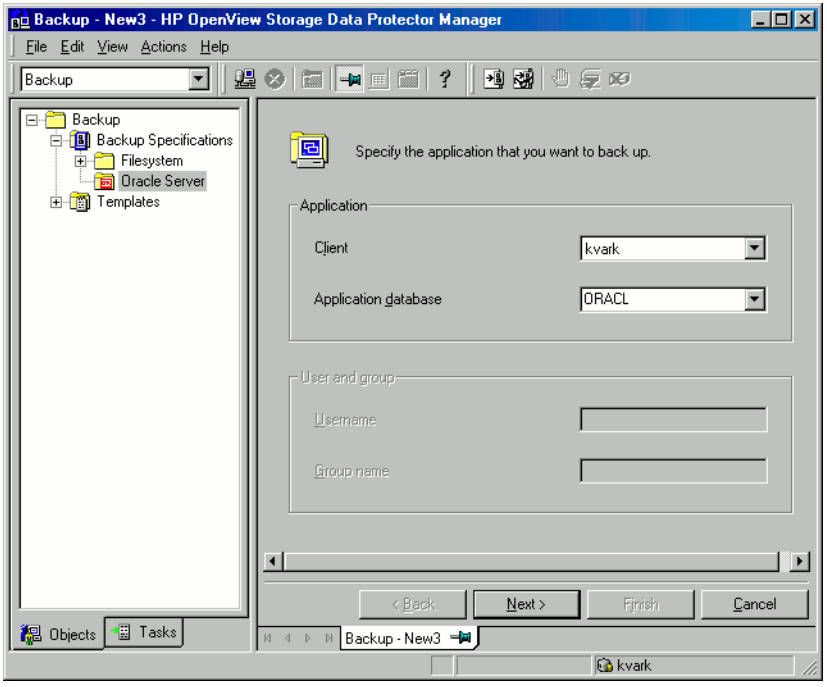

Figure 1-9 Specifying an Oracle Server System (UNIX)

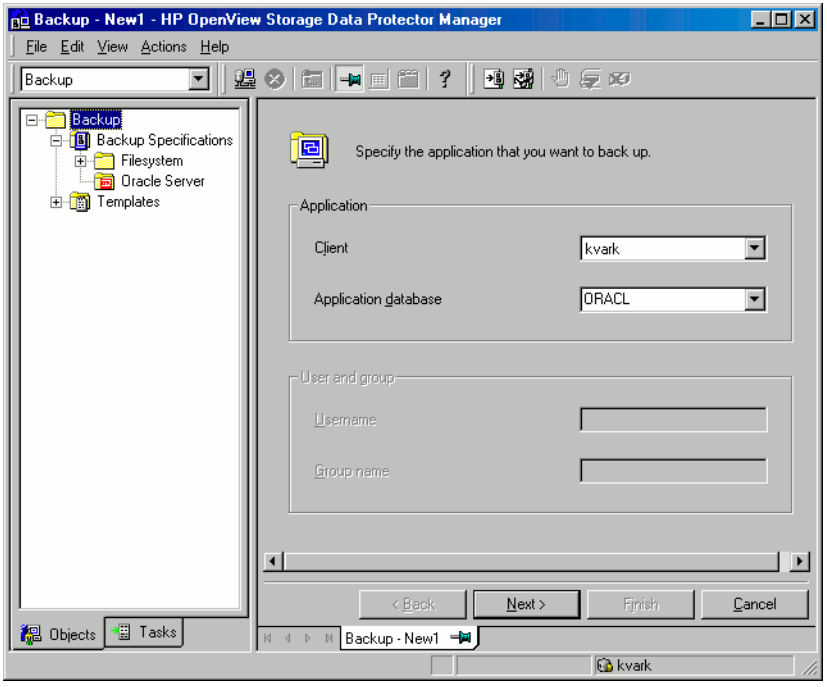

Click `Next`.

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 24.

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.

Figure 1-10 Selecting Backup Objects

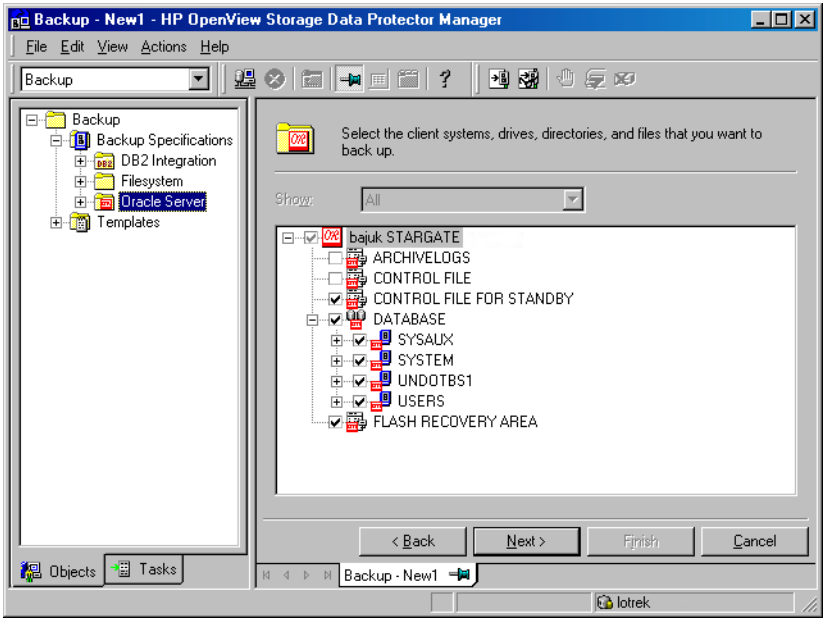

Click `Next`.

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 the `Backup Specification Options` and `Common Application Options`, see the online Help.

   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 1-11), see Table 1-2 on page 44 or online Help.

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 1-11 Oracle Specific Options

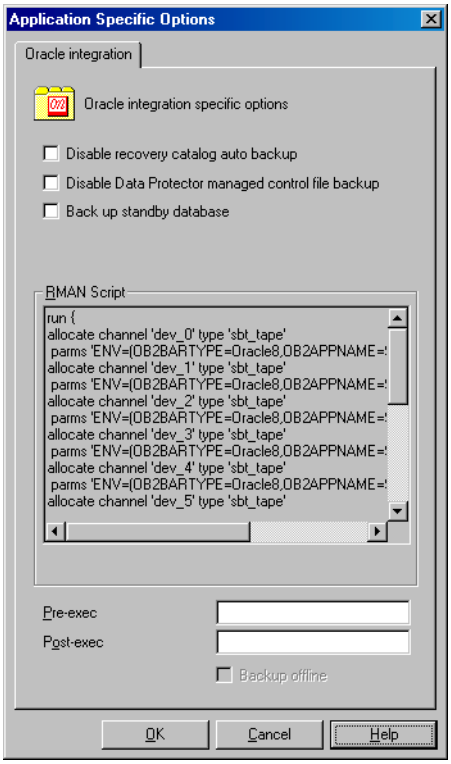

Click `Next`.

9. Optionally, schedule the backup. For more details, see “Scheduling a Backup” on page 58.

   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. Oracle does not allow full stops in backup piece. 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 1-12 Saving the Backup Specification

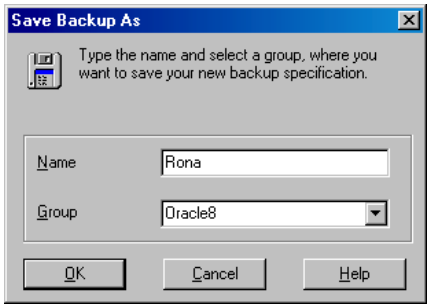

Click OK.

To start the backup, see “Backing Up an Oracle Database” on page 55.

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 21 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 52 for details.

Table 1-2 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.

Table 1-2 Oracle Backup Options

`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 47.
`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 UNIX, see "Examples of Pre-Exec and Post-Exec Scripts on UNIX" on page 46. 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 1-11 on page 43), 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_Name><DB_NAME>_%s:%t:%p>.dbf'
  ```

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

  — 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 support 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 1-13 Previewing a Backup

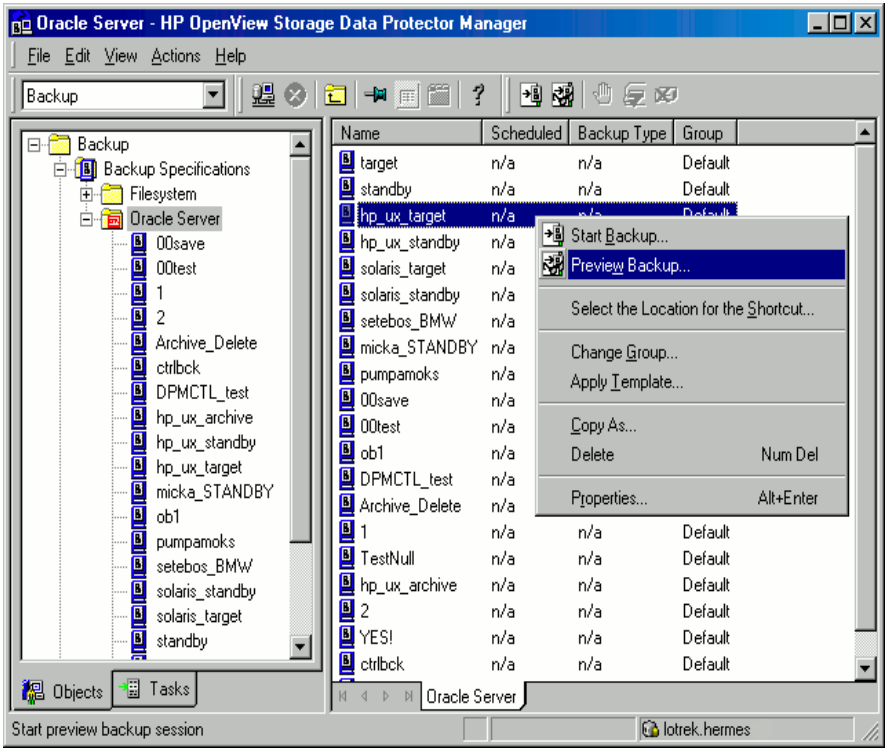

Testing Using the CLI

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 and Solaris: `/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_l qist <backup_specification_name> -test_bar`

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

Backing Up an Oracle Database

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 a 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 a Backup” on page 58.
- 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 60.

- Start a backup on the Oracle server using either Oracle Recovery Manager or Oracle Enterprise Manager. See “Starting Oracle Backup Using RMAN” on page 62.

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 recovery of the database 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 a Backup

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 a 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 `HP OpenView Storage Data Protector Manager`, switch to the `Backup` context.
2. In the Scoping Pane, expand `Backup`, `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 1-14 on page 60. See the RMAN documentation for details on incremental backup levels.

Figure 1-14 Scheduling Backups

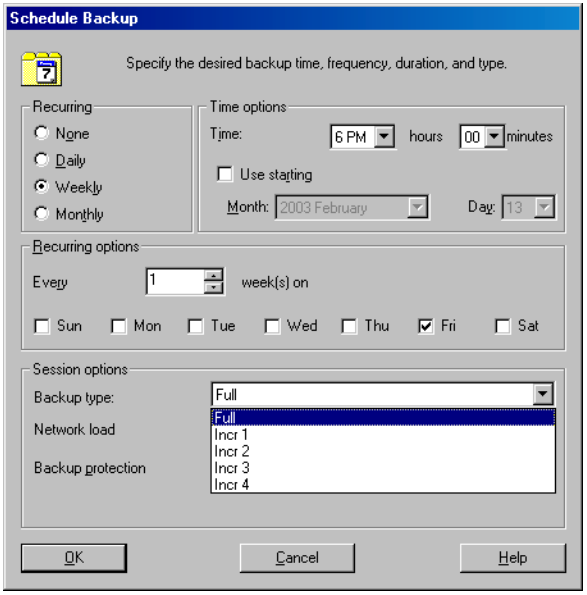

Click `OK` and then `Apply` to save the changes.

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 `HP OpenView Storage Data Protector Manager`, switch to the `Backup` context.
2. In the Scoping Pane, expand `Backup`, `Backup Specifications`, and then `Oracle Server`.
3. Right-click the backup specification and select `Start Backup`.

   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 1-14 on page 60. See the RMAN documentation for details on incremental backup levels.

Click OK.

Figure 1-15 Starting an Interactive Backup

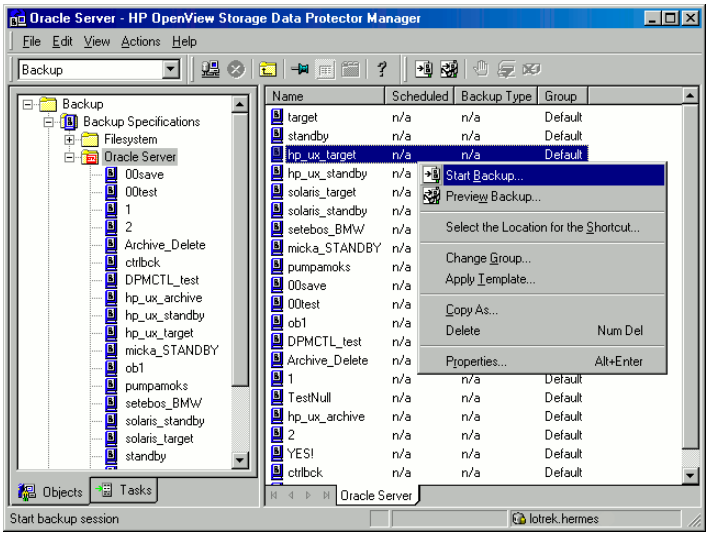

Starting a Backup Using the CLI

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

   Windows: `<Data_Protector_home>\bin`

   HP-UX and Solaris: `/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 "Configuring an Oracle Backup" on page 36 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:

  a. Run `ORAUSER.COM` using `$@OMNI$ROOT:[LOG]LOGIN.COM`.

  b. 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. Note that if Data Protector is linked with Oracle, Data Protector will perform the backup to the tape in the second case.

   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.

   To use Data Protector backup media, specify the channel type `SBT_TAPE`.

3. Specify the `parms` operand:

   ```
   parms'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.

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:

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'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'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';
allocate channel 'dev_1' type 'sbt_tape'
parms'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';
allocate channel 'dev_2' type 'sbt_tape'
parms'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'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';
allocate channel 'dev_1' type 'sbt_tape'
parms'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';
```

```
allocate channel 'dev_2' type 'sbt_tape'
parms'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'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';

allocate channel 'dev_1' type 'sbt_tape'
parms'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';

allocate channel 'dev_2' type 'sbt_tape'
parms'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'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';

allocate channel 'dev_1' type 'sbt_tape'
parms'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';

allocate channel 'dev_2' type 'sbt_tape'
```

```
parms'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'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';
allocate channel 'dev_1' type 'sbt_tape
'parms'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';
allocate channel 'dev_2' type 'sbt_tape'
parms'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'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';
allocate channel 'dev_1' type 'sbt_tape'
```

```
parms'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';

allocate channel 'dev_2' type 'sbt_tape'
parms'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'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';

allocate channel 'dev_1' type 'sbt_tape'
parms'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';

allocate channel 'dev_2' type 'sbt_tape'
parms'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';

backup

incremental level 0

format 'ora1<ORACL_%s:%t>.dbf'

database;

}
```

Restoring an Oracle Database

You can restore the database objects using:

- Data Protector GUI. See “Restoring Oracle Using the Data Protector GUI” on page 72.
- RMAN. See “Restoring Oracle Using RMAN” on page 93.

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 85.

MS 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 1-16.

Figure 1-16 Taking the Oracle Resource Group Offline

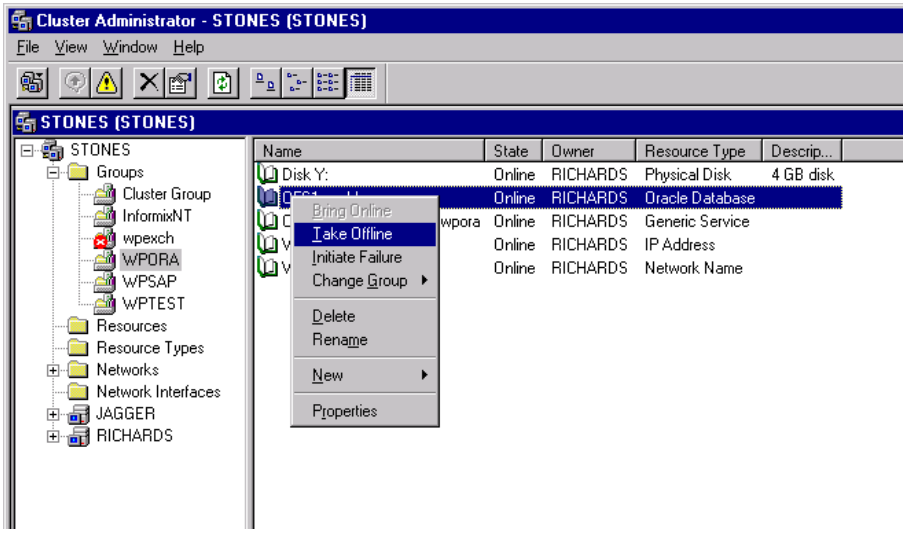

Verify that you have set the Prevent Failback option for the Oracle resource group and Do not restart for the *<DB_NAME>*.world resource, which is an Oracle Database resource.

Figure 1-17 Checking Properties

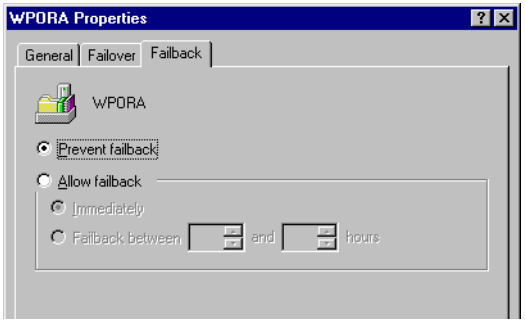

MC/ServiceGuard Clients

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'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 `Mount` state if the whole database is being restored, or in `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. If you want to perform additional actions, you cannot edit the RMAN restore script, but you can perform 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 139.

- 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 139.

  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 1-3 Required Database States

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

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

To put the database into the correct state, run:

```
sqlplus /nolog
SQL>connect <user>/<password>@<service> as sysdba;
SQL>shutdown immediate;
```

To put the database into NoMount state, run:

```
SQL>startup nomount;
```

To put the database into Mount state, run:

```
SQL>startup mount;
```

Restoring the Recovery Catalog Database

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

To restore the recovery catalog database:

1. Ensure that the recovery catalog database is in the `Open` state.
2. 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 1-18 Recovery Catalog Settings Dialog

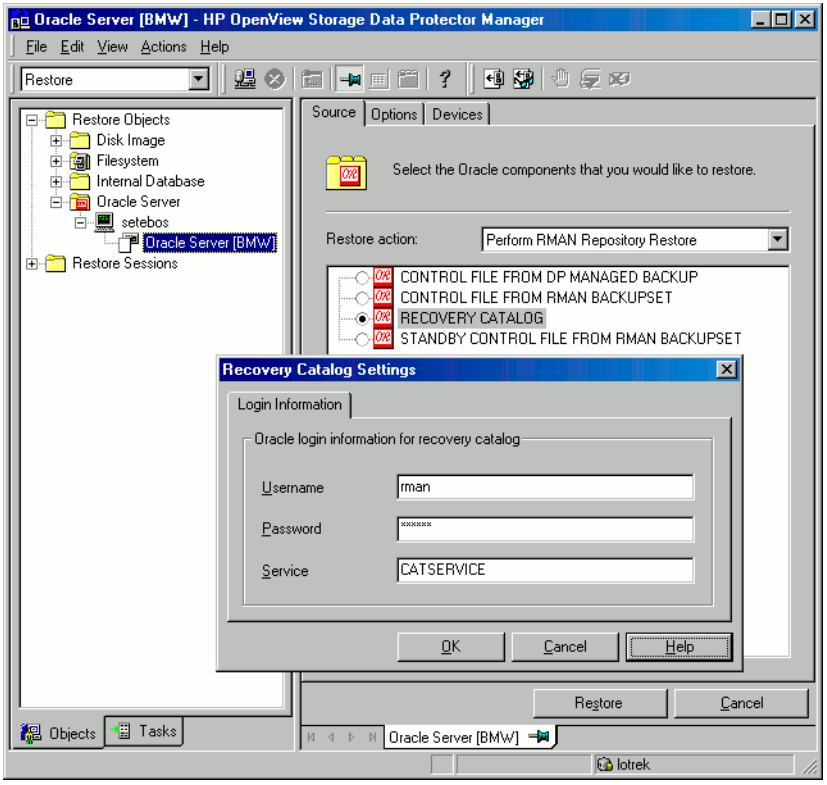

5. In the `Options` page:

   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 88.

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 (ctrl<*DB_NAME*>.dbf) are restored to:

  Windows: <*Data_Protector_home*>\tmp

  HP-UX and Solaris: /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 73.
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 88 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 74 and “Restoring the Control File” on page 76.

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 73.
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 88.

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 93.

Figure 1-19 Source Page

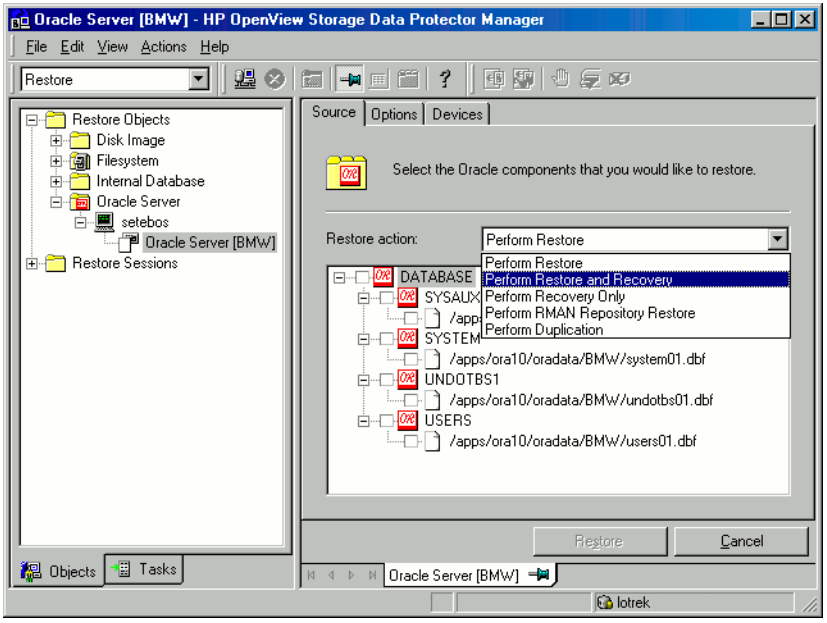

6. In the Results Area, select objects for restore.

   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 88 for information.

Figure 1-20 Options Page

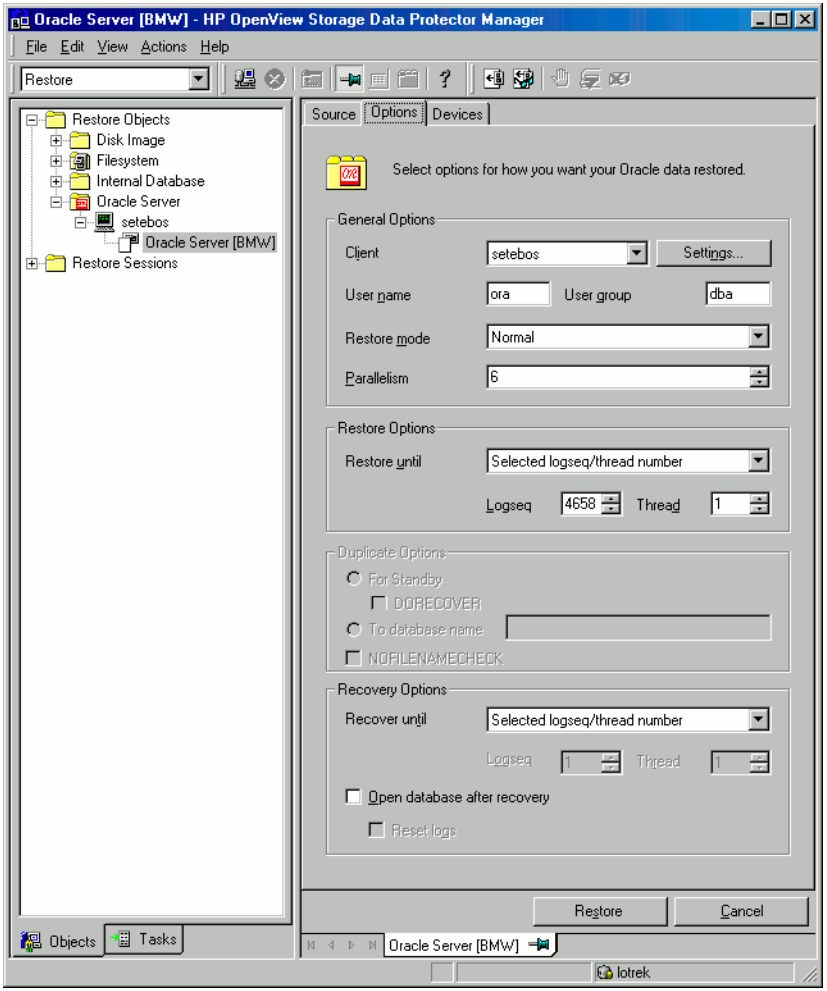

8. In the `Devices` page, select the devices to be used for the restore. You can restore using a device other than that used for backup, although Data Protector defaults to the original device on which the backup was made. To change the device from which an item is restored, select your desired device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

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

Figure 1-21 Devices Page

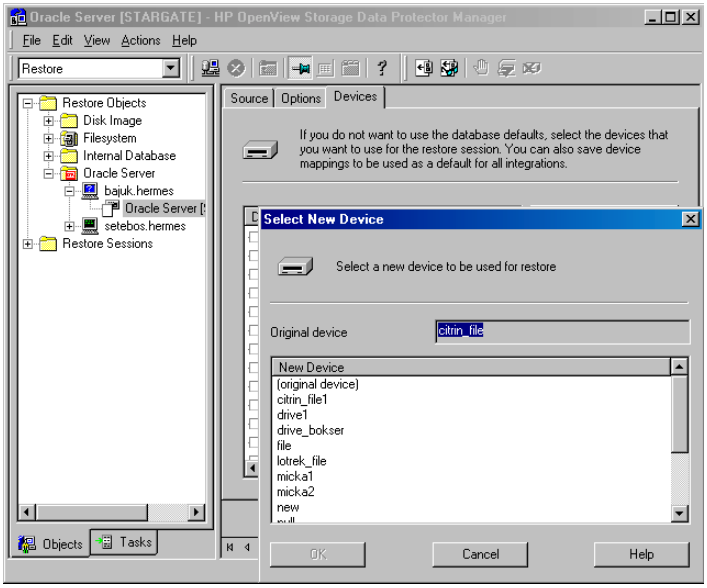

9. Click `Restore`.

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 72.

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 72.

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 73.
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 90 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 1-22 Oracle Duplicate Options

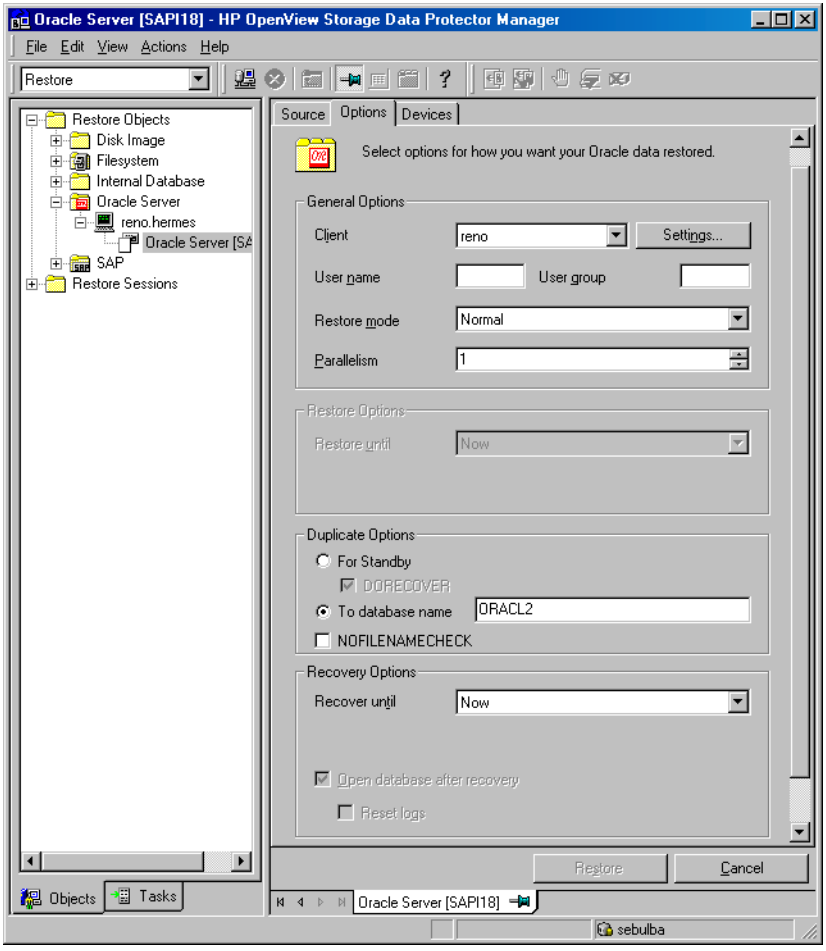

6. Click `Restore`.

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. Users 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. The user can select which redo logs to apply according to SCN number, logseq, or 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 93.

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 21 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 with version of Data Protector older than A.05.00.

- 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 using Data Protector version A.05.10.

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. If you do not enter a value, the number of parallel streams defaults to 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_Login>
```

```
RMAN> 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.
- If a backup of a control file is used in recovery or restore and recovery.

Do not reset the logs:

- After a complete recovery 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 `Restore until` option is set to `Now`, a warning is displayed, stating that you should reset the logs only if you use a backup of the 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 97
- "Example of Point-in-Time Restore" on page 98
- "Example of Tablespace Restore and Recovery" on page 100
- "Example of Datafile Restore and Recovery" on page 102
- "Example of Archive Log Restore" on page 105

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 "Restoring an Oracle Database" on page 70 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 databaseor 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 databaseor 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.

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
'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=<DB_NAME>)';
restore database;
recover database;
sql 'alter database open';
release channel 'dev1';
}
```

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

1. Create a file `restore_database` in the `/var/opt/omni/tmp` (UNIX systems) or `<Data_Protector_home>\tmp` 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 nocatalog`
- On UNIX: `<ORACLE_HOME>/bin/rman target sys/manager@PROD nocatalog`
- On OpenVMS: `rman target sys/manager@PROD nocatalog`

2. Start the point-in-time restore:

```
run{
allocate channel 'dev1' type 'sbt_tape' parms
'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=<DB_NAME>)';
set until time 'Mar 14 2004 11:40:00';
restore database;
recover database;
sql 'alter database open';
release channel 'dev1';
}
```

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
     '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
     '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
'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
'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
'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
'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
'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 and Solaris:

   ```
   /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:

   ```
   $@OMNI$ROOT:[BIN]OMNI$CLI_SETUP.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 OpenView Storage 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 72.
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 an Oracle Backup and Restore

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, proceed as follows:

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, proceed as follows:

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”.

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 and Solaris: `/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 on UNIX and OpenVMS Systems

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.

After you have uninstalled the Data Protector Oracle integration on the Oracle server system, proceed as described in the sections “Removing the Data Protector Oracle Integration Link on HP-UX Systems” on page 112 or “Removing the Data Protector Oracle Integration Link on Solaris and other UNIX Systems” on page 113.

Removing the Data Protector Oracle Integration 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 Data Protector Oracle Integration 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 Data Protector Oracle Integration Link on OpenVMS Systems

Oracle 8i

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

Oracle 9i

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.

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
'ENV=(OB2MAINTENANCE=1)';

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

release channel;
```

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 lists general checks and verifications, and problems you might encounter when using the Data Protector Oracle integration.

For general Data Protector troubleshooting information, see the *HP OpenView Storage 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 OpenView Storage 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.

Using Oracle After Removing the Data Protector Oracle Integration on UNIX Systems

After uninstalling the Data Protector Oracle integration on an Oracle server system, the Oracle server software is still linked to MML. You must rebuild (Oracle 8) or re-link (Oracle 8i/9i/10g) 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 UNIX and OpenVMS Systems” on page 112 for more information on how to make the Oracle server functional again.

General Troubleshooting

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

Problem

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 `svrmgrl` 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 `svrmgrl` binary, rather it changes its permissions. Once the `svrmgrl` binary is renamed, Data Protector will use Oracle 9i `sqlplus`, as it should, to complete the operations correctly.

Checking Prerequisites Related to the Oracle Side of the Integration on UNIX Systems

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

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

   Export *<ORACLE_HOME>* and *<DB_NAME>* 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>"
     ```

Start SQL*Plus from the *<ORACLE_HOME>* directory:

```
bin/sqlplus /nolog
```

Start SQL*Plus and type:

```
connect <Target database login>;
select * from dba_tablespaces;
exit
```

If this fails, open the Oracle target database.

2. **Verify that you can access the recovery catalog (if used) as follows:**

   Export *<ORACLE_HOME>* and *<DB_NAME>* as described on page 117.

   Start SQL*Plus from the *<ORACLE_HOME>* directory:

   ```
   bin/sqlplus /nolog
   ```

   Start SQL*Plus and type:

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

   If this fails, open the recovery catalog.

3. **Verify that the TNS listener is correctly configured for the Oracle target database and for the recovery catalog database. This is required for properly establishing network connections:**

   Export *<ORACLE_HOME>* as described on page 117.

   Start the listener from the *<ORACLE_HOME>* directory:

   ```
   bin/lsnrctl status <service>
   exit
   ```

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

   Export *<ORACLE_HOME>* as described on page 117.

Start SQL*Plus from the *<ORACLE_HOME>* directory:

```
bin/sqlplus /nolog
```

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 TNS configuration file (`TNSNAMES.ORA`).

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

   Export *<ORACLE_HOME>* as described on page 117.

   Start SQL*Plus from the *<ORACLE_HOME>* directory:

   ```
   bin/sqlplus /nolog
   ```

   Start SQL*Plus and type:

   ```
   connect <Target_Database_Login> as SYSDBA
   exit
   ```

   and

   ```
   bin/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 `init`*<DB_NAME>*`.ora` file.

5. **If you use the recovery catalog database, verify that the target database is registered in the recovery catalog**:

   Export *<ORACLE_HOME>* as described on page 117 and start SQL*Plus:

   ```
   bin/sqlplus /nolog
   ```

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 details about how to register an Oracle target database in the recovery catalog database.

6. **Verify backup and restore directly to disk using an RMAN channel type disk.**

   If you use the recovery catalog:

   Export <*ORACLE_HOME*> as described on page 117 and start RMAN:

   Oracle 9i/10g:

   ```
   bin/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:

   Export <*ORACLE_HOME*> as described on page 117 and start RMAN:
   `bin/rman target` <*Target_Database_Login*> `nocatalog cmd_file=rman_script`

   An example of the RMAN 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.

Checking Prerequisites Related to the Oracle Side of the Integration on Windows Systems

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

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

   Set *<ORACLE_HOME>* and *<DB_NAME>*:

   Start SQL*Plus from the *<ORACLE_HOME>* directory:

   ```
   sqlplus /nolog
   ```

   Start SQL*Plus and type:

   ```
   connect <user>/<password>@<service> as sysdba
   select * from dba_tablespaces;
   exit
   ```

   If this fails, open the Oracle target database.

2. **Verify that you can access the recovery catalog (if used) as follows:**

   Set the *<ORACLE_HOME>* and the *<DB_NAME>*.

   Start SQL*Plus from the *<ORACLE_HOME>* directory:

   ```
   sqlplus /nolog
   ```

   Start SQL*Plus and type:

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

   If this fails, open the recovery catalog.

3. **Verify that the TNS listener is correctly configured for the Oracle target database and for the recovery catalog database. This is required for properly establishing network connections:**

   From the *<ORACLE_HOME>* directory run the following command:

   `bin\lsnrctl status <service>`

   If this fails, startup the TNS listener process and see 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`, click `Administrative Tools, Services`.

Figure 1-23 Checking the Status of the Oracle Listener

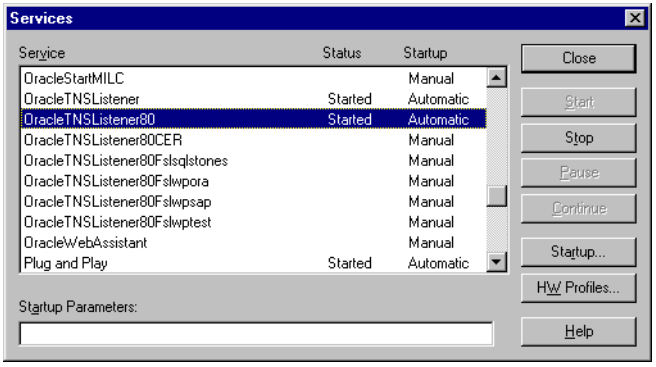

   a. The status of the respective listener service in the `Services` window should be `Started`, otherwise you must start it manually.

   b. From *<ORACLE_HOME>* run:

      `sqlplus /nolog`

      Start SQL*Plus and type:

      `connect <Recovery_Catalog_Login>`

      `exit`

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

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

   Set the *<ORACLE_HOME>* directory.

   Start SQL*Plus from the *<ORACLE_HOME>* directory:

   ```
   bin\sqlplus /nolog
   ```

   Start SQL*Plus and type:

   ```
   connect <Target_Database_Login> as SYSDBA
   exit
   ```

   Connect to the recovery catalog:

   ```
   bin\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 `init`*<DB_NAME>*`.ora` file.

5. **If you use the recovery catalog database, verify that the target database is registered in the recovery catalog**:

   ```
   bin\sqlplus
   ```

   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 details about how to register an Oracle target database in the recovery catalog database.

6. **Verify backup and restore directly to disk using an RMAN channel type disk.**

   If you use the recovery catalog:

   Set <*ORACLE_HOME*> *and* start RMAN from the <*ORACLE_HOME*> directory:

   Oracle 9i/10g:

   ```
   bin\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:

   Set <*ORACLE_HOME*> *and* start RMAN from the <*ORACLE_HOME*> directory:

   ```
   bin\rman target <Target_Database_Login> nocatalog
   cmd_file=rman_script
   ```

   An example of the RMAN 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.

Configuration Problems on UNIX Systems

IMPORTANT If you have encountered any errors up to this point when performing the procedures described in the previous section, please contact Oracle support. The respective tests must be done before you even start checking the Data Protector Oracle configuration.

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

   See the *HP OpenView Storage Data Protector Installation and Licensing Guide* for details.

2. **Verify that the Data Protector Oracle integration Media Management Library (MML) is linked with the Oracle executable**

   Use the following command to check if the `libob2oracle8.sl (libob2oracle8_64bit.sl)` file is linked with the Oracle 8 executable. Note that on HP-UX IA-64 and Solaris systems, the extension for MML is `.so`, and on AIX, the extension is `.a`.

   Export *<ORACLE_HOME>* and *<DB_NAME>* as described on page 117.

   On HP-UX:

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

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

   On Solaris:

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

   On other UNIX:

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

   On IBM AIX systems:

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

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

   On Linux systems:

   `/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, then enable MML dynamic path as follows:

```
/usr/bin/chatr +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:

Figure 1-24 Output of the `ldd` command 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
```

3. **Perform a filesystem backup of the Oracle Server system**

   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.

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

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

4. **Verify the permissions of the current user account**

   Your user account should enable you to perform an Oracle backup or restore with Data Protector. Use the `testbar2` utility to check the permissions:

   `/opt/omni/bin/testbar2 -perform:checkuser` (HP-UX and Solaris systems) or

   `/usr/omni/bin/testbar2 -perform:checkuser` (other UNIX systems).

   If the user account holds all required permissions, you will receive only NORMAL messages displayed on the screen. See also “Configuring Oracle Users on UNIX and OpenVMS” on page 21.

5. **Examine the system errors**

   The system errors are reported in the /var/opt/omni/log/debug.log (HP-UX and Solaris systems) or `/usr/omni/log/debug.log` (other UNIX systems) file on the Oracle Server system.

Configuration Problems on Windows Systems

IMPORTANT If you have encountered any errors up to this point when performing the procedures described in the previous section, please contact Oracle support. The respective tests must be done before you even start checking the Data Protector Oracle configuration.

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

   See the *HP OpenView Storage Data Protector Installation and Licensing Guide* for details.

2. **Verify that the Data Protector Oracle integration Media Management Library (MML) is loaded**

   Once the setup has finished, you need to restart the Oracle services and verify that the `<Drive_Label>:\<%SystemRoot%>\system32\orasbt.dll` MML is loaded. Proceed as follows:

   a. Switch to the `<DriveLabel>:\<%SystemRoot%>\system32` directory and right-click orasbt.dll.

   b. Select `Properties` and click the `Version` tab from the `orasbt.dll Properties` window. In the `Description` field, you should see the file described as a part of the Data Protector integration.

3. **Perform a filesystem backup of the Oracle Server system**

   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.

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

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

4. **Verify the inet startup parameters:**

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

   a. In the `Control Panel`, go to `Administrative Tools`, `Services`.

   b. 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 1-25 Checking the `Inet` Start-Up Parameters:

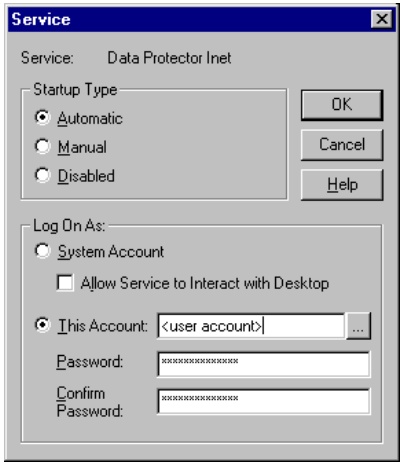

5. **Examine the system errors**

   The system errors are reported in the *<Data_Protector_home>*\log\debug.log file on the Oracle Server system.

Backup Problems on UNIX Systems

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

1. **Check your Oracle Server configuration**

   To check the configuration, log in to the Oracle server system as the user `root` or as the Oracle user that is identified as described in “Configuring Oracle Users on UNIX and OpenVMS” on page 21. The identified Oracle user and the user `root` must also be added to Data Protector `admin` or `operator` group. Then check the configuration as described in “Checking the Configuration” on page 33.

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 Oracle Server system. Check the `/etc/opt/omni/client/cell_server` (HP-UX and Solaris 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:

On HP-UX and Solaris systems:

```
/opt/omni/bin/testbar2 -type:Oracle8 -appname:<DB_NAME>
-bar:<backup_specification_name> -perform:backup
```

On other UNIX systems:

```
/usr/omni/bin/testbar2 -type:Oracle8 -appname:<DB_NAME>
-bar:<backup_specification_name> -perform:backup
```

Switch to the Data Protector Manager and 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:

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 OpenView Storage Data Protector Troubleshooting Guide* for instructions on troubleshooting devices.

Data Protector reports “Export of the Recovery Catalog Database Failed” when backing up Oracle 9i

Problem

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.

Data Protector reports “Cannot allocate/attach shared memory”

Problem

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` option in the `/opt/omni/.omnirc` file to 1 in order to use the memory windowing properly, and restart the failed backup session. See the *HP OpenView Storage Data Protector Troubleshooting Guide* for details on using `omnirc` options.

Backup Fails After a Point in Time Restore and Recovery

Problem

Backup fails after a point in time restore and recovery was performed and 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:

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.

Backup of Archive Logs on RAC Cannot be Performed

Problem

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
'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=<DB_NAME>,OB2BARLIST=RAC
_arch)' connect <username>/<passwd>@<INSTANCE 1>;
allocate channel 'dev_2' type 'sbt_tape'
 parms
'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;
}
```

Backup Problems on Windows

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

1. **Check your Oracle Server configuration**

   To check the Oracle Server configuration, see “Checking the Configuration” on page 33.

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 Oracle Server system. Check the *<Data_Protector_home>*`\Config\client\cell_server` file, which contains the name of the Cell Manager system. Then run:

   ```
   <Data_Protector_home>\bin\testbar2 -type:Oracle8
   -appname:<DB_NAME> -bar:<backup_specification_name>
   -perform:backup
   ```

   Switch to the Data Protector Manager and 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:

   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 OpenView Storage Data Protector Troubleshooting Guide* for instructions on troubleshooting devices.

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

Problem

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.

Backup Fails After a Point in Time Restore and Recovery

Problem

Backup fails after a point in time restore and recovery was performed and 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.

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 an object exists on the backup media**

   This can be done by running the following command on the Oracle server system:

   - On HP-UX and Solaris: `/opt/omni/bin/omnidb -oracle8 "<object_name>" -session "<Session_ID>" -media`
   - On other UNIX: `/usr/omni/bin/omnidb -oracle8 "<object_name>" -session "<Session_ID>" -media (other UNIX systems)`
   - On Windows: `<Data_Protector_home>\bin\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.

2. **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 Oracle Server system. Check the `/etc/opt/omni/client/cell_server` (HP-UX and Solaris systems), `/usr/omni/config/cell/cell_server` (other UNIX systems), or `<Data_Protector_home>\Config\client\cell_server` (Windows systems) file, which contains the name of the Cell Manager system.

   Test Data Protector internal data transfer using the `testbar2` utility:

HP-UX and Solaris

```
/opt/omni/bin/testbar2
-type:Oracle8
-appname:<DB_NAME>
-perform:restore
-object:<object_name>
-version:<object_version>
-bar:<backup_specification_name>
```

Other UNIX

```
/usr/omni/bin/testbar2
-type:Oracle8
-appname:<DB_NAME>
-perform:restore
-object:<object_name>
-version:<object_version>
-bar:<backup_specification_name>
```

Windows

```
<Data_Protector_home>\bin\testbar2 -type:Oracle8
-appname:<DB_NAME>
-perform:restore
-object:<object_name>
-version:<object_version>
-bar:<backup_specification_name>
```

IMPORTANT The hostname should not be specified in the `object` option. It is automatically provided by `testbar2`.

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:

Run the `omnidb` command to view the objects in the database.

3. **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
  ```

4. **Check your environment variables.**

   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."
   ```

   To resolve this problem do as follows:

   a. Close Data Protector.

   b. Set the environment variable on the system where the Cell Manager resides:

      ```
      OB2_ORARESTHOSTNAME = <target Oracle host>
      ```

   c. Restart Data Protector and try to restore the database items again.

   d. When the restore is complete, close Data Protector and re-set the following environment variable:

      ```
      OB2_ORARESTHOSTNAME = <empty>
      ```

   e. Restart Data Protector.

5. **Try using the RMAN CLI to restore the database items.**

   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 about using the CLI see “Restoring Oracle Using RMAN” on page 93.

6. **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;
   ```

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 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.

Problem **"Binary util_orarest failed" 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 failed. Cannot get information from the
remote host.
```

Action 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 backupset 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>,OB2
BARHOSTNAME=<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 a database recovery:

```
run {
allocate channel 'dev_0' type 'sbt_tape';
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, `libclntsh.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 `libclntsh.sl`. See *OracleMetaLink, DOC ID: 113395.1* for information on how to uninstall the Radius Authentication Adapter.

2 Integrating SAP R/3 and Data Protector

In This Chapter

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

The chapter is organized into the following sections:

Introduction

Data Protector integrates with the SAP R/3 Database Server to offer online backup of your SAP R/3 databases.

If the SAP R/3 system uses an Oracle database, then the Data Protector SAP R/3 integration can be used for backup. If any other database is used by SAP, then the corresponding Data Protector integration of that database (for example, Informix) must be used instead.

Refer to the *HP OpenView Storage Data Protector Product Announcements, Software Notes, and References* or http://www.hp.com/support/manuals for up-to-date information about platforms supported by the integration.

The online backup concept is now widely accepted because it addresses the business requirements of high application availability. During backup, the database is online and actively used. The backup is performed quickly and efficiently, with the least possible impact on database performance.

The SAP R/3 part of the integration provides storage management utilities. These utilities communicate with Data Protector via the Data Protector backint executable, which complies with the SAP R/3 backup interface.

Advantages

Using Data Protector with the SAP R/3 Database Server offers several advantages over using SAP R/3 alone:

- Central Management for all backup operations

  You can manage backup operations from a central point. This is especially important in large business environments.

- Backup Management

  Backed up data can be duplicated during or after the backup to increase fault tolerance of backups, to improve data security and availability, or for vaulting purposes.

- Media Management

Data Protector has an advanced media management system that allows you to keep track of all media and the status of each medium, set the protection for stored data, fully automate operations as well as organize and manage devices and media.

- Scheduling

  Data Protector has a built-in scheduler that allows you to automate backups to run periodically. With the Data Protector scheduler, the backups you configure run unattended at the periods you specify.

- Local versus Network Backups

  When configuring an SAP R/3 backup, the location of devices is completely transparent to the user. They can be connected to the SAP R/3 Database Server or any other Data Protector clients on the network.

- Device Support

  Data Protector supports a wide range of devices, from standalone drives to complex multiple drive libraries. Refer to the *HP OpenView Storage Data Protector Product Announcements, Software Notes, and References* or http://www.hp.com/support/manuals for an up-to-date list of supported devices and other information.

- Reporting

  Data Protector has reporting capabilities that allow you to receive information about your backup environment. You can schedule reports to be issued at a specific time or attached to a predefined set of events, such as the end of a backup session or a mount request.

- Monitoring

  Data Protector has a feature that allows you to monitor currently running sessions and view finished sessions from any system that has the Data Protector User Interface installed.

  All backup sessions are logged in the built-in IDB, providing you with a history of activities that can be queried at a later time.

Prerequisites and Limitations

This section provides you with a list of prerequisites and limitations you must be aware of before using the integration.

Prerequisites

- The database used by SAP R/3 must be an Oracle database. If any other database is used by SAP, then the corresponding Data Protector integration of that database (for example, Informix) must be used instead.
- You need a license to use the Data Protector SAP R/3 integration. Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide* for information about licensing.
- Before you begin, make sure that you have correctly installed and configured the SAP R/3 Database Server and Data Protector systems. Refer to the:
  - *HP OpenView Storage 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 OpenView Storage Data Protector Installation and Licensing Guide* for instructions on how to install Data Protector on various architectures and how to install the Data Protector SAP R/3 integration.
  - *SAP R/3 System Online Documentation* for instructions on how to install and configure the SAP R/3 database and the SAP R/3 backup and restore tools (BRBACKUP, BRRESTORE, and BRARCHIVE).
- The SAP R/3 database *user* used by this integration to connect to the target SAP R/3 database during the backup must have the SYSDBA privilege granted. Refer to the Oracle documentation for more information on user privileges in Oracle.

  The operating system `root` user on the Oracle Server also has to be added to either the Data Protector `admin` or `operator` user group.

Limitations

Refer to the *HP OpenView Storage Data Protector Product Announcements, Software Notes, and References* for a list of general Data Protector limitations. This section describes limitations specific for this integration.

- Do not use double quotes (" ") in object-specific pre-exec and post-exec commands.
- Do not configure RMAN Offline SAP R/3 backups using the `Internal` user since the backups will not work.

Integration Concept

This integration links SAP R/3 backup utilities (BRTOOLS) with Data Protector. SAP R/3 backup utilities provide an interface between an SAP R/3 Database Server and media management applications, like Data Protector. They enable the backup or restore of the following SAP R/3 data objects:

- data files
- control files
- online redo logs
- offline (archived) redo logs
- SAP R/3 logs and parameter files

Because SAP R/3 Database Servers run on top of Oracle databases, the SAP R/3 backup objects are very similar to those of Oracle. The main difference is that SAP R/3 backup utilities hide the database from Data Protector, which sees those objects as plain files.

Version 4.5 and higher of the SAP R/3 backup utilities allows Oracle data files to be backed up directly using the Oracle Recovery Manager (hereafter referred to as **RMAN mode**), as well as using the Data Protector Oracle Integration (hereafter referred to as the **backint mode**). This is very useful because RMAN supports incremental backups, and thus the backup time and the amount of backed up data can be significantly reduced.

SAP R/3 Backup Utilities

SAP R/3 backup utilities are the following:

- BRBACKUP

  This utility performs online and offline backup of control files, data files, and online redo log files. Additionally, BRBACKUP saves the profiles and logs relevant for a particular backup session.

- BRARCHIVE

  This utility performs backups of the offline (archived) redo logs, written by Oracle to the archiving directory.

- BRRESTORE

  This utility restores the backed up data using the BRBACKUP and BRARCHIVE utilities.

These backup utilities can be started directly using Data Protector, or interactively using SAPDBA, which is an SAP R/3 administration utility.

NOTE

Data Protector supports all SAP R/3 backup utilities options, except for the `-a` and `-b` options. In order for Data Protector to support also the `-a` and `-b` options, set the `OB2BRTNOSECU` `omnirc` variable to `1`. For more information about the `omnirc` file, refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide*.

Data Protector Integration Software

The Data Protector integration software consists of the following components, as depicted in Figure 2-1 on page 151.

- The `backint` program is a backup interface between the Data Protector software and the SAP R/3 backup and restore tools.

  It is started using BRBACKUP or BRARCHIVE during a backup session, and BRRESTORE during a restore session.
- The `sapback` program performs the actual backup of files.
- The `saprest` program performs the actual restore of files.
- The Data Protector Database Library links Data Protector and Oracle Server software. This is required only if SAP R/3 is backed up in the RMAN mode.
- The `omnisap.exe` program is used by Data Protector to start the SAP R/3 backup tools.
- The `testbar2` utility checks the Data Protector part of the integration.
- The `util_sap.exe` program is used by Data Protector to configure the integration.
- The configuration file on the Cell Manager system contains data needed by Data Protector to run backups and restores.

Figure 2-1 **SAP R/3 Backup Concept**

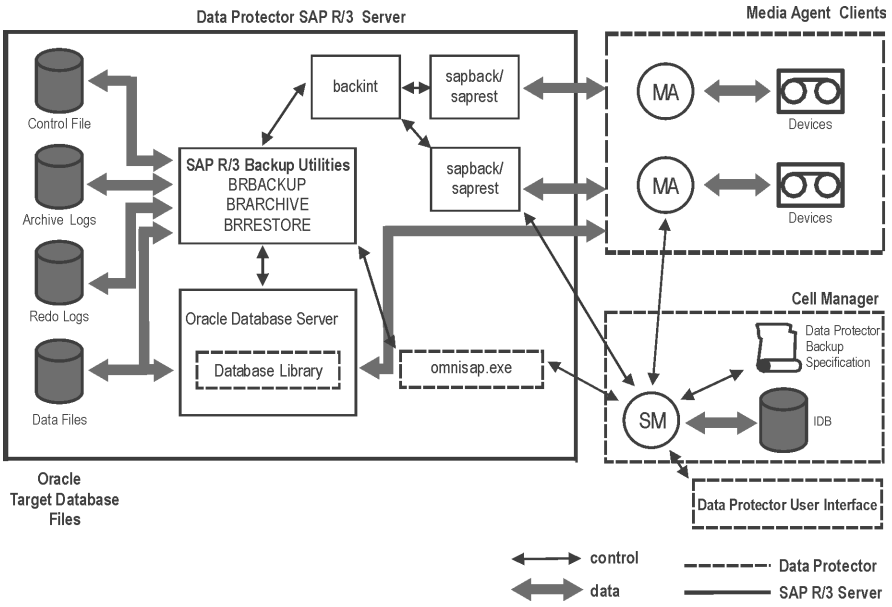

Legend

SM — The Data Protector Session Manager, which is the Data Protector Backup Session Manager during backup or the Data Protector Restore Session Manager during restore.

Database Library — The interface between SAP R/3 Server processes and Data Protector.

IDB — The IDB, which stores information about Data Protector sessions, such as session messages, and information about objects, data, used devices, and media.

MA — The Data Protector General Media Agent.

SAP R/3 Architecture

Depending on the backup mode, there are two possible backup scenarios (backint mode or RMAN mode) that can be used.

Backup Flow Using Backint

The backup session undergoes the following stages if the backup is performed in backint mode. See Figure 2-2 for details.

NOTE

It is not possible to perform an incremental backup in backint mode.

Figure 2-2 SAP R/3 Architecture: Backint Mode

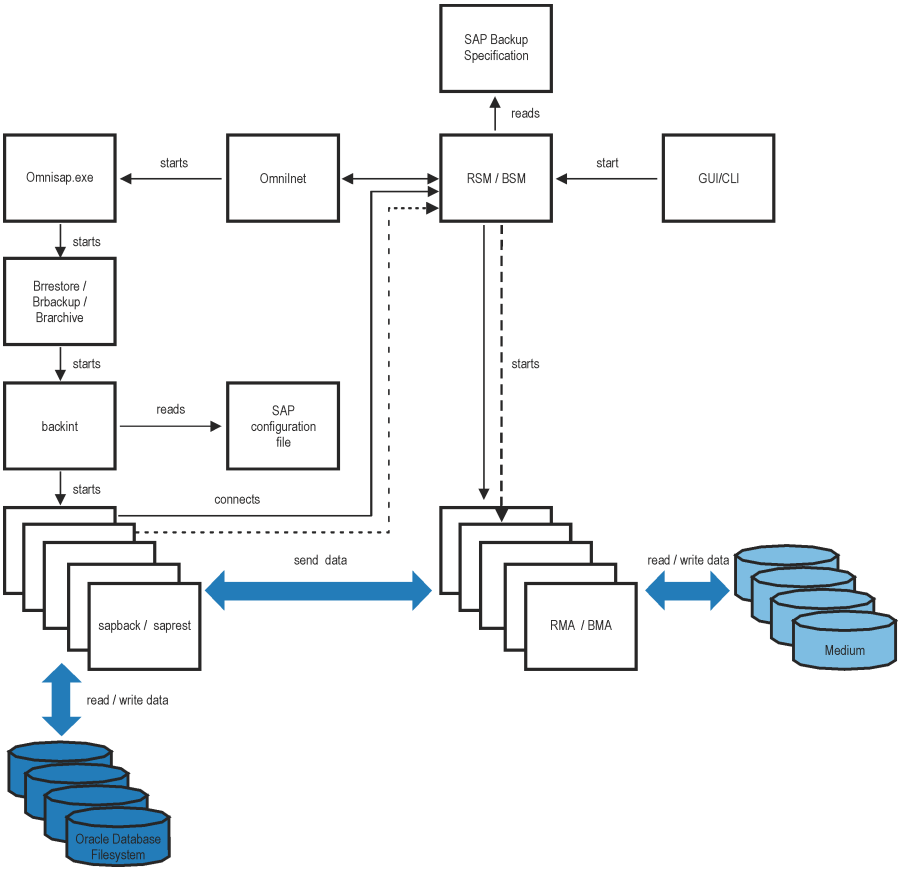

Legend

BSM	Data Protector Backup Session Manager
RSM	Data Protector Restore Session Manager

BMA	Data Protector Backup Media Agent
RMA	Data Protector Restore Media Agent
GUI/CLI	Data Protector User Interface

1. The backup session can be started using the Data Protector GUI, or interactively using the SAP R/3 utilities.

   If the backup session is started using the Data Protector User Interface (or using the scheduler), then the Backup Session Manager (BSM) is started. The BSM then reads the appropriate Data Protector backup specification, checks if the devices are available, and starts the `omnisap.exe` program on the SAP R/3 Database Server.

   The `omnisap.exe` program exports the appropriate environment variables and starts either the BRBACKUP or BRARCHIVE utilities. These utilities then initiate the first backint command to back up the Oracle Target Database's data files and the control files (BRBACKUP) or to back up archived redo log files (BRARCHIVE).

   If the backup is started interactively using the SAPDBA program, then the BRBACKUP or BRARCHIVE utilities are started directly.

2. BRBACKUP does the following:

   - Automatically 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 before the backup.

     The archived redo log files are written to the archiving directory by Oracle and are backed up later using BRARCHIVE.
   - Writes the BRBACKUP log during the backup session, with information about the backup file and the backup ID. These logs must be available in order to determine the location of the database files and archived redo log files during restores.
   - Sets the tablespace mode (BEGIN / END BACKUP) in the case of online backup using backint.

     In this way, the SAP R/3 puts the tablespace in backup mode just before it is backed up, and puts the tablespace back in normal mode immediately after the backup is completed. The tablespaces are therefore in backup mode for a minimal amount of time.

3. The backint program obtains the SAP R/3 configuration from the Cell Manager, divides the files for backup into subsets (provided that the specified concurrency is greater than 1) and starts the `sapback` program for each subset. Each `sapback` process connects to the BSM, which then starts General Media Agents on the corresponding client systems and establishes a connection between the `sapback` processes and General Media Agents.

   Data transfer can begin at this point. The `sapback` processes read data from disks and send it to General Media Agents. The first backint program stops as soon as all `sapback` processes have finished and control is returned to the parent process, either the BRBACKUP or BRARCHIVE utility.

   The second backint command is initiated by either the BRBACKUP or BRARCHIVE command. This command attempts to back up the SAP R/3 log files and parameter files (in the case of BRBACKUP), or the archived redo logs (in the case of BRARCHIVE) that have been created since the first backint command.

   If new archived redo logs have been created, they are backed up and another backint command is started. Otherwise, the SAP R/3 log files and the parameter files are backed up, and the second backint program is started using BRBACKUP.

   Therefore, more than two backint commands may be initiated by BRARCHIVE, while there are only two backint commands initiated by BRBACKUP.

   If archive logs were backed up, omnisap creates a copy of the control files either in the directory defined by the SAPBACKUP variable, or in `/var/opt/omni/tmp` (on UNIX) or `<Data_Protector_home>\tmp` (on Windows) if the variable is not set. The control file is then backed up by the backint utility using sapbackup.

NOTE The total number of sapback processes started in one session using Data Protector is limited to 256.

4. General Media Agents finish transferring data when all the `sapback` processes are complete. When all of the General Media Agents have finished data transfer, the BSM waits for a timeout

(`SmWaitForNewClient` `omnirc` global variable) and completes the backup session, as long as no backint is started within this time frame.

Backup Flow Using Recovery Manager

A backup session using RMAN mode differs from a backup session in backint mode in step 3. See Figure 2-3 on page 156 for details.

BRBACKUP starts RMAN, which then connects to the Data Protector Database Library via the Oracle Server processes. The Database Library provides a connection to the Data Protector BSM, which starts General Media Agents and establishes a connection between the Oracle Server and General Media Agents.

The data transfer begins at this point. The Oracle Server sends data to General Media Agents, which then write the data to the media.

Once the Oracle Target Database's data files have been written to the media, the respective Oracle Server processes are completed, and so, subsequently, is RMAN. The backup control is now returned to BRBACKUP, which starts the first backint command to back up the Oracle Target Database's control file and the SAP R/3 log files. Archive logs are backed up in the same manner as in backint mode.

Figure 2-3 **SAP R/3 Architecture: RMAN Mode**

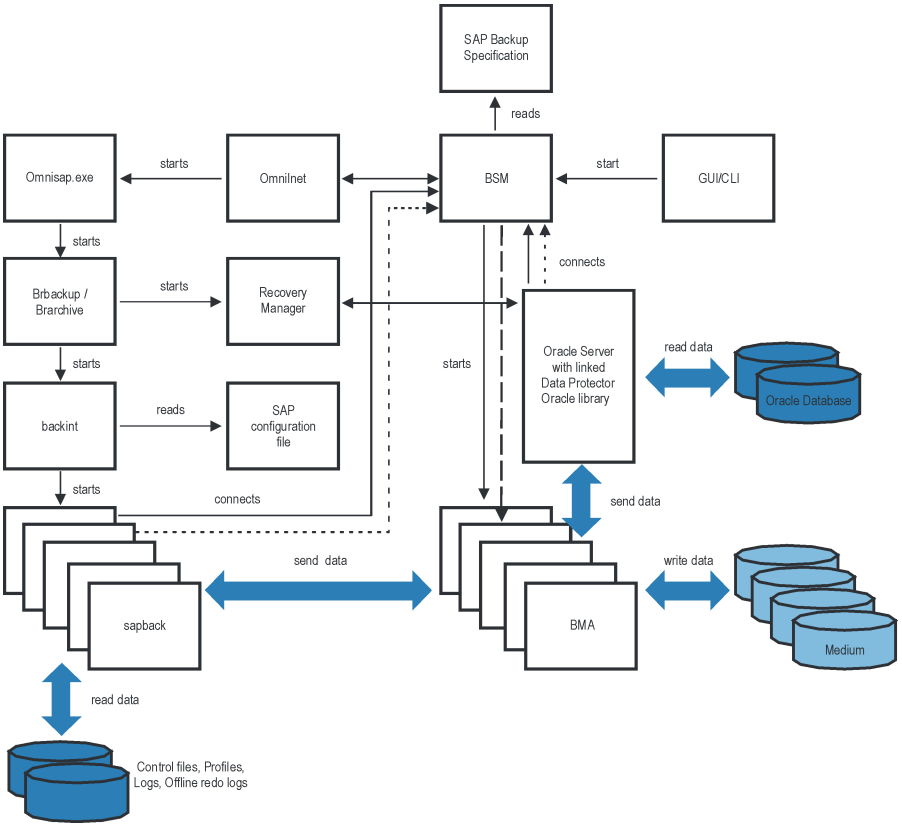

Legend

BSM	Data Protector Backup Session Manager
BMA	Data Protector Backup General Media Agent
GUI/CLI	Data Protector User Interface

Restore Flow Using Backint

SAP R/3 restore can be initiated using Data Protector, or interactively using the SAP R/3 utilities. However, only a standard filesystem restore is performed using Data Protector.

The restore session proceeds according to the following stages if the restore is performed in backint mode.

1. Using the SAPDBA utility, the objects to be restored are selected.
2. The BRRESTORE first checks whether the required free disk space is available to allow the files to be restored. It then starts the first backint command to restore the Oracle Target Database's data files. The backint command reads the SAP R/3 configuration file, divides the files for restore into subsets (provided that the specified concurrency is greater than 1) and starts the `saprest` process for each subset.

   The first `saprest` process starts the Data Protector Restore Session Manager (RSM), while the subsequent `saprest` processes connect to the same RSM. In addition, the `saprest` process checks whether the specified objects have been backed up.

   The RSM checks the availability of the restore devices, starts General Media Agents and establishes a connection between the `saprest` processes and General Media Agents. Data transfer begins at this stage. Data is sent from the media to the target disks. The General Media Agent finishes as soon as all `saprest` processes connected to it are completed.
3. When all the General Media Agents have finished, the RSM waits for a timeout (`SmWaitForNewClient` global variable) and completes the restore session, if no backint is started within this time frame.

Restore Flow Using Recovery Manager

A restore session using RMAN differs from a restore session using the backint mode in the step 2 as follows:

BRRESTORE starts RMAN in order to restore the Oracle Target Database data files. RMAN then connects to the Data Protector Database Library via the Oracle Server processes.

Data Protector SAP R/3 Configuration File

Data Protector stores the SAP R/3 integration parameters for every configured SAP R/3 instance 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
- 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>';

Environment={

[<ENV var1>='<value1>';]

[<ENV var2>='<value2>';

  ...]

}

SAP_Parameters={<bckup_spec_name>=('-concurrency <# of
concurrency>' | '-time_balance' | '-load_balance' |
'-manual_balance' | '-channels <#_of_RMAN_channels>');

}

speed={

AVERAGE=1;

'<filename>'=<# of seconds needed to backup this file>;

}

compression={'<filename>'=<size of the file in bytes after the
compression>;

}

manual_balance={<backup_specifiaction_name>={'<filename>'=<device_
number>;

}

 }
```

Example

This is an example of the file:

```
ORACLE_HOME='/app/oracle805/product';
ConnStr='EIBBKIBBEIBBFIBBGHBBOHBBQDBBOFBBCFBBPFBBCFBBIFBBGFBBDGBBB
FBBCFBBDFBBCFBB';

BR_directory='/usr/sap/ABA/SYS/exe/run'; SAPDATA_HOME='/sap';

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 configuration of the SAP R/3 instance in Data Protector 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 and Solaris 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 <sap_instance> [-local <filename>]

util_cmd -getopt[ion] [SAP <sap_instance>] <option_name>
[-sub[list] <sublist_name>] [-local <filename>]

util_cmd -putopt[ion] [SAP <sap_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 NLS_LANG parameters for the SAP R/3 instance ICE, 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 INSTANCE.txt' -sublist Environment

<Data_Protector_home>\bin\util_cmd -putopt SAP ICE
NLS_LANG 'AMERICAN_AMERICA.US7ASCII' -sublist Environment

<Data_Protector_home>\bin\util_cmd -putopt SAP ICE
NLS_LANG "'10'" -sublist Environment
```

HP-UX and Solaris

```
/opt/omni/lbin/util_cmd -putopt SAP ICE OB2OPTS '-debug \
1-200 INSTANCE.txt' -sublist Environment

/opt/omni/lbin/util_cmd -putopt SAP ICE NLS_LANG \
'AMERICAN_AMERICA.US7ASCII' -sublist Environment

/opt/omni/lbin/util_cmd -putopt SAP ICE BR_TRACE "'10'"
-sublist Environment

/usr/omni/bin/util_cmd -putopt SAP ICE OB2OPTS '-debug \
1-200 INSTANCE.txt' -sublist Environment
```

Other UNIX

```
/usr/omni/bin/util_cmd -putopt SAP ICE NLS_LANG \
'AMERICAN_AMERICA.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 SAP R/3 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 and Solaris: `/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 SAP R/3 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 and Solaris: `/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 SAP R/3 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 and Solaris: `/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

Configuration Overview

Configuring the Data Protector SAP R/3 integration consists of these steps:

1. If you intend to use the Oracle Recovery Manager to backup the SAP R/3 database files, install and configure the Data Protector Oracle integration. When the Data Protector Oracle integration is configured, it is recommended to run a test Data Protector Oracle backup using the Oracle Recovery Manager.
2. Configure the SAP R/3 user (on UNIX systems only).
3. Configure the SAP R/3 Database Server.

Configuring an SAP R/3 User in Data Protector (UNIX Systems Only)

On UNIX systems, to start an SAP R/3 backup session, you need an operating system logon on the system where an SAP R/3 Database Server is running.

In addition, this user has to be registered in the Oracle database and identified by SAP R/3 through the operating system identification.

This means that Oracle Server does not request connection information from an application started under such user account, but only checks whether the user is registered in the database.

Refer to the SAP R/3 and Oracle documentation for further information about different types of connections, about roles and privileges of Oracle database administrators, and about security issues that should be considered.

Further on, this user is allowed to backup and restore an SAP R/3 database. In order to start a backup of an SAP R/3 database using Data Protector, this user has to become the owner of the Data Protector backup specification.

As the owner of the backup specification, the user has to be added to either the Data Protector `admin` or `operator` user group.

Such a user is the user ora<*SID*> from the group sapsys; or, you can identify such a user by running the following command on the SAP R/3 Database Server system:

```
ps -ef|grep ora_pmon_<ORACLE_SID>
```

or

```
ps -ef|grep ora_lgwr_<ORACLE_SID>
```

Figure 2-4 Finding the Oracle User

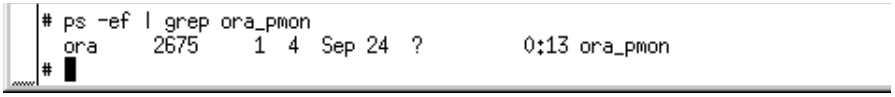

It can be seen from the example above that the user ora has sufficient privileges within the SAP R/3 database to backup and restore the SAP R/3 database. Therefore, this user has to be added to the corresponding Data Protector user group (admin or operator) and have to become the owner of the backup specification, so that the user is able to backup the SAP R/3 database using Data Protector.

IMPORTANT

Additionally, the operating system root user on the SAP R/3 Server also has to be added to either the Data Protector admin or operator user group.

After the two users are added, Data Protector sessions can be started under the user account with all the privileges required to perform an SAP R/3 database backup with Data Protector.

Sometimes SAP administrators want to enforce more security and allow restores to be performed only by using a specific user account (for example SAP administrator). In this case, this user should also be configured as a Data Protector user and have to be added to either the operator or admin group.

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

Configuring an SAP R/3 Database Server

Before You Begin

It is recommended that you configure and run a Data Protector test filesystem backup of the SAP R/3 Database Server (a client system in the Data Protector cell).

In case of problems, this type of backup is much easier to troubleshoot than the integration itself.

A test filesystem backup includes installing a Disk Agent on the SAP R/3 Database Server. Any device can be used for the test purposes only. Configure a standard filesystem backup, which can include one directory only. The test should include a partial restore to the SAP R/3 Database Server as well.

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

Configuring the SAP R/3 Database Server involves preparing the environment for performing backups. The environment parameters such as the Oracle home directory and the connection string to the Oracle Target Database are saved on the Cell Manager. The database must be online during the configuration procedure.

Cluster-Aware Clients on Windows/UNIX

You also need to edit the Data Protector `omnirc` file on each cluster node and specify the name of the cluster node in the SAPLOCALHOST variable. Below you see an example of the `omnirc` file:

```
# SAP R/3 related entries for clustering
# SAPLOCALHOST=<cluster_node_name>
```

NOTE

Make sure that the `SAPLOCALHOST` variable is not defined in the `Environment` section of the Data Protector SAP R/3 configuration file. Refer to “Data Protector SAP R/3 Configuration File” on page 158 for information on how to do that.

Cluster-Aware Clients on UNIX

Configure the Data Protector SAP R/3 integration on only one cluster node, since the Data Protector SAP R/3 configuration file resides on the Cell Manager. Use the virtual hostname when configuring the

integration. However, you need to create a link to the Data Protector `backint` interface on all other nodes. Enter the following command on all other nodes:

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

In a cluster environment, the environment variable `OB2BARHOSTNAME` must be defined as the virtual hostname before running the configuration 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>`

Add the SAP R/3 group dba user to Data Protector for the virtual server and for every node in the cluster.

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

For information on the Data Protector Cell Manager package configuration (if you want to install and configure the Data Protector Cell Manager in the MC/SG cluster), see the online Help index “MC/ServiceGuard integration”.

Cluster-Aware Clients on Windows

The client configuration must be performed on only one of the cluster nodes per one SAP R/3 server, since the Data Protector SAP R/3 configuration file resides on the Cell Manager.

However, the Data Protector `backint` program needs to be manually copied to the correct location on all other nodes. On every other node, copy the `<Data_Protector_home>\bin\backint.exe` to the directory where the SAP R/3 backup utilities reside.

NOTE

Each SAP R/3 instance must be configured separately.

NOTE

Make sure to set any Oracle and SAP R/3 related environment variables needed for the Oracle and SAP R/3 databases to function properly (for example, the Oracle `NLS_LANG` environment variable) on the SAP R/3 Database Server. Refer to the Oracle and SAP R/3 documentation for more information.

Data Protector Inet User Account on Windows

On Windows, set the service startup account of the `Data Protector Inet` service as an SAP administrator account. To configure the `Data Protector Inet` service startup account, go to `Control Panel`, then `Administrative Tools, Services`. Double-click the `Inet` service to configure it. This user must also be included in the `ORA_DBA` local group on the system where SAP R/3 instance is running.

Configuration of an SAP R/3 Database Server is performed using the *<Data_Protector_home>*`\bin\util_sap.exe` (Windows systems), `/opt/omni/lbin/util_sap.exe` (HP-UX and Solaris systems), or `/usr/omni/bin/util_sap.exe` (other UNIX systems) command.

On Windows, configuration can be started remotely using the Data Protector GUI from any Data Protector Windows client within the same Data Protector cell, or locally on the SAP R/3 Database Server.

The util_sap.exe Command

Use the `util_sap.exe` command to get the information you may need to configure your SAP R/3 Database Server. This will:

- List all Oracle instances on a particular system.

  ```
  util_sap.exe -APP
  ```

- List the tablespaces that belong to a particular Oracle instance:

  ```
  util_sap.exe -OBJS0 <ORACLE_SID>
  ```

- List the database files that belong to a particular tablespace of the Oracle instance:

  ```
  util_sap.exe -OBJS1 <ORACLE_SID> <TABLESPACE>
  ```

Using the CLI - UNIX Systems Only

On UNIX, to configure an SAP R/3 Database Server, execute the following command with root privileges on the SAP R/3 Database Server:

NOTE

Each instance must be configured separately.

`util_sap.exe -CONFIG <ORACLE_SID> <ORACLE_HOME> \ <targetdb_connection_string> <SAPTOOLS_DIR> \ [<SAPDATA_HOME>]`, where:

- *<ORACLE_SID>*

  is the is the name of the Oracle database instance to be configured

- *<ORACLE_HOME>*

  is the directory in which Oracle binaries are installed

- *<targetdb_connection_string>*

  is the login information to the target database of the *<user_name>*/*<password>*@*<service>* format, described in “Glossary” on page G-1.

  The *<user_name>* is the name by which a user is known to Oracle Server and to other users. Every user is identified by a password, and both must be entered to connect to an Oracle database. This user is, by default, used by `brbackup` and `brarchive` during backup. To define a different user when backing up, use the `-u` *<user_name>* as a `BR Backup` SAP R/3 backup option. See “SAP R/3 Specific Backup Options” on page 186.

NOTE

The user *<user_name>* is visible during backup when the `ps -ef` command is run.

- *<SAPTOOLS_DIR>*

  is the directory in which SAP R/3 backup utilities are stored. SAP recommends to install SAP R/3 backup utilities on both local nodes in the cluster in case the application is cluster-aware.

- *<SAPDATA_HOME>*

Directory where SAP R/3 database files are installed. This is an optional parameter. By default, it is set to *<ORACLE_HOME>*.

Using the GUI

To configure an instance of the SAP R/3 Database Server, perform the following steps using the Data Protector GUI:

1. In the `Data Protector Manager`, switch to the `Backup` context.
2. In the Scoping Pane, expand `Backup`, then expand `Backup Specifications`, and right-click `SAP R/3`.
3. Click `Add Backup`. In the `Create New Backup` dialog box, double-click the `Blank SAP Backup` template or any of the pre-defined templates.

   The properties of a particular backup template can be seen in the corresponding pop-up window.
4. In the Results Area of the next page of the wizard, enter the following information:

   - Name of the SAP R/3 Database Server you want to configure. If the application is cluster-aware, select the virtual server of the SAP R/3 resource group.
   - Name of the Oracle Server instance (`ORACLE_SID`) on which the SAP R/3 Database Server is running.
   - On UNIX, enter also the UNIX user name and user group of the SAP R/3 user, as described in “Configuring an SAP R/3 User in Data Protector (UNIX Systems Only)” on page 165.

Figure 2-5 **Specifying the SAP R/3 Database Server and the Oracle SID on Windows**

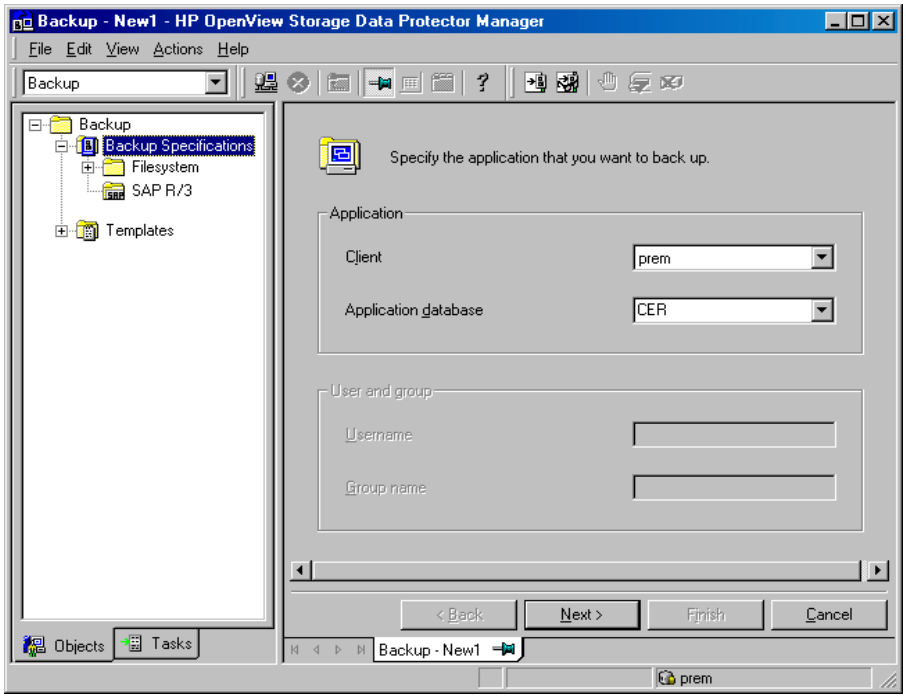

Figure 2-6 **Specifying the SAP R/3 Database Server and the Oracle SID on UNIX**

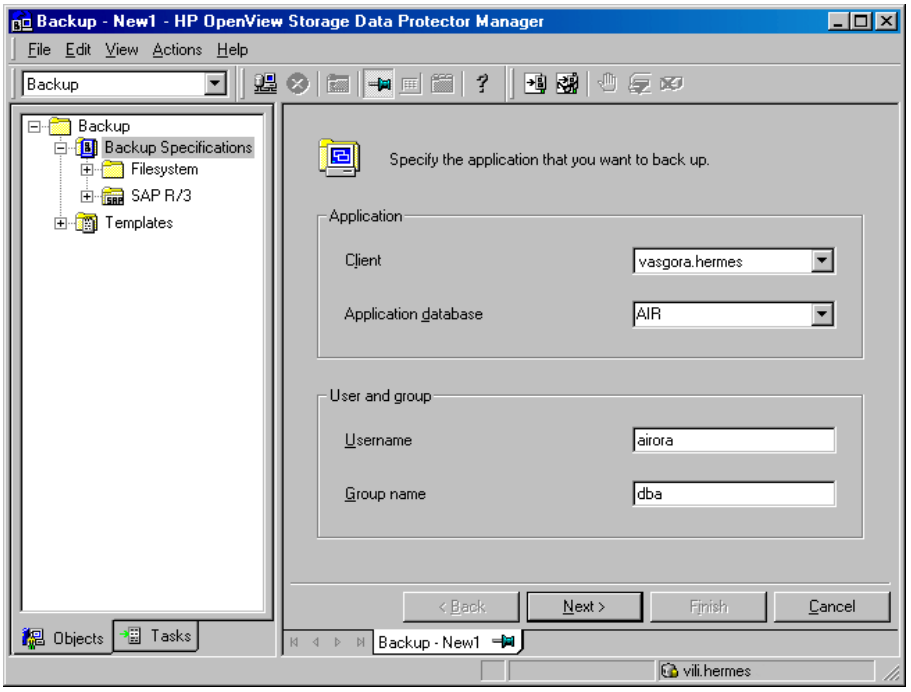

Once you have provided the required information, click `Next`. If the selected system is configured for the first time, the configuration window is displayed.

Figure 2-7 Configuring an SAP R/3 Database Server on Windows

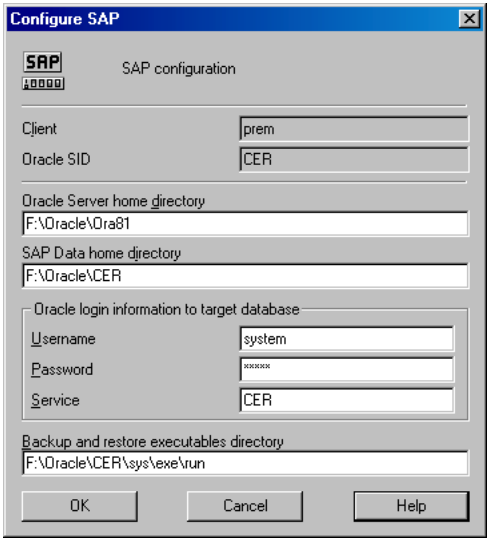

Figure 2-8 Configuring an SAP R/3 Database Server on UNIX

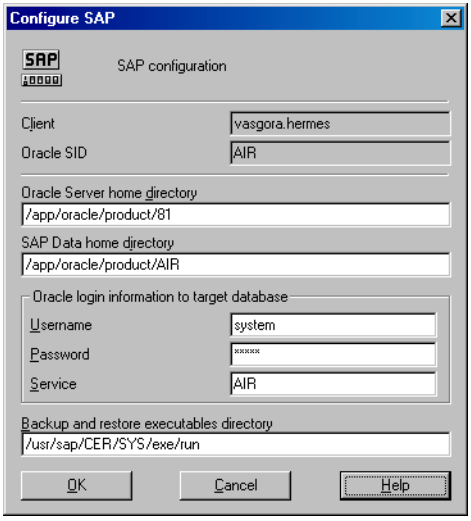

5. Enter the following information in the `Configure SAP` dialog box:
   - The Oracle Server home directory. If not specified, this is set to the default Oracle home directory.
   - SAP data home directory (if not entered, this is set to *<ORACLE_HOME>*)
   - The connection string to the Oracle Target Database.

     See "Glossary" on page G-1 for more information on login connection strings.
   - The directory where the SAP R/3 backup utilities are stored. By default, the utilities reside in the `\\<SAP_system>\sapmnt\`*<ORACLE_SID>*`\sys\exe\run` (Windows systems) or `/usr/sap/`*<ORACLE_SID>*`/SYS/exe/run` (UNIX systems) directory.

What Happens?

The following happens after saving the configuration.

Data Protector starts the `util_sap.exe` file on the SAP R/3 Database Server, which performs the following:

1. Saves the configuration parameters in the Data Protector integration configuration on the Cell Manager in the `/etc/opt/omni/server/integ/config/SAP/`*<client_name>*`%`*<ORACLE_SID>* file (UNIX Cell Manager), or in the *<Data_Protector_home>*`\Config\server\integ\config\sap\`*<client_name>*`%`*<ORACLE_SID>* file (Windows Cell Manager).
2. On UNIX, it creates a soft link for backint from the directory in which SAP R/3 utilities are stored to `/opt/omni/lbin` (HP-UX and Solaris systems) or `/usr/omni/bin` (other UNIX systems).
3. On Windows, copies the backint program from the *<Data_Protector_home>*`\bin` directory to the directory in which the SAP R/3 backup utilities reside.

Checking the SAP R/3 Configuration - Data Protector GUI

To check the configuration of your SAP R/3 Database Server, proceed as follows:

1. Right-click the SAP R/3 Database Server system.
2. Click `Check Configuration`.

   If the configuration is successful, you should receive a message confirming that the integration was properly configured.

NOTE The target database must be online during the check.

The configuration can also be also checked if you have already created and saved a backup specification for a particular SAP R/3 Database Server. Proceed as follows:

1. In the `Data Protector Manager`, switch to the `Backup` context.

   In the Scoping Pane, `expand Backup, Backup Specification`, then `SAP R/3`.

2. In the Results Area, double-click the backup specification, then select `Properties`.

3. In the `Source` property page, right-click the name of the SAP R/3 Database Server, then click `Check Configuration`.

Figure 2-9 Checking the SAP R/3 Configuration

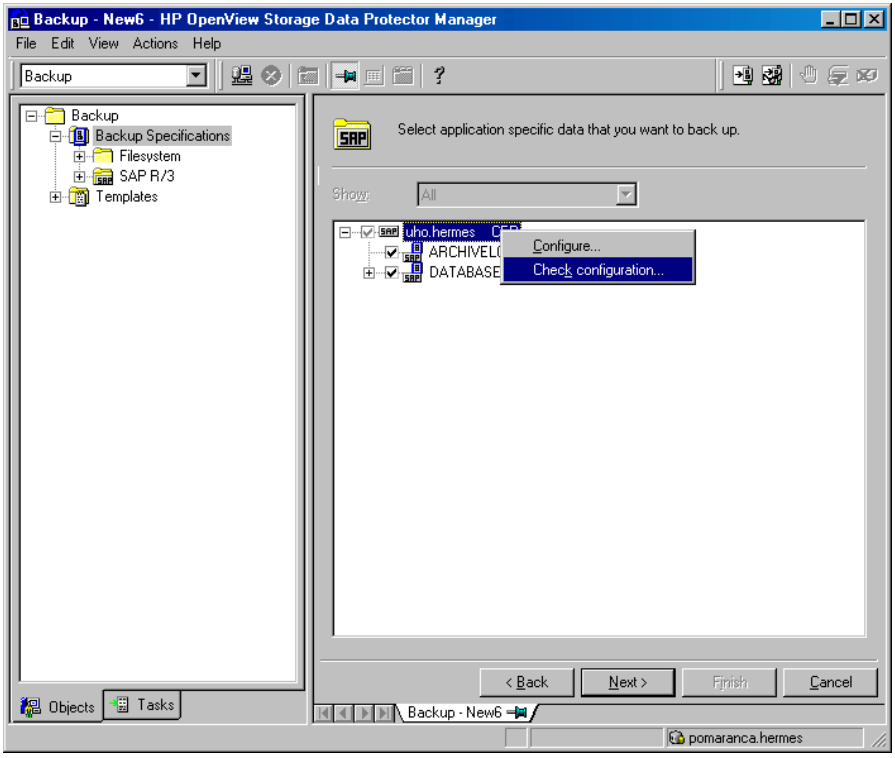

You can also (re)configure an SAP R/3 Database Server by right-clicking it and selecting `Configure`.

Checking the SAP R/3 Configuration - Data Protector CLI

To check the SAP R/3 configuration, start the following command on the client:

`util_sap.exe -CHKCONF <ORACLE_SID>`.

Data Protector verifies the configuration by attempting to connect to the SAP R/3 Database Server using the information that was specified and saved during the configuration.

In case of an error, the error number is displayed in the form *RETVAL*<*error number*>.

On UNIX, to get the error description, start the `/opt/omni/lbin/omnigetmsg 12 <error number>` (HP-UX and Solaris systems) or `/usr/omni/bin/omnigetmsg 12 <error number>` (other UNIX systems) command.

Configuring an SAP R/3 Backup

To configure an SAP R/3 backup, perform the following steps:

1. Configure the devices you plan to use for a backup. See the online Help index keyword “configuring devices” for instructions.
2. Configure media pools and media for a backup. Refer to the online Help index keyword “creating media pools” for instructions.
3. Create a Data Protector SAP R/3 backup specification. See “Creating a Data Protector SAP R/3 Backup Specification” on page 180.
4. Create or modify the parameter file on the SAP R/3 Database Server. See “Creating or Modifying the Parameter File on the SAP R/3 Database Server” on page 190.
5. If you plan to use Recovery Manager for backup, you must do some additional configuration steps. See “Backing Up Using Recovery Manager” on page 191.

Creating a New Template

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

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

To create a new backup template, proceed as follows in the `Data Protector Manager`:

1. In the `Context List`, select `Backup`.
2. In the Scoping Pane, expand `Templates` and right-click `SAP R/3`.
3. Click `Add Template`. Follow the wizard to define the appropriate backup options in your template.

   You can also modify any of the existing pre-defined templates.

Creating a Data Protector SAP R/3 Backup Specification

To create an SAP R/3 backup specification, proceed as follows:

1. In the `HP OpenView Storage Data Protector Manager`, switch to the `Backup` context.
2. In the Scoping Pane, expand `Backup`, and then `Backup Specifications`. Right-click `SAP R/3` and then click `Add Backup`. The `Create New Backup` dialog box is displayed.
3. In the `Create New Backup` dialog box, double-click `Blank SAP Backup` to create a backup specification without predefined options, or use one of the pre-defined templates given below:

Brarchive_CopyDeleteSave	Creates a second copy of the offline redo logs, saves them, deletes them after the backup, and then archives the newly-created redo logs.
Brarchive_Save	Backs up the offline redo logs.
Brarchive_SaveDelete	Backs up the offline redo logs, and then deletes them after the backup.
Brarchive_SecondCopyDelete	Creates a second copy of the offline redo logs that have been already archived, and then deletes them after the backup.
Brbackup_Offline	Backs up the shut-down database using backint.
Brbackup_Online	Backs up the active database. The util_file device type is used for backup. All tablespaces are in backup mode (locked) for the duration of the whole backup session. Whole database, particular tablespace or datafile can be backed up using this template.
Brbackup_Util_File_Online	Backs up the active database.

	Each tablespace is switched into backup mode just before the backup and is switched out from backup mode immediately after the backup. As a result, the increase in archived log files is smaller compared to the backup with the util_file device type. However, if the database consists of a large number of small files, the backup can take longer.
Brbackup_RMAN_Offline	Backs up the shut-down database using Oracle RMAN.
Brbackup_RMAN_Online	Backs up the active database. The tablespace is locked for the time of the whole backup using Oracle RMAN.

Click `OK`.

4. In the Results Area, provide the following information:

   - In the `Client` drop-down list, select the SAP R/3 Database server that you want to back up. If the application is cluster-aware, select the virtual server of the SAP R/3 resource group (on Windows) or package (on UNIX).
   - In the `Application database` drop-down list, select the name of the Oracle Server instance (`ORACLE_SID`) on which the SAP R/3 Database Server is running.
   - On UNIX, enter also the SAP R/3 user name and its group name as described in “Configuring an SAP R/3 User in Data Protector (UNIX Systems Only)” on page 165.

   Click `Next`.

5. If the SAP R/3 Database Server is already configured, the `Source` dialog box is displayed. Otherwise, you are prompted to configure it. See “Configuring an SAP R/3 Database Server” on page 167 for details.

6. In the `Source` property page, select the database objects you want to back up. Database objects include archive logs, tablespaces, and data files.

Figure 2-10 Selecting Backup Objects

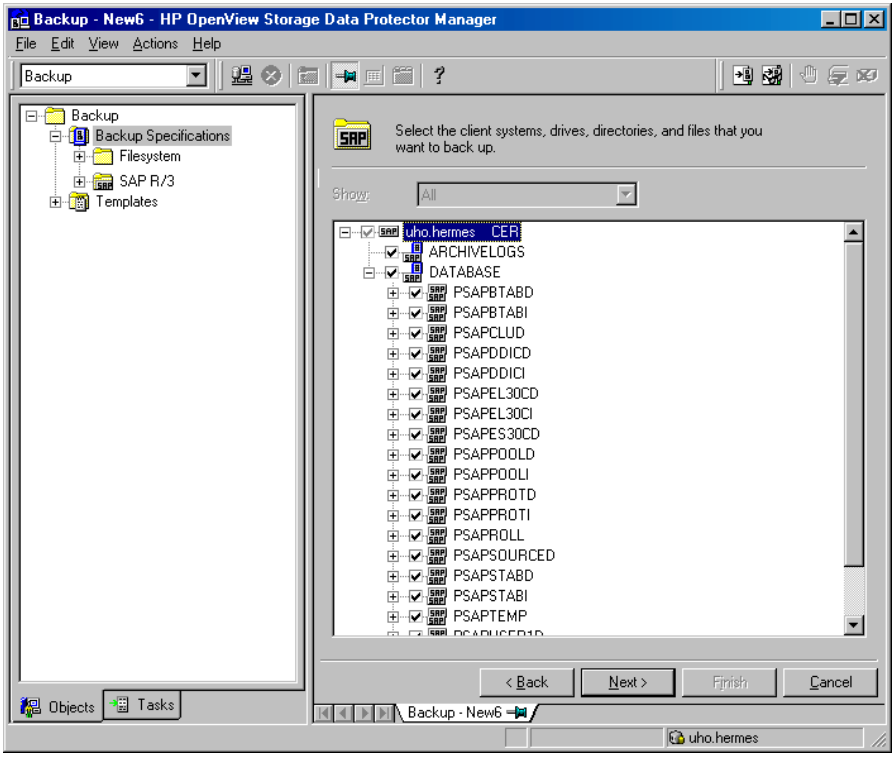

See “Why Archive Redo Logs?” on page 186, for an explanation of the reasons for archiving redo logs, and online Help for details on backup objects.

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.

8. Select the backup options.

   For information on the Backup Specification Options and Common Application Options, refer to the online Help.

   For information on the Application Specific Option (SAP R/3 specific backup options), see "SAP R/3 Specific Backup Options" on page 186 or online Help.

NOTE

The SAP R/3 backup options specified here override the current settings in the init<ORACLE_SID>.sap file.

If you have selected the Blank SAP Backup template and you do not specify any SAP R/3 backup options, the current settings in the init<ORACLE_SID>.sap file define the backup type. In this case, if backup_dev_type=rman_util, ensure that the rman_channels and rman_parms parameters are also specified. For more information, see "Backing Up Using Recovery Manager" on page 191.

Figure 2-11 SAP R/3 Backup Options

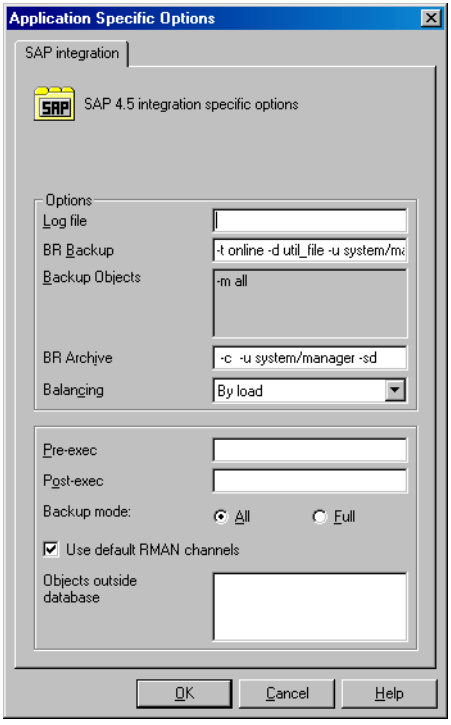

Click `Next`.

9. Optionally, schedule the backup. For more information, refer to “Scheduling a Backup” on page 198.

   Click `Next`.

10. Save the backup specification. It is recommended that you save all SAP R/3 backup specifications in the `SAP` group.

Figure 2-12 Saving the Backup Specification

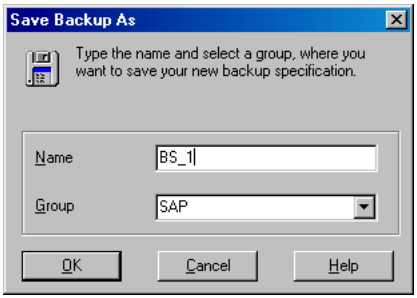

Click OK.

To start the backup, see "Backing Up an SAP R/3 Database" on page 197.

11. On UNIX, after the backup specification is saved, verify that the owner of the backup specification is the specified SAP R/3 user. See "Configuring an SAP R/3 User in Data Protector (UNIX Systems Only)" on page 165 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:
    - On UNIX:
      /etc/opt/omni/server/barlists/sap/*<Backup_Spec_Name>*
    - On Windows:
      <Data_Protector_home>\Config\server\Barlists\SAP*<Backup_Specification_Name>*
13. It is recommended to test the backup specification. See "Testing the Integration" on page 195 for details.

When the backup specification is saved, the SAP configuration, which stores information about parallelism and balancing types, is also automatically saved on the SAP R/3 Database Server.

Note that you can edit backup specifications once you have specified all the backup options.

NOTE

The parallelism of a backup (the number of streams your SAP R/3 database is backed up with) is set automatically. If load balancing is used, the parallelism represents the sum of the device concurrencies defined in the SAP R/3 backup specification. For more information on load balancing, see the online Help index “load balancing”.

The database system of an SAP R/3 system must operate in the `ARCHIVELOG` mode. This prevents the overwriting of online redo log files that have not yet been saved. To protect the archived directory from overflowing, empty the directory regularly.

Why Archive Redo Logs?

The reasons for archiving redo log files are listed below:

- In the event of a failure, consistent database status can only be recovered if all the relevant redo log files are available.
- An online backup of data files is useless if the related redo log files are missing. It is therefore necessary to archive the redo log files generated during the online backup immediately after running BRBACKUP.

SAP R/3 Specific Backup Options

The SAP R/3 specific backup options are specified using the Data Protector GUI in the `Application Specific Options` window. The window can be accessed from the `Options` property page of an SAP R/3 backup specification by clicking the `Advanced` tab.

Log file

Specifies the pathname of the backint log file. By default, this log file is not generated, as Data Protector stores all relevant information about backup sessions in the database. However, the user may decide to enable local logging by specifying a log file pathname.

BR Backup

Enter the BRBACKUP command options. See the *SAP R/3 Online Documentation* for information about

BRBACKUP command options. For example, type `-t online`, for online backup.

Or, type `-u <user_name>` for some other user than default user (usually the user `system`).

Backup Objects

When the backup specification is saved, this field lists the string passed by `omnisap.exe` to the BRBACKUP command.

BR Archive

Enter the BRARCHIVE command options. See the *SAP R/3 Online Documentation* for information about BRARCHIVE command options.

Balancing: By Load

Groups files in subsets by size so that the amount of data on all backup devices is approximately the same. Each subset is backed up by one Data Protector `sapback` program, thus allowing concurrent backup of all subsets.

If this option is set and your backup device uses hardware compression, the size of the backed up file on the medium will not be the same as on the disk. To make Data Protector aware of this, make sure that you specify the size of the backed up file on the medium in the `compression` section of the Data Protector SAP R/3 configuration file. See “Data Protector SAP R/3 Configuration File” on page 158 for information on how to do this.

Balancing: By Time

Groups files in subsets so that backup to all backup devices takes approximately the same time. This depends on the file types, the speed of the backup devices, and external

	influences (such as mount prompts), and is therefore best for environments with large libraries of the same quality. Each subset is backed up by one Data Protector `sapback` program, thus allowing concurrent backup of all subsets of the same type. Data Protector automatically stores backup speed information in the `speed` section of the Data Protector integration configuration file on the Cell Manager. It uses this information to optimize backup time. This type of balancing may lead to non-optimal grouping of files in the case of online backup, or if the speed of backup devices varies significantly among devices.
Balancing: Manual	Manual balancing optimizes backups by allowing you to group files into subsets and back up these subsets using specific devices. See "Manual Balancing of Files into Subsets" on page 193 for more information.
Balancing: None	No balancing is used. The files are backed up in the same order as they are listed in the internal Oracle database structure. To check the order use the Oracle Server Manager SQL command: `select * from dba_data_files`
Pre-exec	Specifies an object pre-exec command with options that will be started on the SAP R/3 Database Server before backup. The command/script is started by Data Protector `omnisap.exe` and has to reside in the *<Data_Protector_home>*`\bin` (Windows systems), `/opt/omni/bin`

	(HP-UX and Solaris systems), or `/usr/omni/bin` (other UNIX systems) directory. Only the filename must be provided in the backup specification.
Post-exec	Specifies an object post-exec command with options that will be started on the SAP R/3 Database Server after backup. The command/script is started by Data Protector `omnisap.exe` and has to reside in the *<Data_Protector_home>*`\bin` (Windows systems), `/opt/omni/bin` (HP-UX and Solaris systems), or `/usr/omni/bin` (other UNIX systems) directory. Only the filename must be provided in the backup specification.
Backup mode	Specifies the type of RMAN backup to be used. This option is disabled if tablespaces and not the whole database are configured to be backed up. If `All` is specified, RMAN backs up the complete database. If `Full` is specified, RMAN performs the Full backup (level 0), thus enabling RMAN incremental backups.
Use default RMAN channels	Enter the concurrency value for your backup. This number overrides the parameter set in the initialization parameter file. This option is valid only if SAP R/3 uses RMAN for backing up the Oracle Target database.

Objects outside database

With this option, you save all non-database files of the SAP R/3 and Oracle environments. This means that the following directory trees can be saved:

`/sapmnt/<ORACLE_SID>`
`/usr/sap/<ORACLE_SID>`,
`/usr/sap/trans/<ORACLE_HOME>`

It is recommended that you save these directories in a separate backup session.

NOTE

Note that the `sapdata<n>` and `saplog` or `origlog/mirrlog` subdirectories of the <*SAPDATA_HOME*> directory should not be saved.

See online Help for details on other specific Data Protector backup options.

Creating or Modifying the Parameter File on the SAP R/3 Database Server

The parameter file is used by SAP R/3 to set specific SAP R/3 backup options in case these options are not yet specified using the backup command. A template for the parameter file is located on the SAP R/3 Database Server as:

- On UNIX: `<ORACLE_HOME>/dbs/init<ORACLE_SID>.sap`

  where `<ORACLE_SID>` represents the identifier for your database.
- On Windows: `<ORACLE_HOME>\database\init<ORACLE_SID>.sap`

To link the Data Protector SAP R/3 Integration Module with the SAP R/3 backup and restore interface, modify the `backup_dev_type` parameter in the parameter file.

You can find this parameter in the following section of the parameter file:

```
# backup device type
# [disk | tape | tape_auto | pipe | pipe_auto | rman_util
| util_file_online | util_file ]
# default: tape
backup_dev_type = util_file
```

You can perform two types of online backups as well as offline backups.

- To start an offline backup, specify the `-t offline` and `-d util_file` BRBACKUP options. You can alternatively specify `backup_dev_type = util_file` and `backup_type = offline` in the SAP parameter file.
- The two types of online backups differ according to the duration in which tablespaces are in backup mode.

  If the `-t online and -d util_file` BRBACKUP options are specified, SAP R/3 puts all tablespaces in backup mode before the backup begins, and puts them back into normal mode after the backup. The same is achieved by specifying `backup_dev_type = util_file` and `backup_type=online` in the SAP parameter file.

  If the `-t online and -d util_file_online` BRBACKUP options are specified, SAP R/3 puts individual objects in backup mode before the backup begins, and puts them back into normal mode after the backup. The same is achieved by specifying `backup_dev_type = util_file_online` and `backup_type=online` in the SAP parameter file.

Refer to the SAP R/3 documentation for more information.

Backing Up Using Recovery Manager

Benefits

Version 4.5 and higher of the SAP R/3 backup utilities allows Oracle data files to be backed up using RMAN mode. RMAN mode is in general transparent to the user. The User Interface remains unchanged and allows the use of new options. The most important benefit of RMAN mode is that the underlying Oracle database can be backed up *incrementally*.

The backup procedure using RMAN mode is very similar to the one for the underlying Oracle database using the Data Protector Oracle integration. The following restrictions must be taken into account when RMAN is used directly:

- The RMAN stores information about backups in the recovery catalog. For security reasons, this catalog should be kept 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 restoration and recovery of data is complicated. It may not be possible 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 on the basis of the backups that have been made.

IMPORTANT

If the SAP R/3 integration is configured with the user `Internal`, the offline SAP R/3 backup using the RMAN fails. Configure the integration using the user `System`.

The integration of RMAN into the BRBACKUP SAP backup utility offers some important benefits:

- The recovery catalog is not used. Information about backups is saved in the control file and SAP log files. After each backup, the control file and SAP log files are saved. When data is restored, the control file is copied back first and then the data files. In case of a disaster, restore SAP log files before restoring any data files.
- Other important files will still be automatically backed up using the backint program.
- All previous SAP backup strategies can still be used with RMAN. However, RMAN cannot be used for offline redo log backups with BRARCHIVE, for standby database backups, or for split mirror backups.

Configuring the SAP R/3 RMAN Backup

To configure the SAP R/3 backup that uses the Oracle RMAN utility for backing up the Oracle Target Database data files:

1. Link the Oracle Server with the Data Protector Database Library. See “Linking Oracle with the Data Protector Oracle Integration Media Management Library (MML) on UNIX” on page 14.
2. Specify `rman_util` as a backup device type using any of the following methods:
   - Select the `Brbackup_RMAN_Offline` or `Brbackup_RMAN_Online` template when creating the backup specification.
   - Specify the `-d rman_util` BRBACKUP parameter when creating

the backup specification.

— In the SAP parameter file, set:

```
backup_dev_type = rman_util
rman_channels = <number_of_channels>
rman_parms =
"ENV=(OB2BARTYPE=SAP,OB2APPNAME=<DB_Name>,
OB2BARLIST=<Backup_Specification_Name>)"
```

Before starting an incremental backup, ensure that the appropriate full backup is done using the following option:

`-m full` (using BRBACKUP) or `backup_mode=full` (in the SAP parameter file)

Incremental Backups if Using RMAN

To start an incremental backup, specify the `Incr` mode in the Data Protector GUI or the `Incr` mode in the CLI, as follows:

```
omnib -sap_list <SAP_Backup_Specification> -barmode incr
```

Manual Balancing of Files into Subsets

Manual balancing allows you to precisely tailor the performance of an SAP R/3 backup by grouping files into subsets that are backed up in parallel. Make sure that:

- You use only one file from the same hard disk at a time.
- The number of files in a subset is equal to or smaller than a concurrency number, that is, the sum of concurrencies of all devices configured in the backup specification.
- If you do not specify all the files, other files that need to be backed up are added to the list automatically using the `load balance` option. Before backup, this list of files is logged in the *<ORACLE_HOME>*`/sapbackup/.*.lst` (UNIX systems) or *<SAPDATA_HOME>*`\sapbackup\*.lst` (Windows systems) file.
- You specify the file subsets in the `manual_balance` section of the Data Protector integration configuration on the Cell Manager in the `/etc/opt/omni/server/integ/config/SAP/`*<client_name>*`%`*<ORACLE_SID>* (UNIX systems) or

<Data_Protector_home>\Config\server\integ\config\sap\<client_name>%*<ORACLE_SID>* (Windows systems) file. See “Data Protector SAP R/3 Configuration File” on page 158.

Creating an SAP /R3 Backup Specification for Manual Balancing

To use manual balancing, you have to edit the SAP R/3 backup specifications. The backup specifications are specified in the /etc/opt/omni/server/barlist/sap directory (UNIX Cell Manager) or in the *<Data_Protector_home>*\Config\server\Barlists\SAP directory (Windows Cell Manager). In the backup specification, define which backup set will be backed up to which device. Use the -restype option followed by the ID numbers of the sets to be backed up by a specific device.

Example

To back up three subsets identified by ID numbers 1, 3, and 4, using a device named device2, specify the following:

```
DEVICE "DEVICE2"
{
  -restype "1 3 4"
}
```

Note that the files in the specified subsets are thus backed up using only the specified device. To optimize backup performance, the number of sets for a device should be equal to the concurrency of the device.

Ensure that all the subsets are specified for backup using a specific device, or they will not be backed up. To ensure that all the subsets are backed up, even if you do not specify them for backup using a specific device, configure one device without the -restype option. All the subsets not configured for backup using a specific device will be backed up on this device.

Save the backup specification before using it.

Example of Configuration

Suppose that you have two devices, *Device_1,* with concurrency 2, and *Device_2,* with concurrency 1. You also have the following manual balance specified in the manual_balance section of the Data Protector integration configuration on the Cell Manager in the /etc/opt/omni/server/integ/config/SAP/*<client_name>*%*<ORACLE_*

SID> file (UNIX systems), or in the *<Data_Protector_home>*\Config\server\integ\config\sap*<client_name>*%*<ORACLE_SID>* file (Windows systems):

```
manual_balance={
SAP-R3={
fileA=0;
fileB=1;
fileC=0;fileD=2;}}
```

Configure your backup specification `SAP-R3` to back up the files `fileA`, `fileC` and `fileD` on device *`Device_1`*, and `fileB` on device *`Device_2`*.

The backup specification then looks like:

```
BARLIST "SAP-R3"
OWNER <user> <group> galeja.zimco.com
DEVICE "DEVICE1"
{
  -restype "0 2"
}
DEVICE "DEVICE2"
{
  -restype "1"
}
CLIENT "ORACLE_SID" galeja.zimco.com
{
  -exec omnisap.exe
-args"-brb -t online -m all"
}
```

Testing the Integration

Once you have created and saved a backup specification, you should test it before running a backup.

Testing Using the Data Protector GUI

Testing Procedure

The procedure consists of checking the Data Protector part of the integration to ensure that communication within Data Protector is established, that the data transfer works properly, and that transactions are recorded either in the recovery catalog (if used) or in a control file. Proceed as follows to test the integration:

1. In the `Data Protector Manager` switch to the `Backup` context.

2. In the Scoping Pane, expand `Backup`, then `Backup Specifications`, `SAP R/3`, and right-click the backup specification you want to preview.
3. Click `Preview Backup` to open the `Start Preview` dialog box. Select the type of backup you want to run as well as the network load. See online Help for a description of these options.

Testing Using the Data Protector CLI

A test can be executed from the CLI on the SAP R/3 Database Server system or on any other Data Protector client within the same cell, provided that the systems have the Data Protector User Interface installed.

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

- On HP-UX and Solaris: `/opt/omni/bin/omnib -sap_list \ <backup_specification_name> -test_bar`
- On other UNIX: `/usr/omni/bin/omnib -sap_list \ <backup_specification_name> -test_bar`
- On Windows: `<Data_Protector_home>\bin\omnib -sap_list <backup_specification_name> -test_bar`

What Happens?

The session messages are displayed on the screen during the command execution, while the following happens:

The `omnisap.exe` program is started, which then starts the Data Protector `testbar` command. This command then checks:

- the communication within Data Protector,
- the syntax of the SAP R/3 backup specification,
- if the devices are correctly specified,
- if the required media reside in the devices.

Backing Up an SAP R/3 Database

To run a backup of an SAP R/3 database, use any of the following methods:

Backup Methods

- Schedule a backup of an existing SAP R/3 backup specification using the Data Protector Scheduler. See “Scheduling a Backup” on page 198.
- Start an interactive backup of an existing SAP R/3 backup specification using the Data Protector GUI or the Data Protector CLI. See “Running an Interactive Backup” on page 200.
- Start an interactive backup on SAP R/3 Database Servers using either the `brbackup` or the `sapdba` command. See “Using SAP R/3 Commands” on page 202.

NOTE

If you use `brbackup` or `sapdba` to start a backup session, you do not receive any Data Protector messages about the progress of the session.

Messages from the Data Protector backup session are logged in the Data Protector database. SAP R/3 messages generated by the `brbackup` or `sapdba` commands are logged to the Data Protector database only if Data Protector is used to start the backup.

Duplicate SIDs

Concurrent backups of systems with the same Oracle SID in the same cell are not supported.

Backup Modes

Configurable backup modes that were used in Data Protector versions earlier than A.03.00 are not supported in the current version of Data Protector. However, their functionality is now supported using templates.

Incremental Backups

Before starting an incremental backup, ensure that the appropriate full backup is done using the following option (note that this is valid for SAP tools version 4.5 and later):

`-m full` (using BRBACKUP) or `backup_mode=full` (in the SAP parameter file).

To start an incremental backup, specify `Incremental mode` in the Data Protector GUI or incr mode in the CLI, as follows:

```
omnib -sap_list <SAP_Backup_Specification> -barmode incr
```

NOTE It is not possible to perform an incremental backup in backint mode.

Aborting a Running Session

In the `Actions` menu, click `Abort` to abort a running SAP R/3 backup session, and then confirm the action.

Figure 2-13 Aborting an SAP R/3 Backup Session

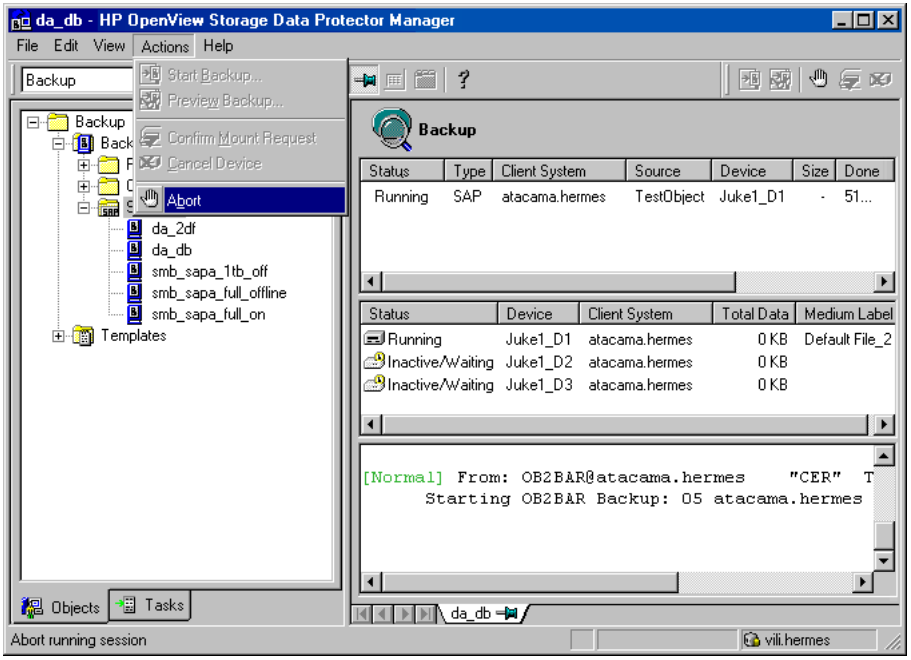

Scheduling a Backup

For more information on scheduling, refer to the online Help index keyword “scheduled backups”.

A backup schedule can be tailored according to your business needs. If you need to keep the database online continuously, then you should back it up frequently, including backup of the archived redo logs, which is required in case you need a 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.

Some examples of scheduling backups of production databases:

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

To schedule an SAP R/3 backup specification, proceed as follows:

1. In the `HP OpenView Storage Data Protector Manager`, switch to the `Backup` context.
2. In the Scoping Pane, expand `Backup`, `Backup Specifications`, and then `SAP R/3`.
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`. See Figure 2-14.

Figure 2-14 Scheduling Backups

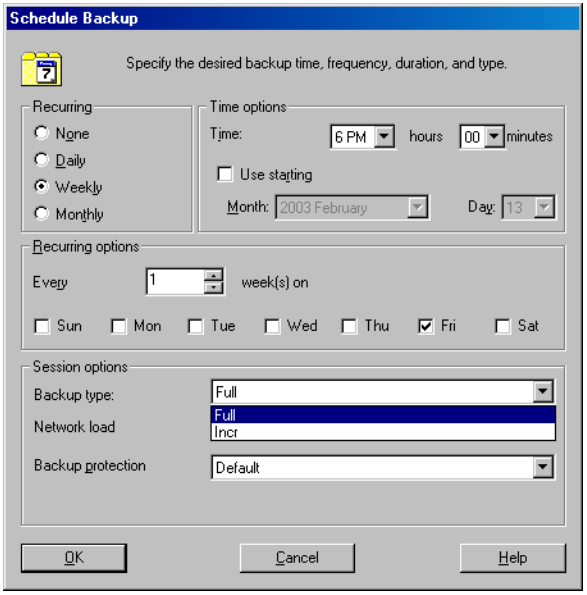

Click `OK` and then `Apply` to save the changes.

NOTE It is not possible to perform an incremental backup in the backint mode.

Running an Interactive Backup

An interactive backup can be performed any time after a backup specification has been created and saved.

Starting a Backup Using the GUI

To start an interactive backup of an SAP R/3 database using the Data Protector GUI, proceed as follows:

1. In the `HP OpenView Storage Data Protector Manager`, switch to the `Backup` context.
2. In the Scoping Pane, expand `Backup`, `Backup Specifications`, and then `SAP R/3`.
3. Right-click the backup specification and select `Start Backup`.

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

Click `OK`.

Figure 2-15 Starting an Interactive Backup

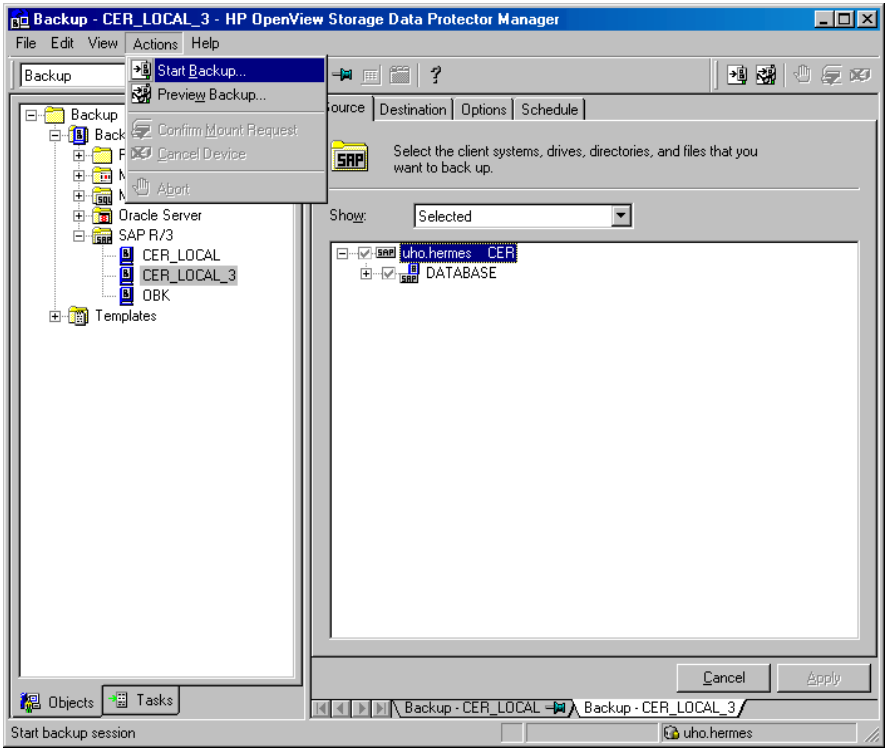

An interactive backup can also be started from the CLI.

Cluster-Aware Clients

In a cluster environment, the environment variable `OB2BARHOSTNAME` must be defined as the virtual hostname before running a backup 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>`

Tru64 Cluster

Before starting a backup on the Tru64 Cluster, 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 init SAP.sap
```

Starting a Backup Using the CLI

Switch to the `/opt/omni/bin` (HP-UX and Solaris systems), `/usr/omni/bin` (other UNIX systems), or *<Data_Protector_home>*`\bin` (Windows systems) directory on an SAP R/3 Database Server system and run the following command:

`omnib -sap_list` *<backup_specification_name>* `[-barmode` *<SapMode>*`]` [*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`

SapMode = {`-full` | `-incr`}

Refer to the `omnib` man page for details.

Example

To start a backup using an SAP R/3 backup specification called `RONA`, run the following command:

```
omnib -sap_list RONA
```

Using SAP R/3 Commands

When you interactively start a backup of your SAP R/3 object using the `brbackup` or `sapdba` commands, Data Protector uses the default SAP R/3 backup specification named `SAP-R3` for backup.

`Data Protector Inet` User Account on Windows

On Windows, before you start backup interactively using the `sapdba` command, and you have at least one device attached to the SAP R/3 Database Server (and specified in the backup specification), you have to

set the service startup account of the Data Protector Inet service to be your logon user account. This does not apply if you initiate the backup or restore using the Data Protector User Interface.

To configure the Data Protector Inet service startup account, perform the following steps:

1. In the Control Panel, go to Administrative Tools, Services.
2. Select the Data Protector Inet service and restart it.

Starting a Backup Using Another Backup Specification

To start a backup using some other SAP R/3 backup specification, you must set the environment variable OB2BARLIST to the appropriate SAP R/3 backup specification name, and OB2APPNAME to the appropriate SAP R/3 backup system ID before starting the backup.

Set the environment variable by entering the following command *before* you enter the brbackup command or sapdba command:

- On UNIX:

  export OB2BARLIST=<*backup_specification_name*>

  export OB2APPNAME=<*ORACLE_SID*>

- On Windows:

  set OB2BARLIST=<*backup_specification_name*>

  set OB2APPNAME=<*ORACLE_SID*>

If you do not set this environment variable, Data Protector assumes that the SAP R/3 backup specification is named SAP-R3.

Restoring an SAP R/3 Database

You can restore SAP R/3 databases in any of the following ways:

- Use the Data Protector GUI or CLI.
- Use SAP R/3 commands.

Considerations

You cannot perform a restore of backups created by the Oracle RMAN using the Data Protector GUI or CLI.

Before you start to restore your data using the Data Protector User Interface, you need detailed information about backed up objects. See the following section on how to find the information you need to restore your data.

If your disk is full before a restore, restoring of a filesystem with SAP R/3 data that was backed up using the `brbackup` command will fail, because the `brrestore` command needs additional disk space for restoring the control file and archive redo log files. How much additional disk space you need depends on the amount of the backed up data.

Localized SAP R/3 Object Names When selecting objects for restore, Data Protector displays the actual names of the files as they are written to the filesystem and not SAP R/3 names, which are displayed when selecting objects for backing up. As a result, if the names contain non-ASCII characters, some of the characters may display different as in the backup specification, depending on your system settings (code pages or locale). This does not impact restore, which is still completed successfully, except on Windows systems where DBCS is not set to the same value as the default Windows character set for non-Unicode programs. See “Troubleshooting on Windows Systems” on page 215..

Note that in UNIX, you must start the GUI in UTF-8 locale in order to be able to switch the encodings.

For example, the database encoding is set to `ja_JP.eucJP` and runs on an HP-UX system. When selecting the objects for backup, the names are displayed correctly if the Data Protector encoding is set to the same encoding as the database, that is `ja_JP.eucJP`. For restore, the

filenames are displayed correctly only if the encoding is set to the encoding of the filesytem and not with the encoding of the database (`ja_JP.eucJP`).

Limitations

- You cannot restore SAP R/3 tablespaces that are located on raw partitions using the Data Protector GUI. Instead, use SAP R/3 restore tools (for example, `brrestore`).

Finding Information Needed for Restore

To find the information needed for a restore, follow the steps below:

Execute the following commands:

- `omnidb -sap`

  to get a list of SAP R/3 objects.

- `omnidb -sap` *<object_name>*

  to get details on a specific object, including the SessionID.

Restoring Using the Data Protector GUI

To restore the SAP R/3 objects using the Data Protector GUI, proceed as follows:

1. In the `HP OpenView Storage Data Protector Manager`, switch to the `Restore` context.
2. In the Scoping Pane, expand `Restore Objects`, `SAP R/3`, and then select the SAP R/3 Database server from which you want to restore. A list of backed up objects is displayed in the Results Area. See Figure 2-16.

Figure 2-16 Restoring SAP R/3 Database Objects

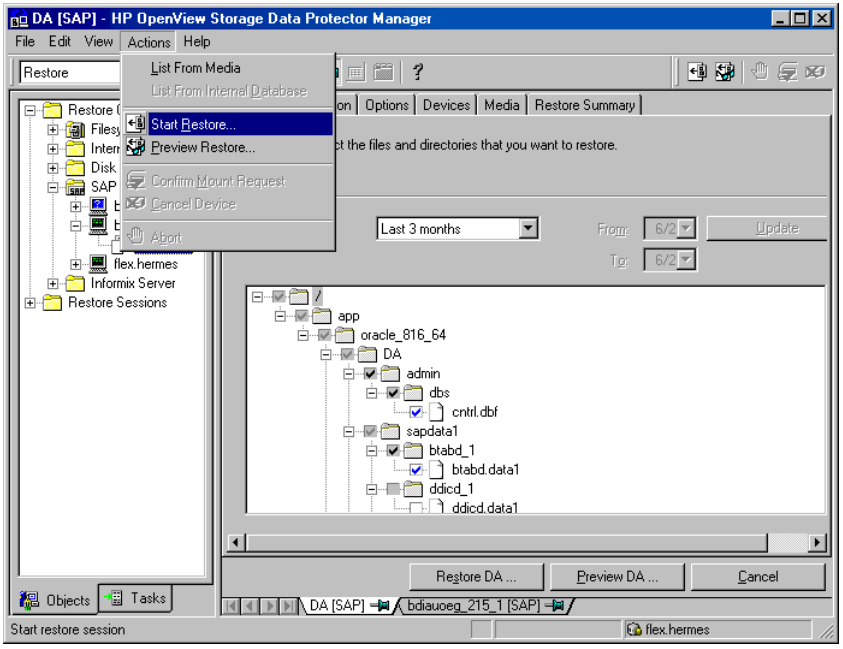

3. Select the backed up SAP R/3 object you want to restore.

   You can also select the search interval for browsing object versions in the Data Protector database by clicking the drop-down list button of the `Search Interval` option. If you select `Interval` in the drop-down list, you can set your own search interval by specifying the `From:` and `To:` options and then clicking the `Update` button.

4. Select the media and devices needed for the restore.

5. Click `Restore` and then `Finish` to start the restore session, or click `Next` to select the `Network Load` and `Report Level` before starting the restore session.

Restoring Using the Data Protector CLI

Cluster-Aware Clients

In a cluster environment, the environment variable `OB2BARHOSTNAME` must be defined as the virtual hostname before starting a restore procedure 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>`

Localized SAP R/3 object names

If the names of backed up objects contain characters that cannot be displayed using the current language group (on Windows) or code page (on UNIX):

✓ Set the environment variable `OB2_CLI_UTF8` to `1`.

✓ ***Windows only:*** Set the encoding used by the terminal to UTF-8.

If not set, backup objects returned by the Data Protector CLI commands (for example `omnidb`) may not be usable when providing the parameters to other Data Protector commands (for example `omnir`).

The `omnir` Command

Using the CLI, execute the following command:

```
omnir -sap <Host:Set> -session <SessionID> [-copyid
<CopyID>] -tree <FileName>
```

where *`FileName`* must be specified as follows:

Windows: Full pathname of the file must be specified using UNIX syntax (slashes and not backslashes) starting with the root directory (`/`), drive letter, and colon. For example: `-tree /c:/oracle/log.dbf`.

UNIX: Full pathname of the file must be specified. For example: `-tree /app/oracle/log.dbf`.

Provide the *`SessionID`* of the backup session. If you want to restore from a specific object copy, provide also the *`CopyID`*, which selects the specific object copy (object mirror or object copy) to be used for restore. By default (if the `-copyid` option is not specified), Data Protector selects the media set to restore from automatically.

Examples

Windows:

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

UNIX:

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

The restore session can be monitored in the `Data Protector Monitor` window, where mount prompts for the required media are also displayed. Refer to the man pages for more information on the Data Protector `omnir` command.

TIP

If you have a sparse file, restore using the `sparse` option to perform a faster restore.

Use any of the following methods to set the `sparse` option:

- Execute the following command: `export OB2SPARSE=sparse` (UNIX systems) or `set OB2SPARSE=sparse` (Windows systems) if the restore is started using the SAP `sapdba` or `brrestore` commands.
- Set `Restore Sparse Files` in the `Restore Options` window if the restore is started using the Data Protector GUI.
- Set restore option `-sparse` if the restore is started using the Data Protector `omnir` command.

Restoring Using the SAP R/3 Commands

The `sapdba` or `brrestore` Commands

You can use `sapdba` or `brrestore` to restore the target database. Both commands use the Data Protector backint interface to restore files backed up using Data Protector.

Prior to restoring the target database, set the `OB2APPNAME` variable:

`export OB2APPNAME=<`*`ORACLE_SID`*`>` (Unix systems)

`set OB2APPNAME=<`*`ORACLE_SID`*`>` (Windows systems)

If you have backups of two different Oracle Servers with the same `ORACLE_SID` but on different SAP R/3 Database Servers, set the `OB2HOSTNAME` variable before starting restore to the name of the SAP R/3 Database Server from which you want to restore:

`export OB2HOSTNAME=<`*`client_name`*`>` (UNIX systems) or `set OB2HOSTNAME=<`*`client_name`*`>` (Windows systems)

See the *SAP R/3 System Online Documentation* for instructions on how to use the `sapdba` or `brrestore` utilities.

Using Another Device

Data Protector supports restore using a device other than the one that was used at backup time.

Restoring Using the Data Protector GUI

If you are performing a restore using the Data Protector GUI, see the online Help index “selecting, devices for restore” for more information on how to perform a restore using another device.

Restoring Using the Data Protector CLI or SAP R/3 Commands

If you are performing a restore using the Data Protector CLI or SAP R/3 commands, specify the new device 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 is a new device.

Note that this file should be deleted after it is used. On Windows, it has to be in UNICODE format.

Example

Suppose you have SAP R/3 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. See also the *HP OpenView Storage 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 troubleshooting section.
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.

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 restoring any other part of the database.

Perform the following steps:

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

   The control files (`ctrl<ORACLE_SID>.dbf`) are by default restored to the directory defined by the SAPBACKUP variable. If the variable is not set, the control files are restored to the following directories:

   - `/var/opt/omni/tmp` (HP-UX and Solaris systems),
   - `/usr/opt/omni/tmp` (other UNIX systems), or
   - `<Oracle_home>\tmp` (Windows systems).

2. Run the following script:

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

Figure 2-17

Where *<TMP_FILENAME>* is the folder to which the control file was restored.

Monitoring an SAP R/3 Backup and Restore

The Data Protector GUI enables you to monitor current or previous backup and restore sessions.

NOTE

Only the Data Protector users in the `Admin` group and those granted the `Monitor` user rights are given access to the Data Protector monitoring functionality.

Monitoring is automatically activated when you start a restore or backup.

Monitoring Current Sessions

During a backup, system messages are sent to both the SAP R/3 Database Server and the Data Protector monitor. Thus, you can monitor a backup session from either the SAP R/3 Database Server or from any Data Protector client in the network where the User Interface is installed.

When it is detected that no more data can be backed up on the media, either because they are not in a device or because they are full, and a mount prompt is issued, the message is sent to the Data Protector monitor only, not to SAP R/3. Change the media and confirm the mount prompt in Data Protector.

To monitor a currently running session using the Data Protector GUI, proceed as follows:

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, proceed as follows:

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.

For detailed information on a completed or aborted session, see “Viewing Previous Sessions”.

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.

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 OpenView Storage 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 OpenView Storage 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.

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 `svrmgrl` 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 `svrmgrl` binary, rather it changes its permissions. Once the `svrmgrl` binary is renamed, Data Protector will use Oracle9i `sqlplus`, as it should, to complete the operations correctly.

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 2-18 Checking the Status of the Oracle Listener

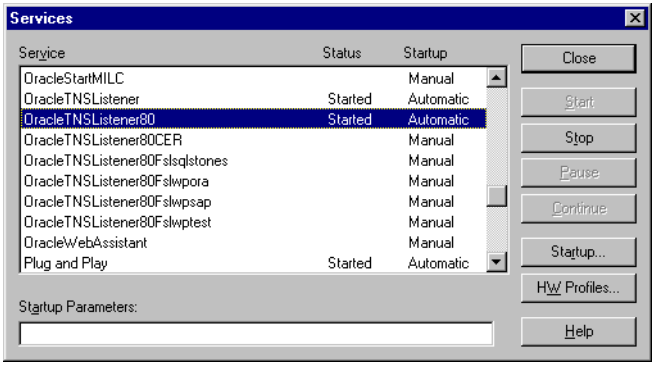

a. The status of the respective listener service in the `Services` window should be `Started`, otherwise you must start it manually.

b. 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 on page 225 and start the Server Manager from the *<ORACLE_HOME>* directory:

   `bin\svrmgrl`

   At the `SVRMGR` 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 `init`*<ORACLE_SID>*`.ora` file.

Prerequisites on the SAP R/3 Side of the Integration

The following verification steps must be performed in order to verify that SAP R/3 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 `init`*<ORACLE_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 R/3 Online Help to learn how to execute backup and restore directly to disk using the SAP R/3 backup utility.

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

4. **Verify that the SAP R/3 backup tools correctly start backint (which is provided by Data Protector):**

   Move the original backint and create a test script named backint.bat in the directory where the SAP R/3 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 R/3 is properly configured for backup using backint; otherwise you have to reconfigure SAP R/3.

See “Configuring an SAP R/3 Database Server” on page 167.

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 OpenView Storage Data Protector Installation and Licensing Guide* for details.

2. **Perform a filesystem backup of the SAP R/3 Database Server:**

   Perform a filesystem backup of the SAP R/3 Database Server system so that you can eliminate any potential communication problems between the SAP R/3 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 R/3 Database Server system.

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

3. **If the SAP R/3 backup utilities are installed in a shared directory, then the inet startup parameter must be specified as described in step 4, 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 R/3 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 2-19 Checking the `Inet` Start-Up Parameters:

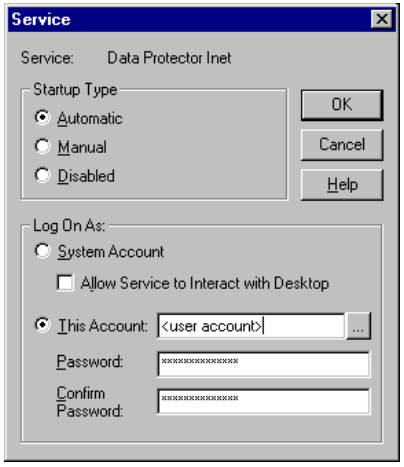

5. **Examine the environment variables:**

   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 R/3 configuration file on the Cell Manager. See "Data Protector SAP R/3 Configuration File" on page 158.

6. **Examine system errors:**

   System errors are reported in the *<Data_Protector_home>*`\log\debug.log` file on the SAP R/3 Server.

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 R/3 Server configuration:**

   To check the configuration, start the following command on the SAP R/3 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 R/3 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 R/3 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 OpenView Storage 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_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 161.
4. Restart the backup.

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 R/3 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 R/3 utilities:

Use backint, so that SAPDBA will 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.

   See “Restoring an SAP R/3 Database” on page 204.

   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 R/3 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 Oracle 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 UNIX and OpenVMS Systems” on page 112 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 on page 225 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 on page 225 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 on page 225 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 `init<ORACLE_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 on page 225 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 on page 225 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 R/3 Side of the Integration

The following verification steps must be performed in order to verify that SAP R/3 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 software with the Database Library provided by SAP R/3 (`libobk.sl`).

   b. Use the same procedure as described for linking the Data Protector Database Library.

      See "Linking Oracle with the Data Protector Oracle Integration Media Management Library (MML) on UNIX" on page 14 for information on how to do this.

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.

   c. You have to define the parameter `init` in the initialization file `init<ORACLE_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 R/3 Online Help to learn how to execute backup and restore directly to disk using the SAP R/3 backup utility.

      Check the error message and resolve this issues before you continue.

4. **Verify that the SAP R/3 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 R/3 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 SAP R/3 user; see "Configuring an SAP R/3 User in Data Protector (UNIX Systems Only)" on page 165:

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

If you receive backint arguments, this means that SAP R/3 is properly configured for backup using backint; otherwise you have to reconfigure SAP R/3.

See "Configuring an SAP R/3 Database Server" on page 167.

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 OpenView Storage Data Protector Installation and Licensing Guide* for details.

2. **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 on page 225.

   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 R/3 Database Server system so that you can eliminate any potential communication problems between the SAP R/3 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 R/3 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 R/3 configuration file on the Cell Manager. See "Data Protector SAP R/3 Configuration File" on page 158.

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 an SAP R/3 User in Data Protector (UNIX Systems Only)" on page 165.

6. **Examine system errors:**

   System errors are reported in the `/var/opt/omni/log/debug.log` (HP-UX and Solaris systems) or `/usr/omni/log/debug.log` (other UNIX systems) file on the SAP R/3 Server.

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 R/3 Server configuration:**

   To check the configuration, start the following command on the SAP R/3 Server system:

   `/opt/omni/lbin/util_sap.exe -CHKCONF <ORACLE_SID>` (HP-UX and Solaris 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 and Solaris 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 R/3 Database Server. Check the `/etc/opt/omni/client/cell_server` (HP-UX and Solaris 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 and Solaris 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 SAP R/3 backup owner as described in the “Configuring an SAP R/3 User in Data Protector (UNIX Systems Only)” on page 165 and that this user belongs to the Data Protector `operator` or `admin` group.

   b. Check that the respective Data Protector user group has the `See private objects` user right enabled.

   c. Create an SAP R/3 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 OpenView Storage 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 and Solaris 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 161.
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 `semmnu` kernel parameter. After you increase the value, rebuild the kernel and reboot the system.

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 an SAP R/3 User in Data Protector (UNIX Systems Only)” on page 165

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 and Solaris systems) or

   `/usr/omni/bin/omnidb -SAP "<object_name>" -session`

"<*Session_ID*>" -media (other UNIX systems)

on the SAP R/3 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 and Solaris systems)

/usr/omni/bin/omnidb -help (other UNIX systems)

You can also do this using the SAP R/3 utilities:

Use backint, so that SAPDBA will also use this command to query:

/opt/omni/lbin/backint -f inquiry -u <*ORACLE_SID*> -i <*input_file*> (HP-UX and Solaris 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.

   See “Restoring an SAP R/3 Database” on page 204.

   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 R/3 Database Server.

   Check the `/etc/opt/omni/client/cell_server` (HP-UX and Solaris 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 and Solaris 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 and Solaris: `/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 R/3 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 R/3 commands (for example, `brrestore`) to restore these tablespaces.

Examples of SAP R/3 Database Restore

This section describes some examples of how you can restore an SAP R/3 database. The following examples are given:

IMPORTANT

The restore of an SAP R/3 database can be performed using SAP R/3 utilities, which are not a part of Data Protector. This section only describes *examples* of how you can perform a restore using the BRRESTORE utility from SAPDBA. The examples provided do not apply to all situations, where the restore is needed. For additional information on how you can restore an SAP R/3 database using the BRRESTORE utility, refer to the SAP R/3 documentation.

Preparing the SAP R/3 Database for Restore

If you are performing a full database restore, you need to know how the backup was performed; whether you have used the Oracle RMAN channels or only BRBACKUP tools. If you used RMAN, use `svrmgrl` (Oracle8/8i) or `sqlplus` (Oracle9i), and RMAN commands to perform the restore. If you have used BRBACKUP utility, use SAPDBA to perform the restore.

If you are performing a partial restore, you can use BRRESTORE tools that come with the SAP R/3 BRBACKUP utility.

The following environment variables must be set before performing the restore:

- ORACLE_SID: system ID of the database instance

  Example: P01

SAPSID refers to the name of the SAP R/3 system, while the DBSID refers to the name of the database instance. When a single instance is installed, SAPSID and DBSID are the same.

- ORACLE_HOME: home directory of the Oracle software is by default *<Oracle_home>*\<DBSID> (Windows systems) or /opt/oracle/*<DBSID>* (UNIX systems).
- SAPDATA_HOME: home directory of the database files is by default *<Oracle_home>*\<DBSID> (Windows systems) or /opt/oracle/<DBSID> (UNIX systems).

IMPORTANT The environment variables ORACLE_SID, ORACLE_HOME and SAPDATA_HOME must always be set.

The following environment variables must only be set if the corresponding paths are different from the default locations:

- SAPARCH: directory for the BRARCHIVE logs is by default *<SAPDATA_HOME>*/saparch (UNIX systems) or *<SAPDATA_HOME>*\saparch (Windows systems).
- SAPBACKUP: directory for the BRBACKUP logs is by default *<SAPDATA_HOME>*/sapbackup (UNIX systems) or *<SAPDATA_HOME>*\sapbackup (Windows systems).
- SAPCHECK: directory for the sapdba -check/analyze logs is by default *<SAPDATA_HOME>*/sapcheck (UNIX systems) or *<SAPDATA_HOME>*\sapcheck (Windows systems).
- SAPREORG: directory for all other SAPDBA logs, as well as shell and SQL scripts is by default *<SAPDATA_HOME>*/sappreorg (UNIX systems) or *<SAPDATA_HOME>*\sappreorg (Windows systems).

  It is also the standard directory for export and unload dump files, if the parameter exireo_dumpdir in the profile init*<DBSID>*.dba is not set.
- SAPTRACE: directory for Oracle trace files and the alert file is *<SAPDATA_HOME>*/saptrace (UNIX systems) or *<SAPDATA_HOME>*\saptrace (Windows systems).
- SAPDATA1: directory of the database data files is by default *<SAPDATA_HOME>*/sapdata1 (UNIX systems) or *<SAPDATA_HOME>*\sapdata1 (Windows systems).

Syntax for SAPDATA<*n*> is: *n=1, ..., 99*. The environment variables SAPDATA<n> must only be defined if directories are on a location other than the default.

- TWO_TASK: identification of a remote database system

  This environment variable must not be set.

Other optional environment variables that can be set:

- LINES: definition of the screen height
- COLUMNS: definition of the screen width
- SAPDBA_DEBUG: setting the trace function for error analysis

Example of Full Database Restore and Recovery

To perform a full database restore and recovery, follow the steps below:

1. Login to the SAPDBA utility. In the SAPDBA select the `m` to display `User and Security` option. Select the `Expert mode` and enter the Expert's password.

Figure 2-20 Starting the SAPDBA in Expert Mode

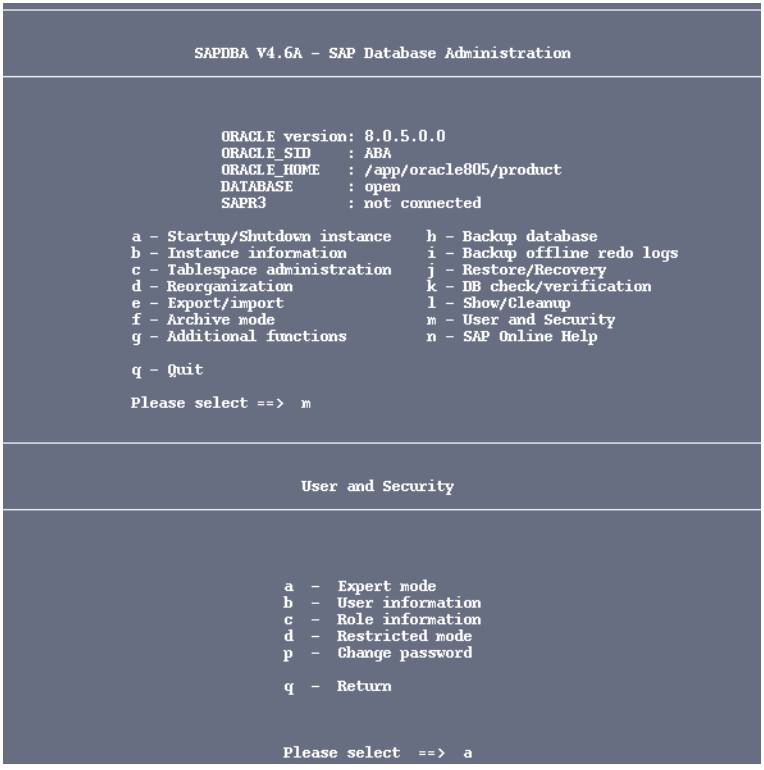

2. When the menu appears, select the `Restore/Recovery` option.

Figure 2-21 Selecting the Restore/Recovery Option

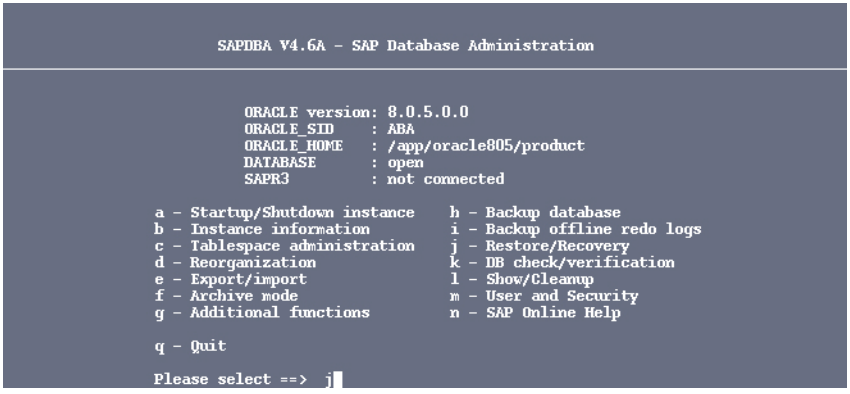

3. When the new menu appears, you can select between different types of restore. Select `Full restore and recovery` option. SAPDBA will check if your database is up and running.

Figure 2-22 Selecting Full Restore and Recovery

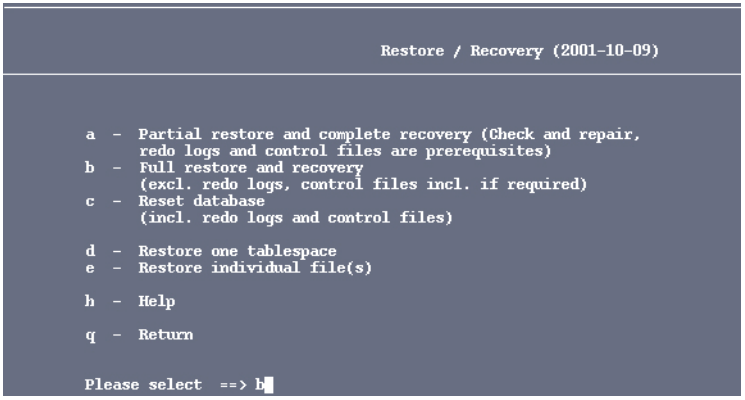

4. After the SAPDBA checks the status of the database, a new window displaying the results appears. Specify the `Select a backup of type` option to select the backup version you want to use to perform the restore.

Figure 2-23 Selecting the Backup Type and Version for Restore

5. Afterwards, enter the full pathname name for the backup tool parameter file.

6. Select the `Start restore and recover` option to start the restore session.

Figure 2-24 Starting the Restore Session

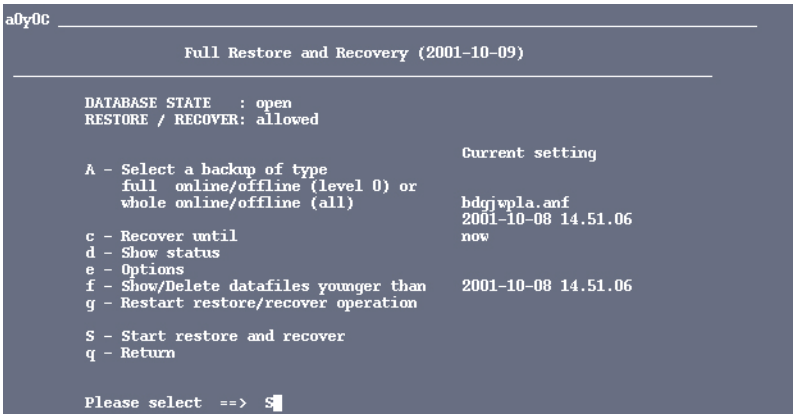

7. Select the `Return to restore procedure and continue`, if you want to specify or modify the restore parameters.

Figure 2-25 Selecting the Return to Restore Process and Continue Option

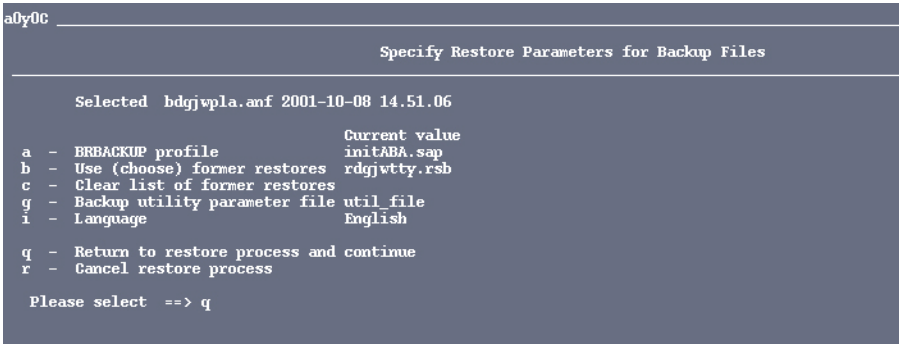

IMPORTANT If an incomplete database recovery was performed or if the control file was recovered, run the `ALTER DATABASE OPEN` command with the `RESETLOGS` option.

After you have opened a database with the RESETLOGS option, it is strongly recommended to perform a whole database backup immediately.

If the database is opened with the RESETLOGS option, the old redo log files are overwritten. Back up the offline redo log files before you open the database.

Example of Partial Restore

To perform a partial restore and recovery, you need to determine whether you need to restore a backup file or an archive redo log. The task of the SAPDBA recovery function is to fix certain media and user errors. When such errors occur, they usually involve the loss of database files, which contain many various types of objects: Oracle Dictionary segments, temporary segments, rollback segments, or user segments (tables and indexes).

SAPDBA utility supports restoring the database after the loss of the following files:

- SAP tablespaces data file (PSAP<*name*>D/I)
- System tablespace files (SYSTEM)
- Rollback tablespace files (PSAPROLL)
- Temporary tablespace files (PSAPTEMP)

The menu option Check (and repair) database only enables the recovery of the database up to the present time.

Example of Lost Files Restore

To restore the lost files, follow the steps below:

1. Define the time period within which you want SAPDBA to search for the backup files. The default value is 30 days. Then select the Start finding backup files menu option. SAPDBA utility uses the BRBACKUP log files to find the backup files.

   If the SAPDBA utility finds backup files, the necessary log sequence number is determined by SAPDBA as follows: SAPDBA searches for the most recent BRBACKUP file for each lost file and then selects the lowest of the respective log sequence numbers.

2. Select the `Show the list of damaged files` to determine the files that need to be restored.

   The SAPDBA utility lists all the lost files and their backup files. Each file shown in the list contains one of the following comments:

   - Backup file: *<name>* on *<tape/disk>*

     `Backed up by` *<name of the external backup program>*

     This means that the file was backed up using the specific program. This comment appears when the parameter `backup_util_name` of the profile `init`*<DBSID>*`.dba` contains the name of the external backup program. Otherwise, the comment is displayed as, for example: `ext. backup utility`.

   - `No restore of a backup file required`

     This means that the existing file can be used.

   - `No backup file found`

     This means that no backup was found for this file in the specified period of time.

3. Select the `Show the list of backup files` option to specify the lost files for which you would like to see the available backup files. Each file that has been lost can have several backup files.

4. Select the `Select a backup file for restore` if you would like to change the proposed backup file, that should be restored. The file that is selected for the restore is flagged with `Selected for restore`.

5. Select the `Select a BRBACKUP run for restore` if you want to change the newest found backup file for each individual file from which the requested files can be restored. You can change this setting, for example, if all the files for restore were backed up in the same backup session and you want to specify only that backup session. The following information is listed:

   - Sequential number of the backup file found
   - Coded timestamp, date and time of the backup
   - The medium on which the backup was performed
   - The number of files found in this backup which are to be restored

6. Select the `Return` option to continue with the recovery process.

The lost files are restored using the SAP utility BRRESTORE.

7. Select the `Start restore of backup files`.

   SAPDBA checks if the files that are to be restored are still available. If these files are still available, an error message is displayed. Confirm that SAPDBA may overwrite these files. If you do not allow SAPDBA to overwrite these files, the restore procedure is terminated at this point.

   SAPDBA checks if there is a backup file for each data file that was lost. If a backup file is missing, the restore procedure is terminated at this point.

   SAPDBA displays the restore parameters. The SAP utility BRRESTORE is started in order to restore the files.

Example of Archive Log Files Restore

To restore the archive log files, follow the steps below:

1. Select the `Restore archive files` option.

   Archive log files are restored using the SAP BRRESTORE utility. If SAPDBA determines that the archiving directory *<Oracle_home>*`/saparch` (UNIX systems) or *<Oracle_home>*`\saparch` (Windows systems) does not have enough space to restore all the necessary redo log files, the redo log files that have already been used will be deleted and the next required redo logs are restored during the subsequent recovery.

2. Select the `Start restore of archive files` option.

   This option is mandatory when the recovery requires offline redo log files that are no longer in the archiving directory. The recovery cannot be started until the necessary archived redo logs are restored.

   SAPDBA displays the following information on the screen:

   - The log sequence number of the first archived file to be restored.
   - The archived files that were found.
   - The maximum size of the archived redo log files.
   - The configured restore parameters which you can change using the `Specify restore parameters` option.

The SAP BRRESTORE utility restores the required files. If the redo logs are still available on the disk, they do not have to be restored.

3. Select `Return` to continue with the recovery process.

3 Integrating SAP DB/MaxDB and Data Protector

In This Chapter

This chapter explains how to configure and use the HP OpenView Storage Data Protector SAP DB/MaxDB integration. It explains the concepts and methods you need to understand in order to back up and restore SAP DB/MaxDB databases using Data Protector.

It is organized into the following sections:

Prerequisites and Limitations

This section provides a list of prerequisites and limitations you must be aware of before using the integration.

Prerequisites

- A license is needed in order to use the Data Protector SAP DB/MaxDB integration. Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide* for information about licensing.
- Before you begin, make sure that you have correctly installed and configured the SAP DB/MaxDB and Data Protector systems. Refer to:
  - *HP OpenView Storage Data Protector Product Announcements, Software Notes, and References* for an up-to-date list of supported versions, platforms, devices, limitations, and other information.
  - *HP OpenView Storage Data Protector Installation and Licensing Guide* for instructions on how to install Data Protector on various architectures and how to install the Data Protector SAP DB/MaxDB integration.
  - SAP DB/MaxDB documentation for information on the SAP DB/MaxDB Server.
- The SAP DB/MaxDB Automatic Log Backup must be activated for an SAP DB/MaxDB instance to enable transactional backup (log backup).

Limitations

Refer to the *HP OpenView Storage Data Protector Product Announcements, Software Notes, and References* for a list of general Data Protector limitations. The following is a list of integration-specific limitations:

- Instance names in UNICODE format are not supported with this integration.
- Pre- and post-exec options on the level of the backup specification are not supported with this integration.
- Preview is not possible for SAP DB/MaxDB restore sessions.

It is assumed that you are familiar with the SAP DB/MaxDB database administration and basic Data Protector functionality.

Introduction

Backup

Data Protector integrates with the SAP DB/MaxDB Database Server to offer an integrated **online backup** of **SAP DB/MaxDB instance**. The following SAP DB/MaxDB instance objects are backed up using the Data Protector SAP DB/MaxDB integration:

- **SAP DB/MaxDB data**,
- **SAP DB/MaxDB configuration** and
- **SAP DB/MaxDB archive logs**

The online backup concept is widely accepted because it addresses the business requirements of high application availability. During backup, the database is online and actively used. The backup is performed quickly and efficiently, with the least possible impact on the database performance. When an SAP DB/MaxDB database is backed up using the Data Protector SAP DB/MaxDB integration, it can be switched to either the `Admin` or to the `Online` mode, depending on the selected options.

Offline backup of SAP DB/MaxDB objects is not integrated. The standard Data Protector filesystem backup can be performed in such a case. For more information on how to perform a Data Protector filesystem backup, online Help index: “standard backup procedure“. It is also not possible to perform an integrated Data Protector SAP DB/MaxDB restore from an offline filesystem backup.

SAP DB/MaxDB data and archive logs are backed up or restored in streams, whereas the SAP DB/MaxDB 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/MaxDB Server, depending on the selected options.

The integration supports SAP DB/MaxDB backup modes, thus it is possible to perform an SAP DB/MaxDB **full backup** (**data backup**), SAP DB/MaxDB **differential backup** (**pages backup**), or SAP DB/MaxDB **transactional backup** (**log backup**).

Table 3-1 on page 253 shows what is actually backed up with regards to the selected SAP DB/MaxDB backup type and Data Protector GUI object.

Table 3-1 SAP DB/maxDB Backup Mode and Data Protector GUI Selections

		SAP DB/MaxDB 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

Restore and Recovery

At the beginning of a restore session, Data Protector switches the SAP DB/MaxDB database to the `Admin` mode. If the database cannot be switched to the `Admin` mode, an error is issued in the Data Protector monitor.

Only a *complete* SAP DB/MaxDB instance can be restored. Using the integration, SAP DB/MaxDB instances can be restored from:

- full backup sessions or
- from a combination of full, differential and transactional backup sessions.

At the end of a restore session, Data Protector switches the SAP DB/MaxDB database to either the `Online` mode or the `Admin` mode, depending on the Data Protector restore and recovery options.

With this integration, **restore** denotes the process of transferring the backed up data (data, archive logs, configuration) from backup media to the system being restored. **Recovery** denotes the process that follows the restore and includes applying **redo logs** (if present on the SAP DB/MaxDB Server) during the process of switching the database to the `Online` mode.

The integration supports SAP DB/MaxDB **migration**, meaning that an SAP DB/MaxDB instance can be restored to an SAP DB/MaxDB Server or instance other than the original. In such a case, if the SAP DB/MaxDB Server has not yet been configured for the Data Protector SAP DB/MaxDB integration, it must be configured before the restore is started. If the instance does not exist, it must be configured before the restore is started. During the migration, the existing data is overwritten and the existing redo logs are deleted.

During the restore or migration, the archive logs on the SAP DB/MaxDB Server are never deleted.

Parallelism

The integration also takes advantage of the concept of SAP DB/MaxDB media and media groups, thus providing parallel backup and restore of SAP DB/MaxDB objects. Several SAP DB/MaxDB media are grouped in an SAP DB/MaxDB media group, which is then backed up or restored in streams. This is referred to as SAP DB/MaxDB parallelism.
SAP DB/MaxDB Parallelism can be utilized only if the value of the Data Protector `Parallelism` option is equal to or lower than the sum of Data Protector concurrency values for all backup devices selected in the backup specification. See “SAP DB/MaxDB Specific Backup Options” on page 270 for more information on the Data Protector `Parallelism` option.

Permissions

Data Protector backup and restore operations on the SAP DB/MaxDB instance require the SAP DB/MaxDB `Restoring backups` (`Recovery`) and `Saving backups` (`Backup`) **permissions**, whereas Data Protector configuration operation requires also the SAP DB/MaxDB `Installation management` (`InstallMgm`) and `Parameter access` (`ParamCheckWrite`) permissions. The backup can be performed in either the SAP DB/MaxDB `Online` or `Admin` mode.

Supported Platforms

Refer to the *HP OpenView Storage Data Protector Product Announcements, Software Notes, and References* or http://www.hp.com/support/manuals for up-to-date information about platforms supported by the integration.

Advantages

Using Data Protector with the SAP DB/MaxDB Database Server offers several advantages over using SAP DB/MaxDB alone:

- Central Management for all backup operations

  You can manage backup operations from a central point. This is especially important in large business environments.

- Media Management

  Data Protector has an advanced media management system that allows you to keep track of all media and the status of each medium, set the protection for stored data, fully automate operations as well as organize and manage devices and media.

- Backup Management

  Backed up data can be duplicated during or after the backup to increase fault tolerance of backups, to improve data security and availability, or for vaulting purposes.

- Scheduling

  Data Protector has a built-in scheduler that allows you to automate backups to run periodically. With the Data Protector scheduler, the backups you configure run unattended at the periods you specify.

- Local versus Network Backups

  When configuring an SAP DB/MaxDB backup using Data Protector, the location of devices is completely transparent to the user. They can be connected to the SAP DB/MaxDB Database Server or any other Data Protector clients on the network.

- Device Support

  Data Protector supports a wide range of devices, from standalone drives to complex multiple drive libraries. Refer to the *HP OpenView Storage Data Protector Product Announcements, Software Notes, and References* or http://www.hp.com/support/manuals for an up-to-date list of supported devices and other information.

- Monitoring

  Data Protector has a feature that allows you to monitor currently running sessions and view finished sessions from any system that has the Data Protector User Interface installed.

  All backup sessions are logged in the IDB and in the SAP DB/MaxDB configuration, providing you with a history of activities that can be queried at a later time.

Integration Concept

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

The integration makes use of the SAP DB/MaxDB database management server and of the `backint` interface for SAP DB/MaxDB to perform backup and restore.

Data Protector Components

The Data Protector integration software consists of the following components:

- The `sapdbbar.exe` module, installed on the SAP DB/MaxDB Server system, which controls activities between the SAP DB/MaxDB Server and Data Protector backup and restore processes.
- The `sapdb_backint` component, installed on the SAP DB/MaxDB Server system, is a binary interface between Data Protector and backup and restore functionality of the SAP DB/MaxDB.
- The DMA (Data Mover Agent) component, installed on the SAP DB/MaxDB 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/MaxDB instance to use with Data Protector and check the instance configuration.

Supported Interfaces

With this integration, an SAP DB/MaxDB database can be backed up or restored using the following interfaces:

- Data Protector GUI or CLI
- SAP DB/MaxDB utilities

Backup Flow

When a backup session is started, the Cell Manager starts the `sapdbbar.exe` module and supplies it with the name of the instance on the SAP DB/MaxDB Server that is to be backed up. The `sapdbbar.exe` module than starts an SAP DB/MaxDB session using the

SAP DB/MaxDB `dbmcli`. The `sapdbbar.exe` module issues `dbmcli` commands that configure SAP DB/MaxDB backup media (parallelism), configure `sapdb_backint` and then start the backup using SAP DB/MaxDB `dbmcli`. SAP DB/MaxDB then starts the configured `sapdb_backint` component. For every SAP DB/MaxDB medium (pipe) `sapdb_backint` starts a DMA, which transfers the data from SAP DB/MaxDB 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.exe` module and DMA. The list of configuration files to be backed up is retrieved through `dbmcli`.

See Figure 3-1 on page 258.

NOTE When running a backup using SAP DB/MaxDB utilities, SAP DB/MaxDB media and pipes must be configured manually.

Restore Flow

When a restore session is started, the Cell Manager starts the `sapdbbar.exe` module, which starts SAP DB/MaxDB `dbmcli`. The `sapdbbar.exe` module issues commands to SAP DB/MaxDB `dbmcli` to configure `sapdb_backint` and SAP DB/MaxDB backup media (parallelism). SAP DB/MaxDB then starts the configured `sapdb_backint`, which starts streaming data to media (pipes) that SAP DB/MaxDB created. For every SAP DB/MaxDB medium (pipe) the `sapdb_backint` starts a DMA, which transfers the data from Data Protector media to SAP DB/MaxDB media (pipes). If SAP DB/MaxDB configuration is being restored, it is the `sapdbbar.exe` module and DMA that perform the restore.

See Figure 3-1 on page 258.

NOTE When running a restore using SAP DB/MaxDB utilities, SAP DB/MaxDB media and pipes must be configured manually.

Figure 3-1 SAP DB/MaxDB Backup and Restore Concept

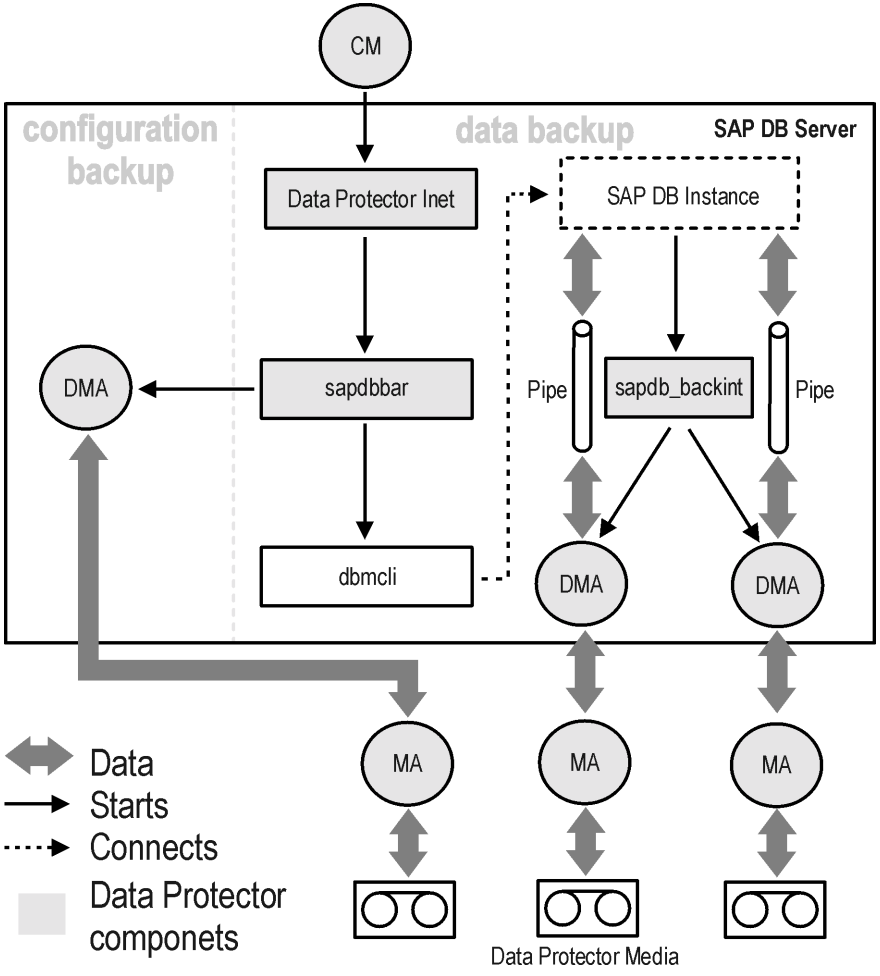

Data Protector SAP DB/MaxDB Configuration File

Data Protector stores the SAP DB/MaxDB integration parameters for every configured SAP DB/MaxDB *instance* in the following file on the Cell Manager:

- `/etc/opt/omni/server/integ/config/SAPDB/<client_name>%<instance_name>` (HP-UX and Solaris systems)
- `<Data_Protector_home>\Config\server\integ\config\SAPDB\<client_name>%<instance_name>` (Windows systems).

The parameters stored in the configuration file are those entered during the configuration of this integration, as described in "Configuring the Integration" on page 263. These parameters are:

- The username of the SAP DB/MaxDB user created or identified as described in "Configuring Users" on page 263.
- The password of the SAP DB/MaxDB user created or identified as described in "Configuring Users" on page 263.
- The SAP DB/MaxDB version.
- The SAP DB/MaxDB independent program path parameter. This parameter is the independent program path directory specified during the installation of the SAP DB/MaxDB application on the SAP DB/MaxDB Server.
- Data Protector SAP DB/MaxDB integration related environment variables.

NOTE The username and the SAP DB/MaxDB independent program path parameter must not contain the single quote character (').

The configuration parameters are written to the Data Protector SAP DB/MaxDB configuration files:

- during configuration of the integration
- during creation of a backup specification

- when the configuration parameters are changed

Configuration File Syntax

The syntax of the file is as follows:

IMPORTANT

To avoid problems with your backups, ensure that the syntax of your configuration file matches the examples.

```
Username='<username>'
Password='<password>'
Version='<SAPDB_version>' //SAP DB version
Home='<SAPDB_independent_program_directory>' //SAP DB
independent program path
```

Example of Configuration File

This is an example of the Data Protector SAP DB/MaxDB configuration file:

```
Username='dba';

Password='FHBBDHBBCHBB';

Version='7.4.3.27';

Home='/opt/sapdb/indep_prog';
```

Setting, Retrieving, and Listing Data Protector SAP DB/MaxDB Configuration File Parameters Using the CLI

Data Protector SAP DB/MaxDB configuration file parameters are normally written to the Data Protector SAP DB/MaxDB configuration files after the completed configuration of the SAP DB/MaxDB instance in Data Protector.

The `util_cmd` Command

You can set, retrieve, or list the Data Protector SAP DB/MaxDB 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 DB/MaxDB client. The command resides in the `/opt/omni/lbin` (HP-UX systems), `/usr/omni/bin/` (other UNIX systems), or in the *<Data_Protector_home>*`\bin` (Windows systems) directory.

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:

UNIX

```
export OB2BARHOSTNAME=<virtual_hostname>
```

Windows

```
set OB2BARHOSTNAME=<virtual_hostname>
```

The util_cmd Synopsis

The syntax of the util_cmd command is as follows:

```
util_cmd -getconf[ig] SAPDB <SAPDB_instance> [-local \
<filename>]

util_cmd -getopt[ion] [SAPDB <SAPDB_instance>] \
<option_name> [-sub[list] <sublist_name>] [-local \
<filename>]

util_cmd -putopt[ion] [SAPDB <SAPDB_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 which a parameter is written to or taken from.

[-local *<filename>*] specifies one of the following:

- When used with the -getconf[ig] option, it specifies a filename that the command output is written to. If the -local option is not specified, the output is written to the standard output.
- When used with the -getopt[ion], it specifies a filename of the file from which the parameter and its value are to be retrieved from and then written to the standard output. If the -local option is not specified, the parameter and its value are retrieved from the Data Protector SAP DB/MaxDB configuration file and then written to the standard output.
- When used with the -putopt[ion] option, it specifies a 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 DB/MaxDB configuration file.

Return Values

The `util_cmd` command displays a short status message after each operation (written 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 or one of the Data Protector SAP DB/MaxDB configuration files is missing on the Cell Manager.

Configuring the Integration

It is assumed that the installation of Data Protector software components on the SAP DB/MaxDB Server system was successful.

Configuration Overview

To run or schedule Data Protector SAP DB/MaxDB integration backups the following task must be performed:

- An SAP DB/MaxDB user with certain permissions must be added to the Data Protector `admin` group. See “Configuring Users” on page 263.
- Backup devices, media, and media pools must be configured. Refer to Data Protector online Help for information on how to do this.
- The SAP DB/MaxDB instance to be backed up must, as a part of the Data Protector backup specification creation, be configured to be used with Data Protector. See “Configuring an SAP DB/MaxDB Backup” on page 264.
- A Data Protector SAP DB/MaxDB integration backup specification must be created. See “Configuring an SAP DB/MaxDB Backup” on page 264.

SAP DB/MaxDB instance can be reconfigured once it has been configured to be used with Data Protector. See “Modifying the Configuration of an SAP DB/MaxDB Instance in Data Protector” on page 272.

Configuring Users

To perform a backup, restore, or other operation on an SAP DB/MaxDB Server, an SAP DB/MaxDB user granted specific SAP DB/MaxDB permissions must be configured on the SAP DB/MaxDB Server and an OS user must be added to Data Protector `admin` group. Follow the steps below to do this:

1. On the SAP DB/MaxDB Server, create or identify an SAP DB/MaxDB user with at least the following SAP DB/MaxDB permissions:
   - `Backup`
   - `Recovery`

- `InstallMgm`
- `ParamCheckWrite`

2. Add the OS user under whose account SAP DB/MaxDB is running to Data Protector `admin` group. Refer to online Help index: “adding users” for more information on how to do this.
3. Add the operating system `root` user on the SAP DB/MaxDB Server to either the Data Protector `admin` or `operator` user group. Refer to online Help, index keyword “adding users” for more information on how to do this.

Configuring an SAP DB/MaxDB Backup

To configure an SAP DB/MaxDB backup, perform the following steps:

1. Configure backup devices, media, and media pools.

   Refer to the online Help for instructions.

2. Create an SAP DB/MaxDB backup specification.

   The Data Protector backup specification is stored on the Cell Manager system and contains information needed for the Data Protector SAP DB/MaxDB integration to perform a backup.

   Once the backup specification is created and saved, you can:

   - schedule unattended backups,
   - start interactive backups, or
   - modify the backup specification.

Creating a Backup Specification

To create a backup specification for backing up SAP DB/MaxDB objects, proceed as follows:

1. In the `HP OpenView Storage Data Protector Manager`, switch to the `Backup` context.
2. In the Scoping Pane, expand `Backup`, and then `Backup Specifications`.
3. Right-click `SAP DB Server` and then select `Add Backup`. The `Create New Backup` dialog box is displayed.

4. Select the `Blank SAPDB Backup` template. With this template, no options and no scheduling is defined, everything is set to default values.

   Click `OK`.

5. In the Results Area, in the `Client` drop-down list, select the client on which the SAP DB/MaxDB Server is running. In a cluster environment, select the virtual hostname for the systems on which the SAP DB/MaxDB Server is running.

   In the `Application database` drop-down list, all the SAP DB/MaxDB instances located on the SAP DB/MaxDB Server and configured for Data Protector are listed. Enter or select the SAP DB/MaxDB instance to be configured.

   On UNIX, enter the user name and the group name for the OS user, under whose account the SAP DB/MaxDB application is running on the SAP DB/MaxDB Server (for example, the `sapdb` user in the `sapsys` group).

Figure 3-2 Selecting an SAP DB Server and Instance on UNIX Systems

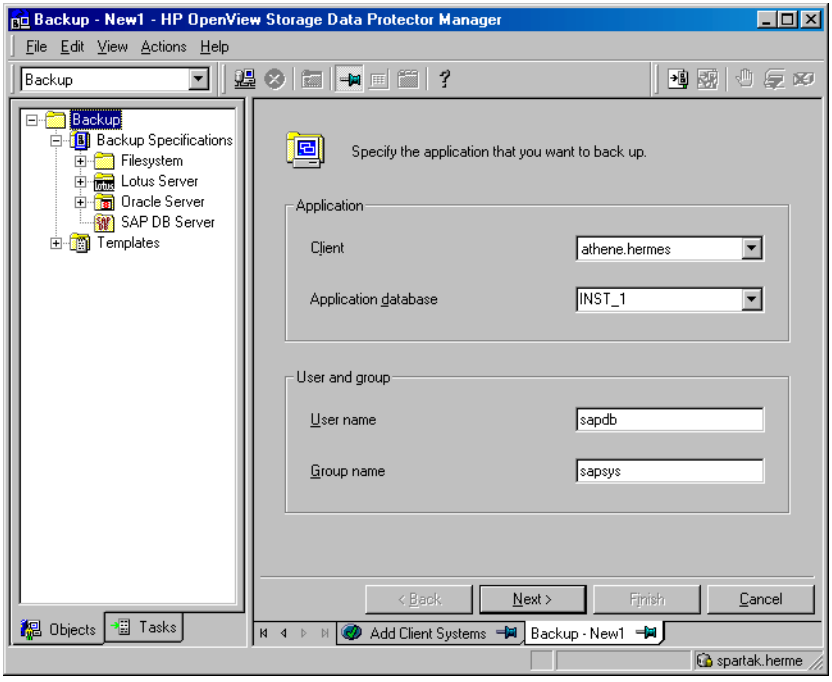

Figure 3-3 Selecting an SAP DB Server and Instance on Windows Systems

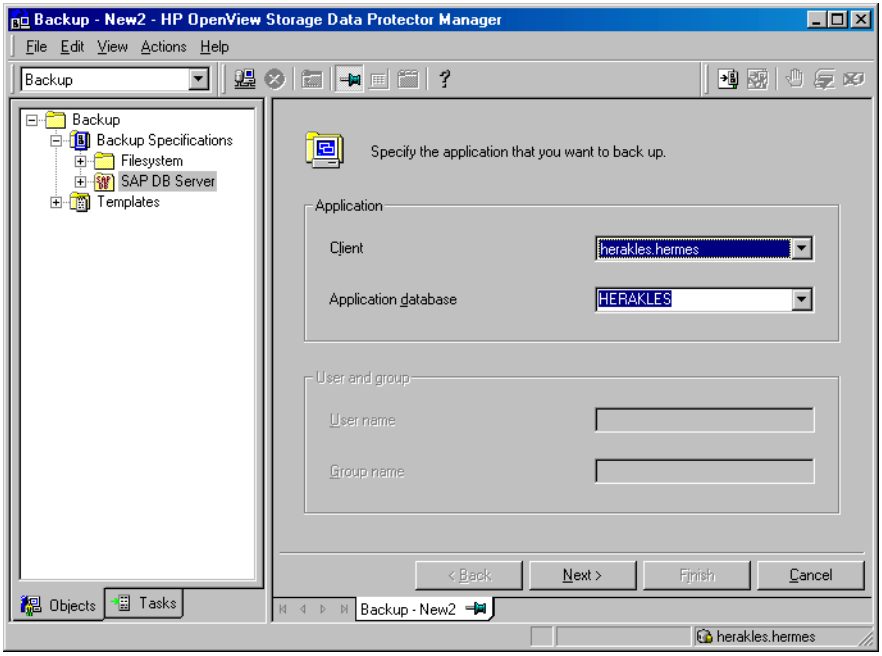

Click `Next`.

If the SAP DB/MaxDB instance you have selected had not yet been configured to be used with Data Protector, the configuration dialog box appears.

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/MaxDB application on the SAP DB/MaxDB Server. You can leave the `Auto-detect` option selected to automatically detect the directory on the SAP DB/MaxDB Server.

Enter the username and the password of the user created or identified as described in “Configuring Users” on page 263.

Click `OK` to confirm the configuration.

6. In the next step of the wizard, select the database objects you want to back up. See Table 3-1 on page 253 for more information on database objects selections.

IMPORTANT

To backup SAP DB/MaxDB archive logs, select the `Data` item in the Results Area. The archive log backup is then triggered by selecting the `Trans` backup type when scheduling the backup or running the backup interactively. See “Backing Up an SAP DB/MaxDB Database” on page 281 for more information on scheduling the backup or running the backup interactively.

Figure 3-4 Selecting SAP DB Objects

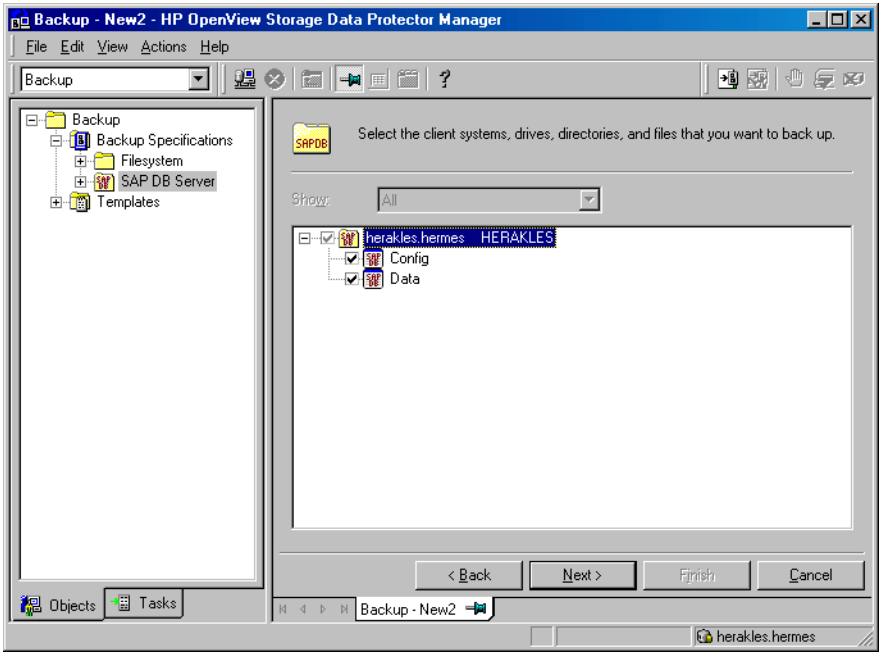

7. Follow the wizard to define devices, options, and schedule.

   Refer to the Data Protector online Help for a description of the backup devices options, backup specification options and common application options.

   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. The minimum number of devices required for mirroring SAP DB/MaxDB integration objects equals the number of devices used for backup.

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

See “SAP DB/MaxDB Specific Backup Options” on page 270 for details about the SAP DB/MaxDB backup options (application specific options).

See “Scheduling Example” on page 282 for information on how to schedule the backup specification.

8. Once you have defined all the backup options, name and save your SAP DB/MaxDB backup specification. It is recommended that you save all the SAP DB/MaxDB backup specifications in the `SAP DB Integration` group.

Figure 3-5 Saving the Backup Specification

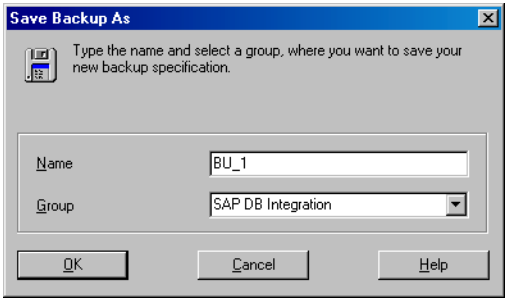

After the backup specification is saved, it can be started either from the Data Protector GUI or the Data Protector CLI, or can be scheduled to run automatically using the Data Protector Scheduler. It can also be modified. See “Backing Up an SAP DB/MaxDB Database” on page 281 for information on how to start a backup using the Data Protector GUI or the Data Protector CLI and on how to schedule a backup specification.

You can examine the newly created and saved backup specification in the Backup context. The backup specification itself is stored in the *<Data_Protector_home>*\Config\server\barlists\sapdb*<backup_specification_name>* file on Windows Cell Manager systems and in the /etc/opt/omni/server/barlists/sapdb/*<backup_specification_name>* file on UNIX Cell Manager systems.

It is recommended that you test the backup specification by clicking the Start Preview button. See "Testing the Integration" on page 278 for a step-by-step procedure. This is an interactive test that does not back up any data. However, as a result of this test, the following file is created on the SAP DB/MaxDB Server system:

- on Windows: *<Data_Protector_home>*\tmp*<Backup_Specification_Name>*_TEST_FILE
- on UNIX: /var/opt/omni/tmp/*<Backup_Specification_Name>*_TEST_FILE

It should be deleted after the test.

You can start an interactive backup that includes data transfer by clicking the Start Backup button.

SAP DB/MaxDB Specific Backup Options

The SAP DB/MaxDB specific backup options can be accessed using the Data Protector GUI by clicking the Options tab and then the Advanced button next to Application Specific Options.

Figure 3-6 Backup Options

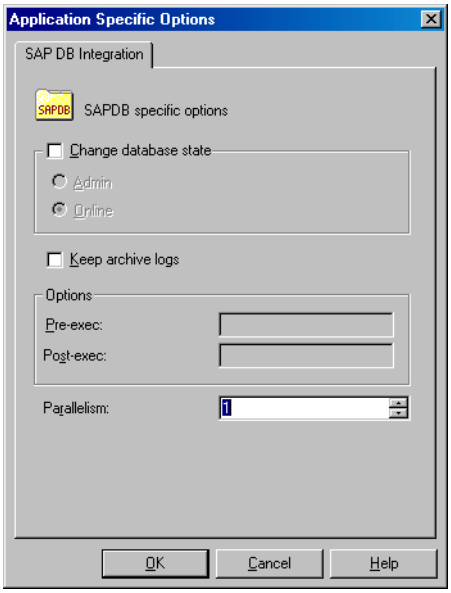

The following are the SAP DB/MaxDB specific backup options:

Change database state Selects the SAP DB/MaxDB database mode during the backup operation. The database can be either switched to the `Admin` or to the `Online` mode. If this option is not set, the database remains in the current mode.

Keep archive logs If this option is selected, the SAP DB/MaxDB archive logs are kept on the SAP DB/MaxDB Server after the backup has finished.

If this option is not selected, the SAP DB/MaxDB archive logs on the SAP DB/MaxDB Server are deleted after the backup has finished.

Parallelism Sets the `Parallelism` option to set the number of SAP DB/MaxDB media created on the SAP DB/MaxDB Server and consequently the number of SAP DB/MaxDB backup data streams. The default value is 1, the maximum value is 32 and the recommended value is the same as the number of SAP DB/MaxDB data volumes to be backed up.

The value of the `Parallelism` option must be equal as or lower than the SAP DB/MaxDB `MAXBACKUPDEVS` parameter.

The value of the `Parallelism` option must also be equal as or lower than the sum of concurrency values for all backup devices selected in the backup specification. For more information on the Data Protector `Concurrency` option, see online Help index: “concurrency“.

Modifying the Configuration of an SAP DB/MaxDB Instance in Data Protector

The parameters that need to be specified during the configuration of an SAP DB/MaxDB instance in Data Protector are the username and the password of the SAP DB/MaxDB user created or identified as described in the section “Configuring Users” on page 263, and the SAP DB/MaxDB independent program path parameter (the latter is, by default, detected automatically). These parameters are also used for establishing the connection to the SAP DB/MaxDB Server system if you start non-backup and non-restore-related operations in Data Protector, such as listing of objects for backup.

The configuration is performed during the creation of a new backup specification, or by modifying an existing backup specification. For the step-by-step procedure on creating an SAP DB/MaxDB backup specification, see “Creating a Backup Specification” on page 264.

If properly configured, the SAP DB/MaxDB user entered is allowed to back up or restore SAP DB/MaxDB Server database objects. In order to start a backup of an SAP DB/MaxDB object using Data Protector, the user must also be the owner of the Data Protector backup specification.

Refer to the SAP DB/MaxDB documentation for further information on different types of connections, roles and authorities of SAP DB/MaxDB database administrators and security issues that must be considered.

The configuration of an SAP DB/MaxDB instance in Data Protector can be modified using the Data Protector GUI or CLI.

Modifying the Configuration of an SAP DB/MaxDB Instance Using the GUI

The procedure below describes the re-configuration of an SAP DB/MaxDB instance in Data Protector in an existing backup specification:

1. In the `HP OpenView Storage Data Protector Manager`, switch to the `Backup` context.
2. In the Scoping Pane, expand `Backup`, `Backup Specifications`, and then `SAP DB Server`. Click an existing backup specification.
3. Right-click the name of the SAP DB/MaxDB instance listed in the `Source` property page, and then select `Configure` from the pop-up menu.
4. 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/MaxDB application on the SAP DB/MaxDB Server. You can leave the `Auto-detect` option selected to automatically detect the directory on the SAP DB/MaxDB Server.

   Enter the username and the password of the user created or identified as described in “Configuring Users” on page 263.

Figure 3-7 SAP DB/MaxDB Configuration

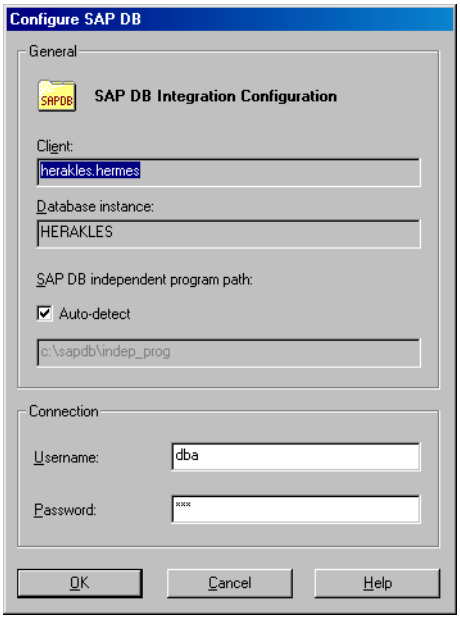

Click OK to confirm the configuration.

Modifying the Configuration of an SAP DB/MaxDB Instance Using the CLI

UNIX

Login to the SAP DB/MaxDB Server as the user under whose account the SAP DB/MaxDB application is running on the SAP DB/MaxDB Server (for example, the `sapdb` user in the `sapsys` group).

In a cluster environment, the environment variable `OB2BARHOSTNAME` must be defined as the virtual hostname before running the `util_sapdb.exe` command from the command line (on the client). The `OB2BARHOSTNAME` variable is set as follows:

```
export OB2BARHOSTNAME=<virtual_hostname>
```

Then execute the following command:

On HP-UX:

```
/opt/omni/lbin/util_sapdb.exe \
[-homedir <SAPDB_independent_program_directory>] \
-config <Instance_Name> <username> <password>
```

On other UNIX:

```
/usr/omni/bin/util_sapdb.exe \
[-homedir <SAPDB_independent_program_directory>] \
-config <Instance_Name> <username> <password>
```

Windows

In a cluster environment, the environment variable `OB2BARHOSTNAME` must be defined as the virtual hostname before running the `util_sapdb.exe` command (on the client). The `OB2BARHOSTNAME` variable is set as follows:

```
set OB2BARHOSTNAME=<virtual_hostname>
```

Execute the following command on the SAP DB/MaxDB Server system:

```
<Data_Protector_home>\bin\util_sapdb.exe \
[-homedir <SAPDB_independent_program_directory>] \
-config <Instance Name> <username> <password>
```

The parameters are defined as follows:

<SAPDB_independent_program_directory> The SAP DB/MaxDB independent program path parameter. This parameter is the independent program path directory specified during the installation of the SAP DB/MaxDB application on the SAP DB/MaxDB Server.

This parameter is optional. If it is not specified, the directory is detected automatically.

<Instance_Name> The name of the SAP DB/MaxDB instance to be configured.

<username> The username of the SAP DB/MaxDB user created or identified as described in “Configuring Users” on page 263.

<password>	The password of the SAP DB/MaxDB user created or identified as described in “Configuring Users” on page 263

NOTE The username and the SAP DB/MaxDB independent program path parameter must not contain the single quote character (‘).

Examples In the example below, the SAP DB/MaxDB independent program path is `/opt/sapdb/indep_prog` (UNIX systems) or `c:\program files\sapdb\indep_prog` (Windows systems), the instance name is `sapdb_inst`, the username is `sapdb_user` and the password is `sapdb_pass`.

UNIX On HP-UX:

```
/opt/omni/lbin/util_sapdb.exe -homedir
<SAPDB_independent_program_directory>/indep_prog -config
sapdb_inst sapdb_user sapdb_pass
```

On other UNIX:

```
/usr/omni/bin/util_sapdb.exe -homedir
<SAPDB_independent_program_directory>/indep_prog -config
sapdb_inst sapdb_user sapdb_pass
```

Windows

```
<Data_Protector_home>\bin\util_sapdb.exe -homedir
"<SAPDB_independent_program_directory>" -config sapdb_inst
sapdb_user sapdb_pass
```

Checking the Configuration of an SAP DB/MaxDB Instance

To check the configuration of an SAP DB/MaxDB instance, use either the GUI or CLI.

Checking the Configuration of an SAP DB/MaxDB Instance Using the GUI

Follow the procedure below to check the configuration of an SAP DB/MaxDB instance using the GUI:

1. In the HP OpenView Storage Data Protector Manager, switch to the Backup context.
2. In the Scoping Pane, expand Backup, Backup Specifications, and then SAP DB Server. Click an existing backup specification.
3. Right-click the name of the SAP DB instance listed in the Source property page, and then select Check configuration from the pop-up menu.

   A dialog box displaying the status of the check is shown. Click OK to close it.

Checking the Configuration of an SAP DB/MaxDB Instance Using the CLI

UNIX

Login to the SAP DB/MaxDB Server system as the user under whose account the SAP DB/MaxDB application is running (for example, the sapdb user in the sapsys group).

In a cluster environment, the environment variable OB2BARHOSTNAME must be defined as the virtual hostname before running the util_sapdb.exe command from the command line (on the client). The OB2BARHOSTNAME variable is set as follows:

```
export OB2BARHOSTNAME=<virtual_hostname>
```

Then execute the following command:

```
/opt/omni/lbin/util_sapdb.exe -chkconf <Instance_Name>
```

(HP-UX systems) or

```
/usr/omni/bin/util_sapdb.exe -chkconf <Instance_Name>
```

(other UNIX systems).

Data Protector attempts to connect to the SAP DB/MaxDB Server system using the information that was specified and saved during the configuration procedure.

In case of an error, the error number is displayed in the form *RETVAL*<*Error_number*>.

To obtain an error description, start the following command on the SAP DB/MaxDB Server system:

On HP-UX:

```
/opt/omni/lbin/omnigetmsg 12 <Error_number>
```

On other UNIX:

```
/usr/omni/bin/omnigetmsg 12 <Error_number>
```

Windows

In a cluster environment, the environment variable `OB2BARHOSTNAME` must be defined as the virtual hostname before running the `util_sapdb.exe` command from the command line (on the client). The `OB2BARHOSTNAME` variable is set as follows:

```
set OB2BARHOSTNAME=<virtual_hostname>
```

To check the configuration, run the following command on the SAP DB/MaxDB Server system:

```
<Data_Protector_home>\bin\util_sapdb.exe -chkconf
<Instance_Name>
```

If an error occurs, it is explained.

Testing the Integration

Once you have created and saved a backup specification, you should test it before running a backup.

Testing Using the Data Protector GUI

Testing Procedure

The testing procedure consists of checking the Data Protector part of the integration to ensure the communication within Data Protector is established and the data transfer works properly. Proceed as follows to test the integration:

1. In the `HP OpenView Storage Data Protector Manager`, switch to the `Backup` context.
2. In the Scoping Pane, expand `Backup`, then `Backup Specifications`, `SAP DB Server` and right-click the backup specification you want to preview.

3. Click `Preview Backup` to open the `Start Preview` dialog box. Select the type of backup you want to run as well as the network load. For a description of these options, press **F1**.

Figure 3-8 Previewing a Backup

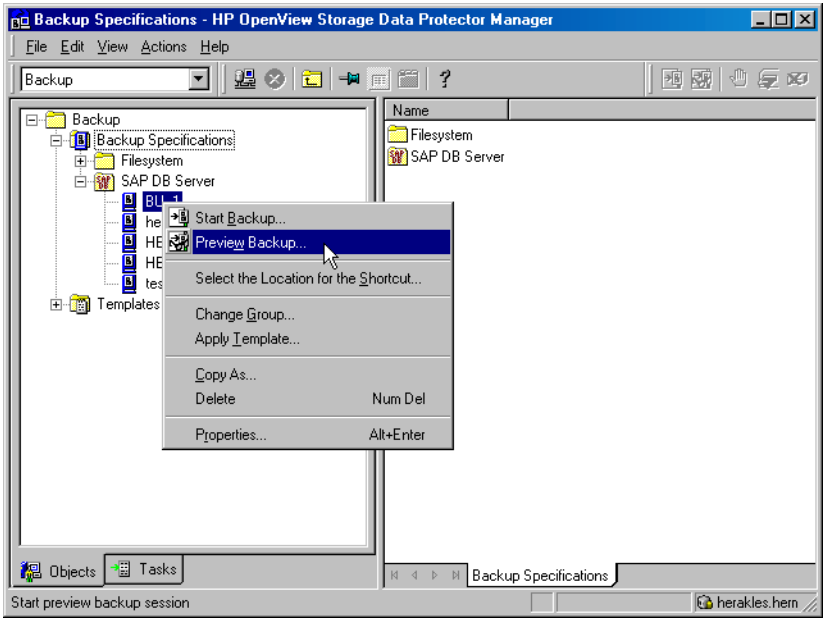

Testing Using the Data Protector CLI

To test a backup specification, run the `omnib` command with the `-test_bar` option.

Execute the following command:

- on HP-UX: `/opt/omni/bin/omnib -sapdb_list` *`<backup_specification_name>`* `-test_bar`
- on other UNIX: `/usr/omni/bin/omnib -sapdb_list` *`<backup_specification_name>`* `-test_bar`
- on Windows: *`<Data_Protector_home>`*`\bin\omnib -sapdb_list` *`<backup_specification_name>`* `-test_bar`

Refer to the `omnib` man page for more information on the `omnib` command.

What Happens?

The session messages are displayed on the screen during the command execution.

The `sapdbbar.exe` program is started, which then starts the Data Protector `testbar2` command. This command checks the following:

- if the communication within Data Protector works properly
- if the syntax of the SAP DB/MaxDB Integration backup specification is correct
- if the devices are correctly configured
- if the required media reside in the devices

After that, the SAP DB/MaxDB part of the preview is started, which checks if all the backup objects are present and are in a correct mode for a backup.

Backing Up an SAP DB/MaxDB Database

During the backup, the database can be in either the Admin or in the Online mode. To perform an offline backup of an SAP DB/MaxDB instance, a regular Data Protector filesystem backup should be configured. See online Help index: “standatd backup procedure“ for more information on filesystem backups.

During a backup in the Online mode, the database is open and available for the other applications. During a backup in the Admin mode, operations on the database are suspended.

Backup Methods

To run a backup, use any of the following methods:

- Schedule a backup of an existing SAP DB/MaxDB backup specification using the Data Protector Scheduler.
- Start an interactive backup using the Data Protector GUI or the Data Protector CLI.

Scheduling an Existing Backup Specification

Data Protector allows you to run unattended backups at specific times or periodically. The powerful Data Protector Scheduler can highly influence the effectiveness and performance of your backup.

For more detailed information on scheduling, refer to the online Help index keyword “scheduled backups”.

To schedule an existing backup specification, perform the following steps in the HP OpenView Storage Data Protector Manager:

Scheduling Procedure

1. In the Context List, select Backup.
2. In the Scoping Pane, expand Backup, then Backup Specifications. Click SAP DB Integration.

   A list of configured backup specifications is displayed in the Results Area.
3. Double-click the backup specification you want to schedule and click the Schedule tab to open the Schedule property page.

4. In the Schedule property page, select a date in the calendar click Add to open the Schedule Backup dialog box.
5. Specify Recurring, Time options, Recurring options, and Session options. See Figure 3-9 on page 283.

IMPORTANT

To backup SAP DB/MaxDB archive logs, the Data item must be selected in the Source property page of the backup specification. Additionally, the Trans backup type under Session options must be selected.

6. Click OK to return to the Schedule property page.
7. Click Apply to save the changes.

Scheduling Example

To schedule a backup specification so as to produce a full backup at 8.00 a.m., and then a differential backup at 1.00 p.m. and at 6.00 p.m. during week days, open the Schedule property page of the backup specification as described in the above procedure, and then proceed as follows:

1. In the Schedule property page, click Add to open the Schedule Backup dialog box.
2. Under Recurring, select Weekly. Under Time options, select the time 8 AM. Under Recurring Options, select Mon, Tue, Wed, Thu, and Fri. Under Session options, select the Full backup type. Click OK.

   See Figure 3-9 on page 283.

Figure 3-9 Scheduling the Backup Specification

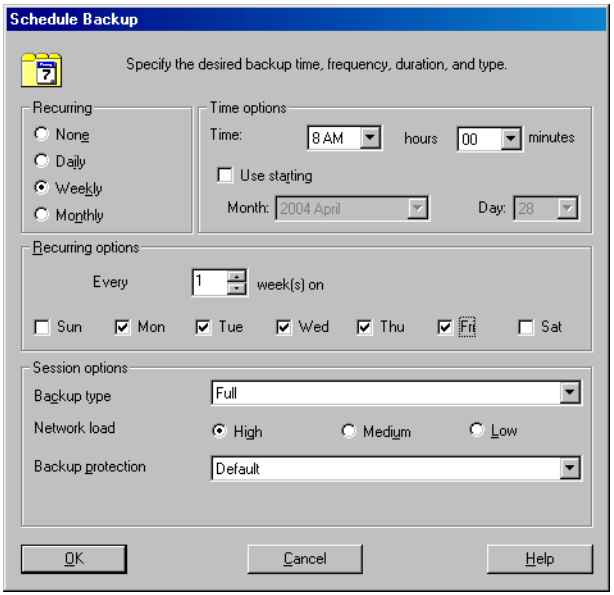

3. Repeat steps 1 and 2 to schedule another backup. Specify options as described, except the time, which should be set to `1 PM`, and the Backup type that should be set to `Diff`.
4. Repeat steps 1 and 2 to schedule another backup. Specify options as described, except the time, which should be set to `6 PM`, and the Backup type that should be set to `Diff`.
5. Click `Apply` to save the changes.

After scheduling your backup, you can have it run unattended or you can still run it interactively, as shown in the next section.

Refer to the online Help index: “scheduled backups” for scheduling details.

NOTE When creating an SAP DB/MaxDB backup specification, you access the Data Protector Scheduler through the Backup Wizard. See “Creating a Backup Specification” on page 264 for information about accessing the Backup Wizard.

Running an Interactive Backup Using the Data Protector GUI

An interactive backup can be run any time after the backup specification has been created and saved.

Backup Procedure

To start an interactive backup of an SAP DB/MaxDB backup object using the Data Protector GUI, perform the following steps:

1. In the `HP OpenView Storage Data Protector Manager`, switch to the `Backup` context.
2. In the Scoping Pane, expand the `Backup`, and then the `Backup Specifications` items.

   Expand `SAP DB Integration`. A list of backup specifications appears.
3. Right-click the backup specification you want to back up, and then select `Start Backup` from the pop-up menu.

   The `Start Backup` dialog box appears.

   Select the backup type and network load. For a description of these options, press **F1**.

IMPORTANT

To backup SAP DB/MaxDB archive logs, the `Data` item must be selected in the `Source` property page of the backup specification. Additionally, the `Trans` backup type must be selected.

Click `OK`.

Figure 3-10 Starting an Interactive Backup

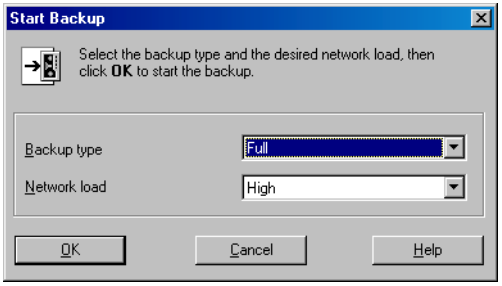

Messages appear in the Results Area as the backup session proceeds. Upon successful completion of the backup session, the `Session completed successfully` message and backup size are displayed.

Running an Interactive Backup Using the Data Protector CLI

You can start an interactive backup from the Data Protector CLI. Switch to the `/opt/omni/bin` (HP-UX systems), `/usr/omni/bin` (other UNIX systems) or *<Data_Protector_home>*`\bin` (Windows systems) directory, and run the following command:

```
omnib -sapdb_list <ListName> [-barmode <sapdbmode>]
[<list_options>] [-preview]
```

The *<ListName>* parameter is the name of the backup specification.

The *<sapdbmode>* parameter specifies the type of the backup.

The *<list_options>* parameters set the level of the protection, the level of the network traffic generated by the session, enables writing a CRC checksum, and disables monitoring of the backup session.

You can select among the following *<sapdbmode>* backup modes:

```
{full | diff | trans}
```

You can select among the following *<list_options>*:

```
-protect {none | weeks n | days n | until date | permanent}
-load {low | medium | high}
-crc
-no_monitor
```

Refer to the `omnib` man page for more information.

Example

To start a full backup using an existing SAP DB/MaxDB backup specification called `TEST`, and to set data protection to 10 weeks, execute the following command:

```
omnib -sapdb_list TEST -barmode full -protect weeks 10
```

Running an Interactive Backup Using SAP DB/MaxDB Utilities

Using this integration, you can run an integrated Data Protector backup of an SAP DB/MaxDB Server from SAP DB/MaxDB utilities.

Prerequisites

- The SAP DB/MaxDB Server must be configured for use with this integration. See “Configuring the Integration” on page 263 for information on how to configure an SAP DB/MaxDB Server for use with this integration.
- A Data Protector SAP DB/MaxDB integration backup specification must be created. See “Configuring an SAP DB/MaxDB Backup” on page 264 for information on how to create a Data Protector SAP DB/MaxDB integration backup specification.

In order to run an interactive backup using SAP DB/MaxDB utilities, the following must be done:

- SAP DB/MaxDB media must be created
- the SAP DB/MaxDB `bsi_env` file must be created

Follow the procedure on the next few pages to run a backup from SAP DB/MaxDB utilities using an existing Data Protector SAP DB/MaxDB integration backup specification. In the procedure, the following conventions are used:

<inst_name> is the name of the instance to be backed up

<name_of_backup_spec> is the name of the Data Protector backup specification to be used for backup

<username>, *<password>* is the connection string for the SAP DB/MaxDB user created or identified as described in “Configuring Users” on page 263

<location> is the location of the `bsi_env` file

<media_group_name> is the name of the SAP DB/MaxDB media group

<medium_name> is the name of the SAP DB/MaxDB medium

<pipe_name> is the name of the SAP DB/MaxDB pipe

<medium_type> is the type of the SAP DB/MaxDB medium

1. On the SAP DB/MaxDB Server create the `bsi_env` file in a directory of your choice. The file must have the read permission set for the OS user under whose account the database runs (the user is described in "Configuring Users" on page 263) and 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
PARAMETERFILE <name_of_backup_spec>
TIMEOUT_SUCCESS 60
TIMEOUT_FAILURE 30
```

UNIX

On 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
```

On 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. Login to the SAP DB/MaxDB database manager as the SAP DB/MaxDB user created or identified as described in "Configuring Users" on page 263. On the SAP DB/MaxDB Server, execute the following command to login:

   ```
   dbmcli -d <inst_name> -u <username>,<password>
   ```

3. In the SAP DB/MaxDB database manager, register the location of the `bsi_env` file created in the step 1. of this procedure as follows:

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/MaxDB media in an SAP DB/MaxDB media group. Execute the following command for every medium to be created:

   ```
   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/MaxDB media 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
```

5. Start the SAP DB/MaxDB utility session by executing the following command:

   ```
   util_connect
   ```

6. Start the backup. The following exemplary command starts the full backup for the media created in the previous step of this procedure:

   ```
   backup_start BACKDP-Data[2] DATA
   ```

7. Observe the progress of the session in the Data Protector `Monitor` context. For more information on how to do this, see “Monitoring an SAP DB/MaxDB Backup and Restore” on page 310.

Restoring an SAP DB/MaxDB Database

An SAP DB/MaxDB database can be restored using the Data Protector GUI or CLI. An integrated restore is performed.

An SAP DB/MaxDB database can be either restored or migrated. Both can be accomplished using any of the following methods:

- The Data Protector GUI: see “Restoring Using the Data Protector GUI” on page 295.
- The Data Protector CLI: see “Restoring Using the Data Protector CLI” on page 297.
- The SAP DB/MaxDB utilities: see “Restoring Using SAP DB/MaxDB Utilities” on page 299.

When performing an SAP DB/MaxDB migration using any of the above methods, some additional tasks must first be done in order to prepare the SAP DB/MaxDB Server or instance. These tasks are described in “SAP DB/MaxDB Migration Prerequisites” on page 294.

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/MaxDB Restore and Recovery Options” on page 304.

At the beginning of a restore session, Data Protector switches the SAP DB/MaxDB 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/MaxDB 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/MaxDB Server and the restored volumes can occur. When performing recovery (when the database is switched to the `Online` mode), SAP DB/MaxDB 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/MaxDB 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/MaxDB Server before the restore are not deleted during the restore. Note that during the SAP DB/MaxDB migration, the redo logs that existed on the SAP DB/MaxDB Server before the restore are deleted during the restore.

When restoring, the existing redo logs on the SAP DB/MaxDB 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/MaxDB 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/MaxDB 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/MaxDB Server (if full or diff backup session is restored).

NOTE The `Use existing archive logs` option is disabled in case of SAP DB/MaxDB 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 3-11 SAP DB/MaxDB Restore Process

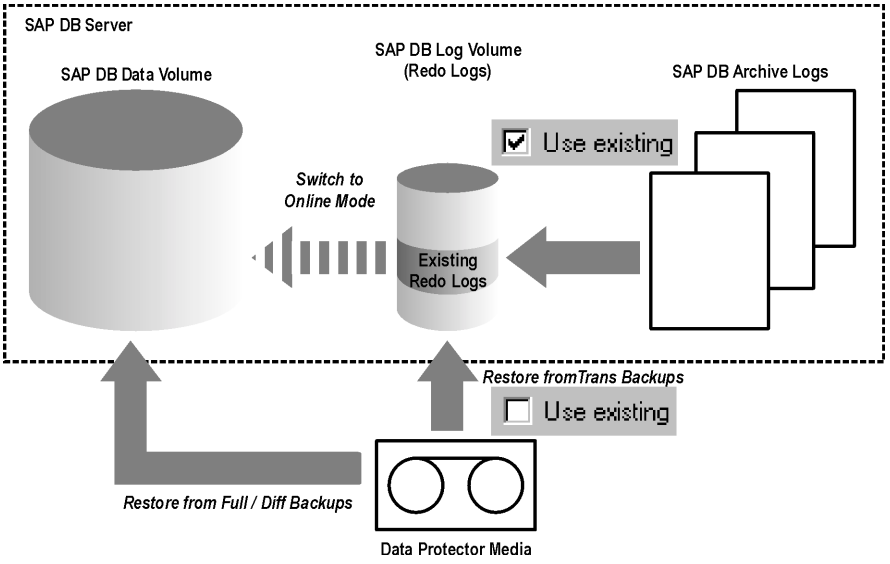

Figure 3-12 **SAP DB/MaxDB Archive Logs Restore Process—Redo Logs Details**

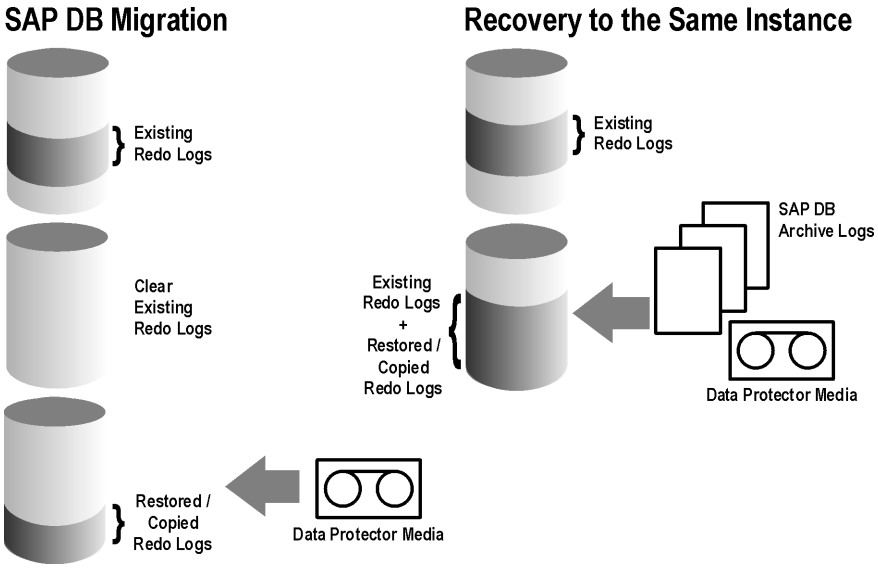

If you select a differential or a transactional backup session to be restored, you can set the integration to:

- 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/MaxDB Server are never deleted.

SAP DB/MaxDB Migration Prerequisites

The integration supports SAP DB/MaxDB migration, meaning that an SAP DB/MaxDB instance can be restored to some other SAP DB/MaxDB Server or instance than the original.

If the SAP DB/MaxDB Server has not yet been configured for the Data Protector SAP DB/MaxDB integration, it must be configured before the restore is started. If the instance does not exist, it must be configured before the restore is started. During the migration, the existing data is overwritten and the existing redo logs are deleted.

Perform the following list of tasks before starting an SAP DB/MaxDB migration:

- Install the Data Protector SAP DB/MaxDB integration on the SAP DB/MaxDB Server to which you want to migrate the backed up SAP DB/MaxDB instance. When this is done, add the SAP DB/MaxDB Server in the Data Protector cell.
- Identify or create an OS user under whose account the SAP DB/MaxDB is running and add it to the Data Protector `admin` group. For information on how to do this, see “Configuring Users” on page 263.
- When performing an SAP DB/MaxDB migration, first configure the instance to which you want to perform the restore. For information on how to configure an instance, see “Modifying the Configuration of an SAP DB/MaxDB Instance in Data Protector” on page 272.

  When performing an SAP DB/MaxDB migration using the Data Protector GUI, the instance does not need to be configured beforehand, it can be configured during the restore process.

Restoring Using the Data Protector GUI

When performing an SAP DB/MaxDB migration, some additional tasks must first be done in order to prepare the SAP DB/MaxDB Server or instance. These tasks are described in “SAP DB/MaxDB Migration Prerequisites” on page 294.

To restore your data, proceed as follows in the `HP OpenView Storage Data Protector Manager`:

1. In the `Context List`, select `Restore`.

   Expand `Restore Objects`, then `SAP DB Server`, and then the SAP DB/MaxDB Database Server which you want to restore.

   Select the SAP DB/MaxDB instance. A list of backed up objects is displayed in the Results Area.

2. Select the SAP DB/MaxDB objects you want to restore.

Figure 3-13 Selecting SAP DB/MaxDB Objects for a Restore

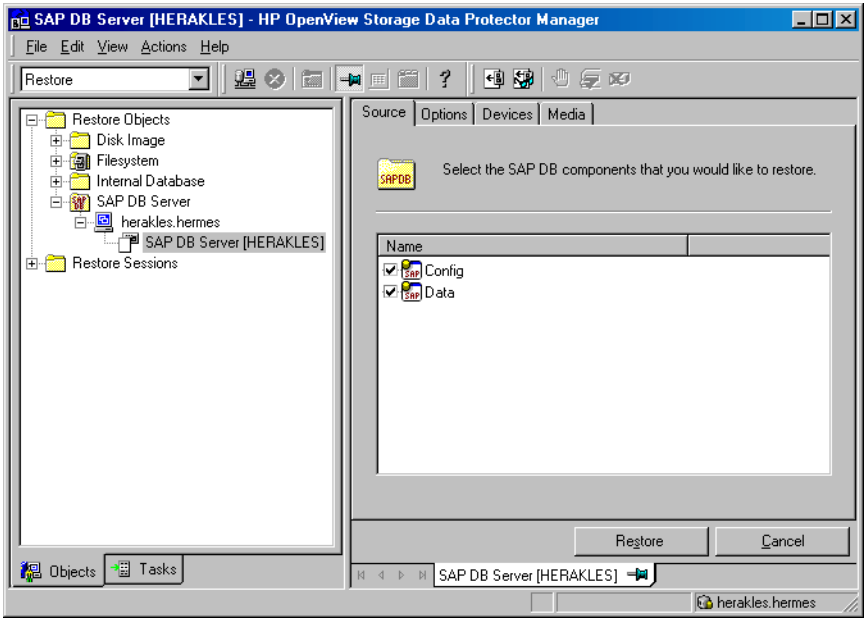

IMPORTANT To restore SAP DB/MaxDB archive logs, select the `Data` item in the Results Area. The archive log restore is then triggered by selecting any of the `Trans` backup sessions in the next step of this procedure.

3. Right-click the `Data` item and select `Properties` from the pop-up menu to open the Properties for Data dialog box.

Figure 3-14 The Properties for Data Dialog Box

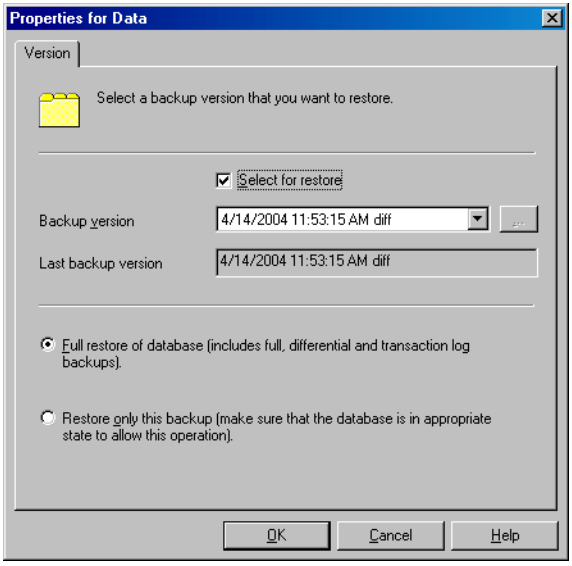

In the Properties for Data dialog box, select the backup session you want to restore in the `Backup version` drop-down list.

IMPORTANT Do not select the backup session for the `Configuration` item. The same session as selected for the `Data` item will be used, regardless of what you select for the `Configuration` item.

If you select a Trans or a Diff backup session, you are given a possibility to:

- Perform a full restore of the database (the Full restore of database option). In this case, the integration automatically determines the chain of needed full, differential or transactional backup sessions when performing a restore. After the restore has finished, the database is, if the Recovery option is selected, switched to the Online mode.
- Restore only the selected backup session (the Restore only this backup option). If a database becomes 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.

Click OK to close the Properties for Data dialog box.

4. Click the Options tab to set the restore and recovery options. See “SAP DB/MaxDB Restore and Recovery Options” on page 304 for more information on SAP DB/MaxDB restore and recovery options.
5. Click Restore to start the restore session.

   When the session starts, messages are displayed in the Results Area. Upon successful completion, a message is issued in the Session Information dialog box.

Restoring Using the Data Protector CLI

When performing an SAP DB/MaxDB migration, some additional tasks must first be done in order to prepare the SAP DB/MaxDB Server or instance. These tasks are described in the“SAP DB/MaxDB Migration Prerequisites” on page 294.

Finding Information Needed for a 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/MaxDB objects.

- `omnidb -sapdb <object_name>`

  to get details on a specific object, including the SessionID.

Restoring

The `omnir` Command Syntax

The following is the syntax of the `omnir` command to be used to restore an SAP DB/MaxDB instance:

```
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 `-sapdb` option selects an SAP DB/MaxDB restore.

The `-barhost` option sets the name of the SAP DB/MaxDB Server that was backed up.

The `-instance` option sets the name of the SAP DB/MaxDB 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/MaxDB Restore and Recovery Options” on page 304. Refer also to the `omnir` man page.

Example

To restore an instance named “inst1” (together with configuration), backed up on an SAP DB/MaxDB Server named “srv1.company.com” from the last backup session and then perform a recovery until the end of logs, enter the following command:

UNIX

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
```

Windows

```
<Data_Protector_home>\omnir -sapdb -barhost
srv1.company.com -instance inst1 -recover -endlogs
```

The restore session can be monitored in the `Data Protector Monitor` window, where mount prompts for the required media are also displayed.

Restoring Using SAP DB/MaxDB Utilities

Using this integration, it is also possible to run an integrated Data Protector restore of an SAP DB/MaxDB Server from SAP DB/MaxDB utilities.

To perform a restore to an existing SAP DB/MaxDB Server instance, see “SAP DB/MaxDB Restore and Recovery” on page 300.

To migrate an SAP DB/MaxDB instance, see “SAP DB/MaxDB Migration” on page 303.

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/MaxDB objects.

- `omnidb -sapdb <object_name>`

  to get details on a specific object, including the SessionID.

SAP DB/MaxDB Restore and Recovery

Follow the procedure on the next few pages to restore and recover a database using SAP DB/MaxDB utilities from existing Data Protector SAP DB/MaxDB 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/MaxDB user created or identified as described in "Configuring Users" on page 263

<location> is the location of the `bsi_env` file

<media_group_name> is the name of the SAP DB/MaxDB media group

<medium_name> is the name of the SAP DB/MaxDB medium

<pipe_name> is the name of the SAP DB/MaxDB pipe

<medium_type> is the type of the SAP DB/MaxDB medium

<SessionID> is the Data Protector session ID of the session to be restored

Restore

1. Skip this step if the `bsi_env` file is already present and configured on the SAP DB/MaxDB Server.

   On the SAP DB/MaxDB 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
```

UNIX

On 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
```

On 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/MaxDB database manager as the SAP DB/MaxDB user created or identified as described in "Configuring Users" on page 263. On the SAP DB/MaxDB Server, execute the following command to login:

   ```
   dbmcli -d <inst_name> -u <username>,<password>
   ```

3. In the SAP DB/MaxDB 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/MaxDB 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/MaxDB media and pipes to be used with Data Protector are already existing on the SAP DB/MaxDB Server.

   Note that to restore a Data Protector SAP DB/MaxDB backup session, the number of SAP DB/MaxDB media and pipes required equals the parallelism value used during the backup session.

Create SAP DB/MaxDB media in an SAP DB/MaxDB 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/MaxDB 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/MaxDB 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.

Recovery

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/MaxDB database manager:

      ```
      db_online
      ```

   b. To recover the database until the specified point in time, execute the following command in the SAP DB/MaxDB 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/MaxDB Migration

When performing an SAP DB/MaxDB migration, some additional tasks must first be done in order to prepare the SAP DB/MaxDB Server or instance. These tasks are described in the“SAP DB/MaxDB Migration Prerequisites” on page 294.

Follow the procedure in the section “SAP DB/MaxDB Restore and Recovery” on page 300 to migrate the SAP DB/MaxDB database using SAP DB/MaxDB utilities from existing Data Protector SAP DB/MaxDB backup session(s). When following the mentioned procedure, *before* executing the `recover_start` command, delete the existing redo logs on the SAP DB/MaxDB Server by executing the following command in the SAP DB/MaxDB database manager:

```
util_execute clear log
```

SAP DB/MaxDB Restore and Recovery Options

SAP DB/MaxDB restore and recovery options can be accessed in the Data Protector GUI by clicking the `Options` tab in the Restore context after an SAP DB/MaxDB object has been selected.

Figure 3-15 SAP DB/MaxDB Restore and Recovery Options

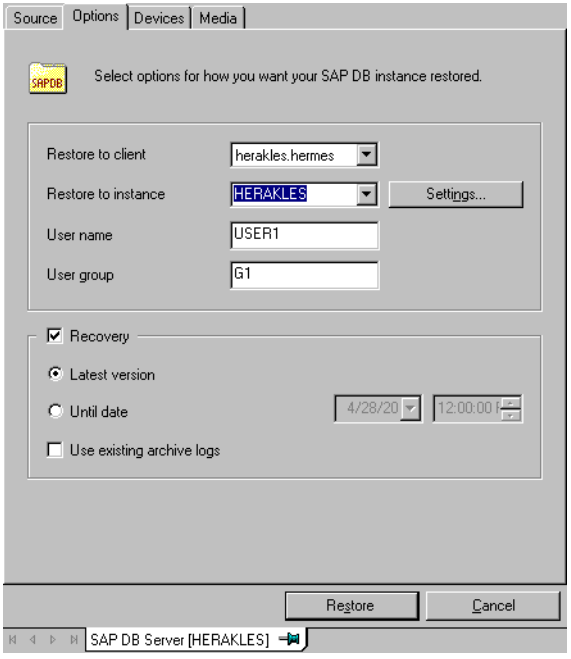

The following are SAP DB/MaxDB specific backup options:

Migration Options

To restore selected SAP DB/MaxDB object to the same SAP DB/MaxDB Server and instance, leave the migration options as they are. Use the migration options only in case of SAP DB/MaxDB migration (when restoring to some other SAP DB/MaxDB 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/MaxDB Server to which you want to restore the database.

When using the CLI, specify the `-destination` option and the name of the SAP DB/MaxDB Server as the *<ClientName>* argument.

The selected SAP DB/MaxDB 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 the Integration” on page 263 for information on how to configure an SAP DB/MaxDB 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 the Integration” on page 263 for information on how to configure an SAP DB/MaxDB 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/MaxDB application is running on the SAP DB/MaxDB 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 step 4. on page 273 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 "Modifying the Configuration of an SAP DB/MaxDB Instance Using the CLI" on page 274.

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/MaxDB Server and the restored volumes can occur. When performing recovery (when the database is switched to the `Online` mode), SAP DB/MaxDB 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, refer to step 3 on page 296. For information on the `-session` option, refer to "Restoring Using the Data Protector CLI" on page 297.

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/MaxDB Server to SAP DB/MaxDB 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/MaxDB 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/MaxDB Server are applied to redo logs.

NOTE The `Use existing archive logs` option is disabled in case of SAP DB/MaxDB 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).

Using Another Device

Data Protector supports restore using a device other than the one that was used at backup time.

If you are performing a restore using the Data Protector GUI, see online Help index: “selecting, devices for restore“ for more information on how to perform a restore using another device.

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 refer to the *HP OpenView Storage 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 troubleshooting section.
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 an SAP DB/MaxDB Backup and Restore

The Data Protector GUI enables you to monitor current or view previous backup and restore sessions.

Monitoring is automatically activated when you start a restore or backup interactively.

Monitoring Current Sessions

To monitor a currently running session using the Data Protector GUI, proceed as follows:

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. See Figure 3-16.

Figure 3-16 Monitoring a Current Session

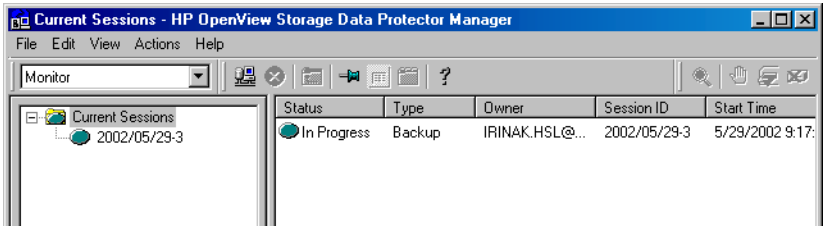

Clearing Sessions

To remove all completed or aborted sessions from the Results Area of the `Monitor` context, proceed as follows:

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.

For detailed information on a completed or aborted session, see "Viewing Previous Sessions".

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. See Figure 3-17.

Figure 3-17 Viewing a Previous Session

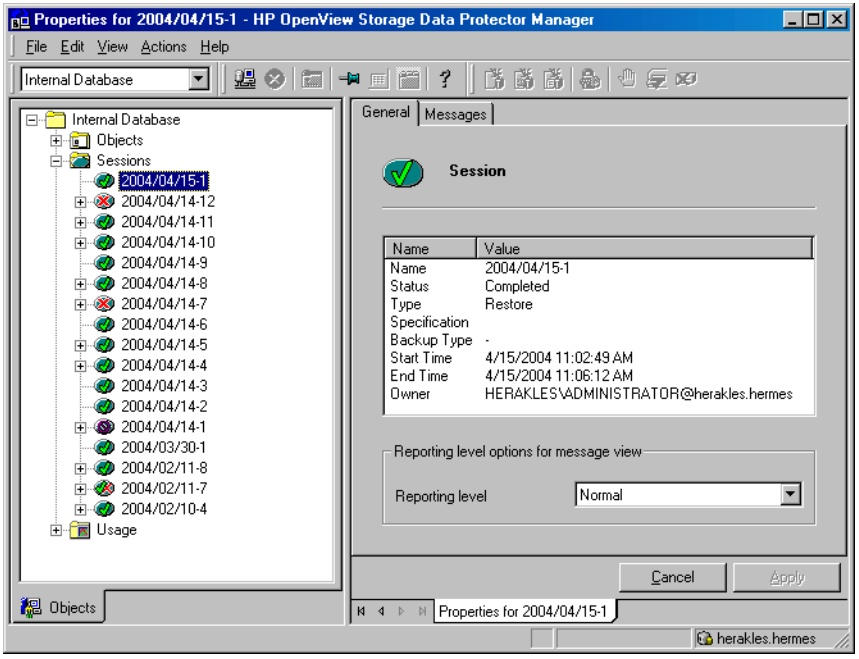

Troubleshooting

This section lists problems you might encounter when using the Data Protector SAP DB/MaxDB integration.

For general Data Protector troubleshooting information, see the *HP OpenView Storage 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 OpenView Storage 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

18:17:18 Error: SAPDB responded with:

-24920,ERR_BACKUPOP: backup operation was unsuccessful
The database was unable to fulfill a request

(-2025, Invalid number of backup devices).
```

Action

Increase the value of the SAP DB/MaxDB `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/MaxDB instance cannot be started after restore.**

Action Using the SAP DB/MaxDB `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/MaxDB instance kernel error file for errors.

  If there was insufficient space for SAP DB/MaxDB 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/MaxDB 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/MaxDB x_server. Refer to the SAP DB/MaxDB 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/MaxDB instance. Refer to the SAP DB/MaxDB 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,utility session is already in use
```

Action Some other user is connected to the SAP DB/MaxDB instance and is performing administrative tasks (utility session). Such SAP DB/MaxDB 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/MaxDB instance as described in the section “Modifying the Configuration of an SAP DB/MaxDB Instance in Data Protector” on page 272.

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 with code 0.
```

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 on the `util_cmd` command, refer to "Data Protector SAP DB/MaxDB Configuration File" on page 259.

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/MaxDB instance object in the Results Pane under the `Source` tab and select the `Set Environment Variables` from the pop-up menu.

SAP DB/MaxDB 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**.

ACSLS *(StorageTek specific term*)
The Automated Cartridge System Library Server (ACSLS) software that manages the Automated Cartridge System (ACS).

Active Directory *(Windows specific term*)
The directory service in a Windows network. It contains information about resources on the network and makes them accessible to users and applications. The directory services provide a consistent way to name, describe, locate, access and manage resources regardless of the physical system they reside on.

AML *(EMASS/GRAU specific term)*
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 *(ZDB specific term)*
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 *(Oracle specific term)*
Also called offline redo log. If the Oracle database operates in the ARCHIVELOG mode, as each online redo log is filled, it is copied to an archived log destination. This copy is the archived redo log. The presence or absence of an archived redo log is determined by the mode the database is using:

- 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 *(Lotus Domino Server specific term*)
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 *<Data_Protector_home>*\Config\Server\dr\asr on a Windows Cell Manager or in /etc/opt/omni/server/dr/asr/ 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.

autochanger
See **library**

autoloader
See **library**

Automatic Storage Management
(Oracle specific term)
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.

BACKINT *(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.

backup API
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.

backup chain
See **restore chain**.

backup device
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.

backup generation
One backup generation includes one full backup and all incremental backups until the next full backup.

backup ID
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.

backup object
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 entity or a disk image (rawdisk).

A backup object is defined by:

- Client name: hostname of the Data Protector client where the backup object resides.
- Mount point: the access point in a directory structure (drive on Windows and mount point on UNIX) on the client where the backup object is located.
- Description: uniquely defines backup objects with identical client name and mount point.
- Type: backup object type (for example filesystem or Oracle).

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 **incremental backup** and **full backup.**

backup set
A complete set of integration objects associated with a backup.

backup set *(Oracle specific term)*
A logical grouping of backed up files that are created using the RMAN backup command. A backup set is a complete set of files associated with a backup. The files can be multiplexed to increase performance. A backup set contains either datafiles or archived logs, but not both together.

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, 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 *(ZDB specific term)*
A system connected to target volumes of one or multiple application systems. The backup system is typically connected to a ZDB disk arraybackup 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 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 *(EMC Symmetrix specific term)*
Business Continuance are processes that allow customers to access and manage instant copies of EMC Symmetrix standard devices.
See also **BCV**.

BC *(HP StorageWorks Disk Array XP specific term)*
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 to the backup system.
See also **HP StorageWorks Disk Array XP LDEV**, **CA**, **Main Control Unit**, **application system**, and **backup system**.

BC EVA *(HP StorageWorks EVA specific term)*
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 *(EMC Symmetrix specific term)*
A protected storage environment solution that has defined specially configured EMC Symmetrix devices as mirrors or Business Continuance Volumes to protect data on EMC Symmetrix standard devices.
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 consistency checkmanual disaster recovery is equivalent to consistencymanual AND checkdisaster 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 **SAPDBA**, **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 **SAPDBA**, **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 **SAPDBA**, **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 *(UNIX specific term)*
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**.

channel *(Oracle specific term)*
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:

- type 'disk'
- type ‘sbt_tape’

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.

circular logging *(Microsoft Exchange Server and Lotus Domino Server specific term*)
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 writers and 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 ...).

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 *(Windows specific term)*
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
A set of DOS and UNIX like commands that you can use in shell scripts to perform Data Protector configuration, backup, and restore, and management tasks.

Command View (CV) EVA *(HP StorageWorks EVA specific term)*
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**.

concurrency
See **Disk Agent concurrency**.

control file *(Oracle and SAP R/3 specific term)*
An Oracle data file that contains entries specifying the physical structure of the database. It provides database consistency information used for recovery.

copy set *(HP StorageWorks EVA specific term)*
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) runs on the Data Protector Cell Manager. It starts and controls the backup and restore sessions. The service is started as soon as Data Protector is installed on the Cell Manager.
CRS runs under the account root on UNIX systems, and under any Windows account. By default, it runs under the account of the user, specified at installation time.

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 *(Oracle and SAP R/3* specific term)
A physical file created by Oracle that contains data structures such as tables and indexes. A data file can only belong to one Oracle database.

data protection
Defines how long the backed up data on media remains protected, that is, Data Protector will not overwrite it. When the protection expires, Data Protector will be able to reuse the media in one of the next backup sessions.
See also **catalog protection**.

Data Protector Event Log
A central repository of all Data Protector related notifications. By default, all notifications are sent to the Event Log. The events are logged in the *<Data_Protector_home>*\log\server\Ob2EventLog.txt file on the Cell ManagerThe Event Log is accessible only to Data Protector users in the Admin group and to Data Protector users who are granted the Reporting and notifications user rights. You can view or delete all events in the Event Log.

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.

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 *(HP StorageWorks EVA specific term)*
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.
See also **copy set**.

database server
A computer with a large database stored on it, such as the SAP R/3 or Microsoft SQL database. A server has a database that can be accessed by clients.

Dbobject *(Informix Server specific term)*
An Informix Server physical database object. It can be a blobspace, dbspace, or logical log file.

DC directory
The Detail Catalog (DC) directory consists of DC binary files, which store information about file versions. It represents the DCBF part of the IDB, which occupies approximately 80% of the IDB. The default DC directory is called the dcbf directory and is located in the <*Data_Protector_home*>\db40 directory on a Windows Cell Manager and in the /var/opt/omni/server/db40 directory on a UNIX Cell Manager. You can create more DC directories and locate them as appropriate to you. Up to 10 DC directories are supported per cell. The default maximum size of a DC directory is 4 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.

delta backup
A delta backup is a backup containing all the changes made to the database from the last backup of any type.
See also **backup types**

device
A physical unit which contains either just a drive or a more complex unit such as a library.

device chain
A device chain consists of several standalone devices configured for sequential use. When a medium in one device gets full, the backup automatically continues on a medium in the next device in the device chain.

device group *(EMC Symmetrix specific term)*
A logical unit representing several EMC Symmetrix devices. A device cannot belong to more than a single device group. All devices in a device group must be on the same EMC Symmetrix unit. You can use a device group to identify and work with a subset of the available EMC Symmetrix devices.

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 Incr1 backup type.
See **incremental backup**.

differential backup *(MS SQL specific term)*
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 *(Windows specific term)*
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 *(Veritas Volume Manager specific term)*
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 *(Informix Server specific term)*
The Informix Server configuration file ixbar.<*server_id*> that resides in the directory <*INFORMIXDIR*>/etc (on Windows) or <*INFORMIXDIR*>\etc (on UNIX). <*INFORMIXDIR*> is the Informix Server home directory and <*server_id*> is the value of the SERVERNUM 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.
See also **MoM**.

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.

exchanger
Also referred to as SCSI Exchanger.
See also **library**.

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.
See also **importing media.**

Extensible Storage Engine (ESE)

(Microsoft Exchange Server specific term)
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 *(HP StorageWorks EVA specific term)*
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 bidirectional 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 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 (*HP StorageWorks Disk Array XP specific term*)
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 numbers**.

flash recovery area *(Oracle specific term)*
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 fnames.dat 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 (media 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 backed upstreamed 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 /etc/opt/omni/server/options directory on HP-UX and Solaris systems and in the *<Data_Protector_home>*\Config\Server\Options directory on Windows systems.

group *(Microsoft Cluster Server specific term)*
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, and Windows) graphical user interface, provided by Data Protector for easy access to all configuration, administration, and operation tasks.

hard recovery *(Microsoft Exchange Server specific term)*
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: /etc/opt/omni/server/Holidays on the UNIX Cell Manager and *<Data_Protector_home>*\Config\Server\holidays 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 **OVO**.

HP OpC
See **OVO**.

HP OpenView SMART Plug-In (SPI)
A fully integrated, out-of-the-box solution which "plugs into" HP OpenView Operations, 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 OpenView Operations (OVO).

HP OVO
See **OVO**.

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** *(HP StorageWorks Disk Array XP specific term)*, **CA** *(HP StorageWorks Disk Array XP specific term)*, 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 **OVO**.

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 so onwhich 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.
See also **exporting media.**

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.
See also **backup types**.

incremental backup *(Microsoft Exchange Server specific term)*
A backup of the Microsoft Exchange Server data that has changed since the last full or incremental backup. With the incremental backup, only the transaction log files are backed up.
See also **backup types**.

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 *(EMC Symmetrix specific term)*
A BCV or SRDF control operation. In BCV control operations, an incremental establish causes the BCV device to be synchronized incrementally and to function as an EMC Symmetrix mirrored medium. The EMC Symmetrix devices must have been previously paired.
In SRDF control operations, an incremental establish causes the target (R2) device to be synchronized incrementally and to function as an EMC Symmetrix mirrored medium. The EMC Symmetrix devices must have been previously paired.

incremental restore *(EMC Symmetrix specific term)*
A BCV or SRDF control operation.
In BCV control operations, an incremental restore reassigns a BCV device as the next available mirror of the standard device in the pair. However, the standard devices are updated with only the data that was written to the BCV device during the time of the original pair split, and the data that was written to the standard device during the split is overwritten with data from the BCV mirror. In SRDF control operations, an incremental restore reassigns a target (R2) device as the next available mirror of the source (R1) device in the pair. However, the source (R1) devices are updated with only the data that was written to the target (R2) device during the time of the original pair split, and the data that was written to the source (R1) device during the split is overwritten with data from the target (R2) mirror.

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 sessions, 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/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**.

integrated security *(MS SQL specific term)*
Integrated security allows the Microsoft SQL Server to use Windows authentication mechanisms to validate Microsoft SQL Server logins for all connections. Using integrated security means that users have one password for both Windows and Microsoft SQL Server. Integrated security should be used in environments where all clients support trusted connections. Connections validated by Windows Server and accepted by Microsoft SQL Server are referred to as trusted connections. Only trusted connections are allowed.

integration object
A backup object of a Data Protector integration, such as Oracle or SAP DB.

Internet Information Server (IIS)

(Windows specific term)
Microsoft Internet Information Server 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
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 *(Sybase specific term)*
A Sybase utility used to perform system administration tasks on Sybase SQL Server.

ITO
See **OVO**.

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 *(Microsoft Exchange Server specific term)*
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.

lock name
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.

log_full shell script *(Informix Server UNIX specific term)*
A script provided by ON-Bar that you can use to start backing up logical log files when Informix Server issues a logfull event alarm. The Informix Server ALARMPROGRAM configuration parameter defaults to the <*INFORMIXDIR*>/etc/log_full.sh, where <*INFORMIXDIR*> is the Informix Server home directory. If you do not want logical logs to be backed up continuously, set the ALARMPROGRAM configuration parameter to <*INFORMIXDIR*>/etc/no_log.sh.

logging level
The logging level determines the amount of details on files and directories written to the IDB during backup, object copying, or object consolidation. You can always restore your data, regardless

of the logging level used during backup. Data Protector provides four logging levels: Log All, Log Directories, Log Files, and No Log. The different logging level settings influence the IDB growth, backup speed, and the convenience of browsing data for restore.

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 *(MS 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 (*Microsoft Exchange Server specific term)*
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 (*Microsoft Exchange Server specific term)*
A part of the Information Store that maintains information about user mailboxes. A mailbox store consists of a binary rich-text .edb file and a streaming native internet content .stm file.

Main Control Unit (MCU) *(HP StorageWorks Disk Array XP specific term)*
An HP StorageWorks XP disk array that contains the primary volumes for the Continuous Access configuration 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 **Enterprise Cell Manager**.

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.

MAPI *(Microsoft Exchange specific term)*
The MAPI (Messaging Application Programming Interface) is the programming interface that lets applications and messaging clients interact with messaging and information systems.

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.

media 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 Appendable, Non-Appendable, or Appendable for incrementals only.

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. *See also* **overwrite**.

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) *(Windows specific term)*
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, 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*)
A Mirror Unit number is 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 *(Windows specific term)*
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
A backup during which an application database cannot be used by the application.

- 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 (~minutes/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 (~seconds). Normal database operation can then be resumed for the rest of the backup process.

See also **zero downtime backup (ZDB)** and **online backup**.

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 *(Informix Server specific term)*
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:

- the onbar 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 *(Informix Server specific term)*
An environment variable that specifies the name of the active ONCONFIG configuration file. If the ONCONFIG environment variable is not present, Informix Server uses the configuration values from the onconfig file in the directory <*INFORMIXDIR*>\etc (on Windows) or <*INFORMIXDIR*>/etc/ (on UNIX).

online backup
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.

- For simple backup methods (non ZDB), backup mode is required for the whole backup period (~minutes/ 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 (~seconds). Normal database operation can then be resumed for the rest of the backup process.

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**.

online redo log *(Oracle specific term)*
Redo logs that have not been archived, but are either available to the instance for recording database activity or are filled and waiting to be archived or reused.
See also **archived redo log**.

OpC
See **OVO**.

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 *(Oracle specific term)*
Oracle Data Guard is Oracle's primary disaster recovery solution. Oracle Data Guard is able to maintain up to nine standby databases, each of which is a real-time copy of the production (primary) database, to protect against corruptions, data failures, human errors, and disasters. If a failure occurs on the production (primary) database, then a failover to one of the standby databases which becomes the new primary database is possible. In addition, planned downtime for maintenance can be reduced because the production processing can be moved from the current primary database to a standby database and back quickly.

Oracle instance *(Oracle specific term)*
Each installation of an Oracle database on one or more systems. One computer system can have several instances of a database running.

ORACLE_SID *(Oracle specific term)*
A unique name for an Oracle Server instance. To switch among Oracle Servers, specify the desired <*ORACLE_SID*>. The <*ORACLE_SID*> is included in the CONNECT DATA parts of the connect descriptor in a TNSNAMES.ORA file and in the definition of the TNS listener in the LISTENER.ORA 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**.

OVO
HP OpenView Operations 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 OVO management servers on HP-UX and Solaris. Earlier versions of OVO were called IT/Operation, Operations Center and Vantage Point Operations.
See also **merging**.

ownership
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: root.sys@<Cell Manager>, 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.

P1S file
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 <*Data_Protector_home*>\Config\Server\dr\p1s directory on a Windows Cell Manager or in /etc/opt/omni/server/dr/p1s directory on a UNIX Cell Manager with the filename recovery.p1s.

package *(MC/ServiceGuard and Veritas Cluster specific term)*
A collection of resources (for example volume groups, application services, IP names and addresses) that are needed to run a specific cluster-aware application.

pair status *(HP StorageWorks Disk Array XP specific term*)
A mirrored pair of disks can have various status values depending on the action performed on it. The three most important status values are:

- COPY - The mirrored pair is currently resynchronizing. 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 resynchronized 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) *(HP StorageWorks Disk Array XP specific term)*

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)**.

protection
See **data protection** and also **catalog protection**.

public folder store (*Microsoft Exchange Server specific term)*
The part of the Information Store that maintains information in public folders. A public folder store consists of a binary rich-text .edb file and a streaming native internet content .stm file.

public/private backed up data
When configuring a backup, you can select whether the backed up data will be:

- 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 *(HP StorageWorks Disk Array XP specific term)*
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 *(HP StorageWorks Disk Array XP specific term)*
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 *(HP StorageWorks specific term)*
The Remote Control Unit acts as a slave of an MCU in a CA configuration. In bidirectional configurations, the RCU can act as an MCU.

RDBMS
Relational Database Management System.

RDF1/RDF2 *(EMC Symmetrix specific term)*
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 *(Oracle specific term)*
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:

- 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 *(Oracle specific term)*
An Oracle database that contains a recovery catalog schema. You should not store the recovery catalog in your target database.

recovery files *(Oracle specific term)*
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) *(Oracle specific term)*
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 *(Oracle specific term)*
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 *(HP StorageWorks Disk Array XP specific term)*
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 *(Windows specific term)*
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 *(Windows specific term)*
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 *(ZDB specific term)*
An image, at a particular point in time, of the data in source volumes that contain user-specified backup objects. Depending on the hardware/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 a virtual copy (for example, a snapshot). From a host's perspective, on a basic UNIX or Windows system, the complete physical disk containing a backup object is replicated. However, if a volume manager is used on UNIX, the whole volume/disk group containing a backup object 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.

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 *(Windows specific term)*
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.

SAPDBA *(SAP R/3 specific term)*
An SAP R/3 user interface that integrates the BRBACKUP, BRARCHIVE, and BRRESTORE tools.

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) *(HP StorageWorks Disk Array XP specific term*)
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).**

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- and Post-exec script is a Data Protector unique identification of any session, including preview sessions. The session key is not recorded in the database, and it is used for specifying options for the omnimnt, omnistat and omniabort CLI commands.

shadow copy *(MS VSS specific term)*
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 changes as the backup process continues, but the shadow copy of the volume remains constant.
See also **Microsoft Volume Shadow Copy service**.

shadow copy provider *(MS VSS specific term)*
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 *(MS VSS specific term)*
A collection of shadow copies created at the same point in time.
See also **shadow copy**.

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.

Site Replication Service *(Microsoft Exchange Server specific term)*
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**.

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 -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.

standard security *(MS SQL specific term)*
Standard security uses the login validation process of the Microsoft SQL Server for all connections. Standard security is useful in network environments with a variety of clients, some of which may not support trusted connections. It also provides backward compatibility for older versions of the Microsoft SQL Server.
See also **integrated security**.

Storage Group
(Microsoft Exchange Server specific term)
A collection of databases (stores) that share a common set of transaction log files. Exchange manages each storage group with a separate server process.

StorageTek ACS library
(StorageTek specific term)
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 *(ZDB specific term)*
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 *(Sybase specific term)*
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 *(Sybase specific term)*
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) *(EMC Symmetrix specific term)*
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 *(Oracle specific term)*
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 *(Sybase specific term)*
The four system databases on a newly installed Sybase SQL Server are the:

- master database (master)
- temporary database (tempdb)
- system procedure database (sybsystemprocs)
- model database (model).

system disk
A system disk is a disk containing operating system files. Microsoft terminology defines the system disk as a disk containing the files required for initial step of boot process.

system partition
A system partition is a partition containing operating system files. Microsoft terminology defines a system partition as a partition containing the files required for initial step of boot process.

System State *(Windows specific term)*
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 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 (*Windows specific term)*
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 *(Oracle specific term)*
In RMAN, the target database is the database that you are backing up or restoring.

target (R2) device *(EMC Symmetrix specific term)*
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 *(Disaster Recovery specific term)*
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 *(ZDB specific term)*
A storage volume to which data is replicated.

Terminal Services *(Windows specific term)*
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 *(MS SQL Server specific term)*
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 *(EMC Symmetrix specific term)*
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 *(Oracle and SAP R/3 specific term)*
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 *(Sybase and SQL specific term*)
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 (*Data Protector specific term)*
Keeps 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 *(Sybase specific term)*
A system table in which all changes to the database are automatically recorded.

transportable snapshot *(MS VSS specific term)*
A shadow copy that is created on the

application system and can be presented to the backup system which performs the backup.
See also **Microsoft Volume Shadow Copy service (VSS)**.

TSANDS.CFG file *(Novell NetWare specific term)*
A file that allows you to specify the names of containers where you want backups to begin. It is text file located in the SYS:SYSTEM\TSA directory on the server where TSANDS.NLM is loaded.

unattended operation
See **lights-out operation**.

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 *(Windows specific term)*
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)

(HP StorageWorks EVA specific term)
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 *(MS SQL Server specific term)*
This is a SQL Server programming interface that allows fast backup and restore of large databases.

virtual disk *(HP StorageWorks EVA specific term)*
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 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.

volser *(ADIC and STK specific term)*
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 mountpoint (*Windows specific term*)
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 **OVO**.

VSS
See **Microsoft Volume Shadow Copy service**.

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 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
(MS VSS specific term)
A process that initiates change of data on the original volume. Writers are typically applications or system services that write persistent information on a volume. Writers also participate in the shadow copy synchronization process by assuring data consistency.

XBSA interface *(Informix Server specific term)*
ON-Bar and Data Protector communicate with each other through the X/Open Backup Services Application Programmer's Interface (XBSA).

XCopy engine *(direct backup specific term)*
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, and 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**.