

HP OpenView Storage Data Protector Integration Guide

for

**Oracle
SAP**

Manual Edition: October 2004



Manufacturing Part Number: B6960-90109

Release A.05.50

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

Legal Notices

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

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

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

UNIX® is a registered trademark of The Open Group.

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

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

Java™ is a U.S. trademark of Sun Microsystems, Inc.

ARM® is a registered trademark of ARM Limited.

1. Integrating Oracle and Data Protector

In This Chapter	2
Overview	3
Prerequisites and Limitations	6
Integration Concept	8
Data Protector Oracle Configuration Files	14
Setting, Retrieving, and Listing Data Protector Oracle Configuration Files Parameters Using the CLI	16
Configuring the Integration	21
Linking Oracle with the Data Protector Database Library on UNIX Systems	22
Configuring an Oracle User in Data Protector on UNIX Systems	29
Configuring the Oracle Server	30
Configuring an Oracle Backup	43
Creating a New Template	43
Creating a Data Protector Oracle Backup Specification	44
Oracle Specific Backup Options	48
Testing the Integration	55
Backing Up an Oracle Database	58
Scheduling a Backup	62
Running an Interactive Backup	63
Starting Oracle Backup Using the Oracle Recovery Manager (RMAN)	65
Restoring an Oracle Database	72
Restoring Oracle Using the Data Protector GUI	74
Restoring Oracle Using RMAN	91
Restoring Oracle Using CLI	103
Restoring Using Another Device	104
Disaster Recovery	105
Monitoring an Oracle Backup and Restore	106
Monitoring Current Sessions	106
Viewing Previous Sessions	107
Using Oracle After Removing the Data Protector Oracle Integration on UNIX Systems	108
Removing the Data Protector Oracle Integration Link on HP-UX Systems	108
Removing the Data Protector Oracle Integration Link on Solaris and other UNIX Systems	109
Oracle8i/9i/10g RMAN Metadata and Data Protector Media Management Database Synchronization	111
Troubleshooting	113
Before You Begin	113

Contents

Using Oracle After Removing the Data Protector	
Oracle Integration on UNIX Systems	113
Setting Up the Environment Variables	114
General Troubleshooting	114
Checking Prerequisites Related to the Oracle Side of the Integration on UNIX Systems	115
Checking Prerequisites Related to the Oracle Side of the Integration on Windows Systems	119
Configuration Problems on UNIX Systems	123
Configuration Problems on Windows Systems	126
Backup Problems on UNIX Systems	128
Backup Problems on Windows	132
Restore Problems	134

2. Integrating SAP R/3 and Data Protector

In This Chapter	142
Overview	143
Prerequisites and Limitations	145
Integration Concept	146
Data Protector SAP R/3 Configuration File	155
Setting, Retrieving, Listing, and Deleting Data Protector SAP R/3 Configuration File Parameters Using the CLI	158
Configuring the Integration	162
Configuring an SAP R/3 User in Data Protector (UNIX Systems Only)	162
Configuring an SAP R/3 Database Server	164
Configuring an SAP R/3 Backup	176
Creating a New Template	176
Creating a Data Protector SAP R/3 Backup Specification	177
SAP R/3 Specific Backup Options	183
Creating or Modifying the Parameter File on the SAP R/3 Database Server	187
Backing Up Using Recovery Manager	188
Manual Balancing of Files into Subsets	190
Creating an SAP /R3 Backup Specification for Manual Balancing	191
Testing the Integration	192
Backing Up an SAP R/3 Database	194
Scheduling a Backup	195
Running an Interactive Backup	197
Using SAP R/3 Commands	199
Restoring an SAP R/3 Database	201

Considerations	201
Finding Information Needed for Restore	201
Restoring Using the Data Protector GUI	201
Restoring Using the Data Protector CLI	203
Restoring Using the SAP R/3 Commands	204
Using Another Device	204
Disaster Recovery	205
Monitoring an SAP R/3 Backup and Restore	207
Monitoring Current Sessions	207
Viewing Previous Sessions	208
Troubleshooting	209
Before You Begin	209
General Troubleshooting	209
Troubleshooting on Windows Systems	210
Troubleshooting on UNIX Systems	220
Examples of SAP R/3 Database Restore	234
Preparing the SAP R/3 Database for Restore	234
Example of Full Database Restore and Recovery	236
Example of Partial Restore	240
Example of Lost Files Restore	241
Example of Archive Log Files Restore	242

3. Integrating SAP DB and Data Protector

In This Chapter	246
Prerequisites and Limitations	247
Overview	248
Integration Concept	252
Data Protector SAP DB Configuration File	255
Setting, Retrieving, and Listing Data Protector SAP DB Configuration File Parameters Using the CLI	256
Configuring the Integration	259
Configuring Users	259
Configuring an SAP DB Backup	260
SAP DB Specific Backup Options	265
Modifying the Configuration of an SAP DB Instance in Data Protector	267
Checking the Configuration of an SAP DB Instance	271
Testing the Integration	273
Backing Up an SAP DB Database	276

Contents

Scheduling an Existing Backup Specification	276
Running an Interactive Backup Using the Data Protector GUI	279
Running an Interactive Backup Using the Data Protector CLI	280
Running an Interactive Backup Using SAP DB Utilities	281
Restoring an SAP DB Database	285
Restore and Recovery Overview	285
SAP DB Migration Prerequisites	289
Restoring Using the Data Protector GUI	289
Restoring Using the Data Protector CLI	292
Restoring Using SAP DB Utilities	294
SAP DB Restore and Recovery Options	298
Using Another Device	303
Disaster Recovery	303
Monitoring an SAP DB Backup and Restore	305
Monitoring Current Sessions	305
Viewing Previous Sessions	306
Troubleshooting	308
SAP DB Troubleshooting	308
SAP DB Cluster-Related Troubleshooting	311

Glossary

Index

Printing History

The manual printing date and part number indicate its current edition. The printing date will change when a new edition is printed. Minor changes may be made at reprint without changing the printing date. The manual part number will change when extensive changes are made.

Manual updates may be issued between editions to correct errors or document product changes. To ensure that you receive the updated or new editions, you should 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

Conventions

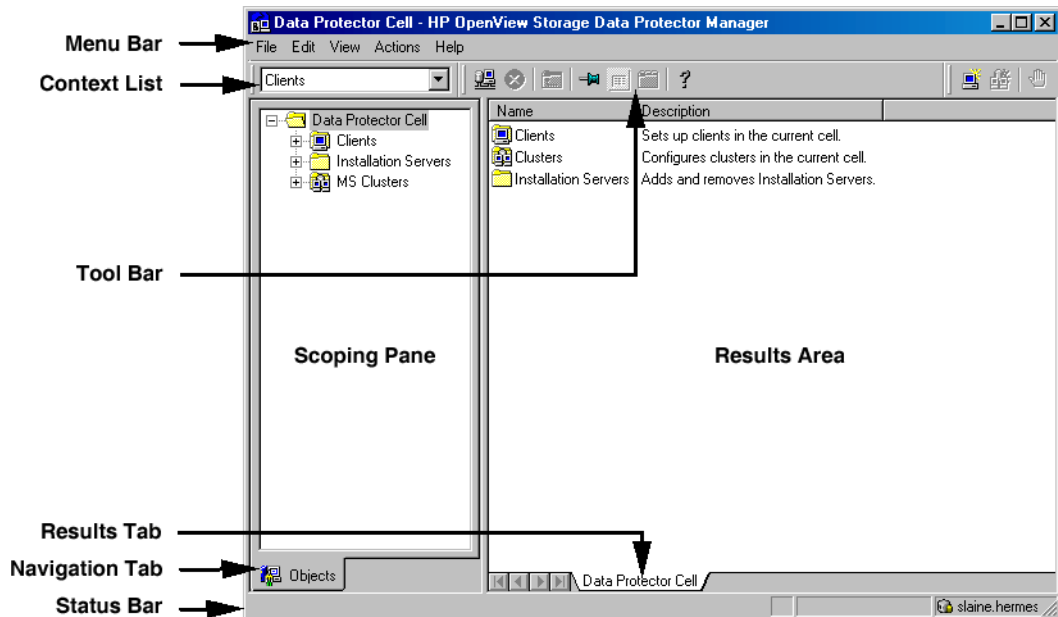
The following typographical conventions are used in this manual.

Table 2

Convention	Meaning	Example
<i>Italic</i>	Book or manual titles, and manual page names	Refer to the <i>HP OpenView Storage Data Protector Integration Guide</i> for more information.
	Provides emphasis	You <i>must</i> follow these steps.
	Specifies a variable that you must supply when entering a command	At the prompt type: <code>rlogin your_name</code> 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 <code>grep</code> command ...
	File and directory names	<code>/usr/bin/X11</code>
	Process names	Check to see if Data Protector <code>Inet</code> is running.
	Window/dialog box names	In the Backup Options dialog box...
	Text that you must enter	At the prompt, type: <code>ls -l</code>
Keycap	Keyboard keys	Press Return .

Data Protector provides a cross-platform (Windows and UNIX) graphical user interface. Refer to the *HP OpenView Storage Data Protector Administrator's Guide* 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://support.openview.hp.com/support.jsp>

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

Information about the latest Data Protector patches can be found at

http://support.openview.hp.com/patches/patch_index.jsp

For information on the Data Protector required patches, refer to the *HP OpenView Storage Data Protector Software Release Notes*.

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

http://ovweb.external.hp.com/lpe/doc_serv/

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://ovweb.external.hp.com/lpe/doc_serv/

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 *HP OpenView Storage Data Protector Administrator's Guide*.

HP OpenView Storage Data Protector Administrator's Guide

This manual describes typical configuration and administration tasks performed by a backup administrator, such as device configuration, media management, configuring a backup, and restoring data.

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 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 7/2000, Exchange Server 5.x, Exchange Server 2000/2003, and Volume Shadow Copy Service*

This manual describes the integrations of Data Protector with the following Microsoft applications: Microsoft Exchange Server 2000/2003, Microsoft Exchange Server 5.x, Microsoft SQL Server 7/2000, 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, IBM DB2, and Lotus Notes/Domino.

- *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, and Network Data Management Protocol.

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, HP OpenView Service Desk, 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.

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 Software Release Notes

This manual gives a description of new features of HP OpenView Storage Data Protector A.05.50. 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.openview.hp.com/products/datapro/spec_0001.html.

Online Help

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

In This Book

The *HP OpenView Storage Data Protector Integration Guide for Oracle and SAP* describes how to configure and use Data Protector with Oracle and SAP applications.

Audience

This manual is intended for backup administrators who are responsible for the planning, setup, and maintenance of network backups. It assumes that 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 141.
- Chapter 3** “Integrating SAP DB and Data Protector” on page 245.
- Glossary** Definition of terms used in this manual.

The integrations of Data Protector with the following database applications are described in the *HP OpenView Storage Data Protector Integration Guide for Microsoft Applications: SQL Server 7/2000, Exchange Server 5.x, Exchange Server 2000/2003, and Volume Shadow Copy Service*:

- Microsoft SQL Server 7.0/2000
- Microsoft Exchange Server 5.x
- Microsoft Exchange Server 2000/2003
- Microsoft Volume Shadow Copy Service

The integrations of Data Protector with the following database applications are described in the *HP OpenView Storage Data Protector Integration Guide for IBM Applications: Informix, DB2, and Lotus Notes/Domino*:

- Informix
- IBM DB2 UDB
- Lotus Notes/Domino

The integrations of Data Protector with the following database applications are described in the *HP OpenView Storage Data Protector Integration Guide for Sybase, Network Node Manager, and Network Data Management Protocol*:

- Sybase
- Network Node Manager
- Network Data Management Protocol

The integrations of Data Protector ZDB integrations with the following applications or operating system services are described in the *HP OpenView Storage Data Protector Zero Downtime Backup Integration Guide*:

- Oracle
- SAP R/3
- Microsoft SQL Server 7.0/2000
- Microsoft Volume Shadow Copy Service
- Microsoft Exchange Server 2000/2003

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:

- “Overview” on page 3
- “Prerequisites and Limitations” on page 6
- “Integration Concept” on page 8
- “Data Protector Oracle Configuration Files” on page 14
- “Configuring the Integration” on page 21
- “Configuring an Oracle Backup” on page 43
- “Backing Up an Oracle Database” on page 58
- “Restoring an Oracle Database” on page 72
- “Monitoring an Oracle Backup and Restore” on page 106
- “Using Oracle After Removing the Data Protector Oracle Integration on UNIX Systems” on page 108
- “Oracle8i/9i/10g RMAN Metadata and Data Protector Media Management Database Synchronization” on page 111
- “Troubleshooting” on page 113

Overview

Data Protector offers offline as well as online backup of the Oracle Server instances. In order 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.

Backup Types

You can perform the following types of backups using the Data Protector Oracle integration:

- 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 file

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

Why use the Data Protector User Interface?

Using the Data Protector Oracle integration with the Data Protector GUI or Recovery Manager (RMAN) offers several advantages over using RMAN alone:

- Central Management

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

- Media Management

Data Protector has an advanced media management system that allows users to monitor media usage, set the protection for stored data, as well as organize and manage devices in media pools.

- 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 the administrator to automate backups to run periodically. With the Data Protector Scheduler, the backups you configure run unattended at specified times, as long as the devices and media are properly set.

- Device Support

Data Protector supports a wide range of devices, from files and standalone drives to complex multiple drive libraries. See the *HP OpenView Storage Data Protector Software Release Notes* for a complete list of supported backup devices.

- 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 the administrator 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 Data Protector database and in the Oracle recovery catalog database.

The administrator is thus provided 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

- You need a license to use the Data Protector Oracle 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 Oracle Server and Data Protector systems. Refer to the:
 - *HP OpenView Storage Data Protector Software Release Notes* or http://www.openview.hp.com/products/datapro/spec_0001.html 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.
 - *Oracle8i/9i Recovery Manager User's Guide and References* for Oracle concepts and backup/recovery strategies.
 - *Oracle8i/9i Backup and Recovery Guide* for the configuration and use of Recovery Manager, as well as for Oracle backup terminology and concepts.
 - *Oracle8 Enterprise Manager User's Guide* for information about backup and recovery with the Oracle Enterprise Manager, as well as information about Oracle8 Server Manager and about Oracle9i/10g SQL Plus.
-
- An instance of Oracle must be created and configured on the host to which you want to restore the database.
- If you want to restore tablespaces or datafiles, the Oracle database to which they belong must exist.
- The database must be in `Mount` state if the whole database is being restored, or in `NoMount` state if just the control file is being restored.

- The Oracle database *user* used by this integration to connect to the target Oracle 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 .

- On Windows, the `<ORACLE_HOME> \bin` directory must be included in the PATH system environmental variable.
- It is assumed that you are familiar with the Oracle database administration and the basic Data Protector functionality.
- In case of XXXXXXXXXX, 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. Refer to “Backup of Archive Logs on RAC Cannot be Performed” on page 131.

Limitations

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

- When restoring a database to a host other than the one on which the database originally resided, the instance name on the host must be the same as the original database instance name.
- When using RMAN scripts in Oracle backup specifications, double quotes (") must not be used - single quotes (') should be used instead.
- You cannot edit RMAN scripts before they are executed when performing a restore operation.

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 utility. 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 works with Oracle to manage all aspects of backup, restore, and recovery operations on the Oracle target database. The following functionality is offered:

- Database startup and shutdown
- Backups (backup and copy)
- Recovery (restore and recovery)
- Catalog maintenance, catalog analysis and display
- Stored script maintenance, miscellaneous operations

How Does the Integration Work?

The Data Protector Oracle Integration agent for Oracle uses RMAN functionality to direct the Oracle server processes on the target database to perform backup, restore, and recovery operations. Furthermore, it maintains the required information about the Target Databases in the recovery catalog, the Oracle central repository of information, or in the control file of a particular Target Database.

The Data Protector Oracle Integration agent for Oracle uses the information in the recovery catalog to determine how to execute the requested backup and restore operations.



- The physical schema of the Oracle Target Databases

- Archived Redo Logs
- Runtime information on backup and restore jobs.

You can back up and restore Oracle control files, datafiles, and Archived Redo Logs using the Data Protector integration into RMAN.

The interface from the Oracle server processes to Data Protector is provided by the Data Protector Database Library (Media Management Layer or Media Management Library), 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

Oracle full and incremental backup types (up to incremental level 4) can be performed using this integration. As opposed to a full backup, where all data blocks per data file are backed up, [REDACTED]

[REDACTED] The data being backed up, regardless of the backup type (full or incremental), is selected and controlled by Oracle.

NOTE

[REDACTED] since the Data Protector incremental backup concept is different from the Oracle incremental backup concept.

A backup that includes all data files 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. [REDACTED] whereas for offline backup, objects need to be prepared for backup by 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 the **Data Protector Database Library**, 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

Before a backup session begins, the Oracle Target Database instance is switched into backup mode.

A Data Protector scheduled or interactive backup is triggered by the Data Protector

[REDACTED]

This user must be defined as the owner of the Data Protector Oracle8 backup specification. Further on,

[REDACTED]

[REDACTED] and then monitors the backup process.

[REDACTED]

[REDACTED]

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.

Once the backup has finished, the Oracle database is switched into normal mode.

Restore Flow

A restore session can be started from the Data Protector GUI, by issuing the RMAN restore command from the RMAN command line, or from the Oracle Enterprise Manager graphical user interface. You must specify which objects are to be restored.

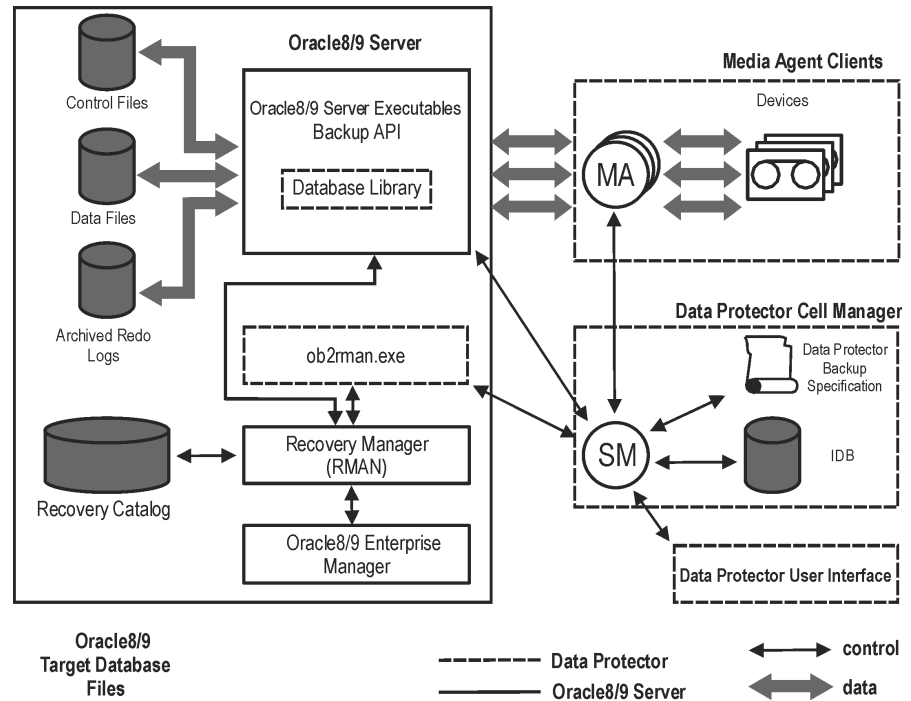
A restore from the Data Protector user interface is triggered by the Data Protector Restore Session Manager, which starts the `ob2rman.exe` script (UNIX systems) or binary (Windows systems). The `ob2rman.exe` script (UNIX systems) or binary (Windows systems) prepares the environment to start the restore, 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 the Database Library. The Database Library 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 the Database Library 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 the Database Library. 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 12, and the related terms are explained in the following table.

Figure 1-1 Data Protector Oracle Integration Concept



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.
- Database Library** The Data Protector 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.

Data Protector Oracle Configuration Files

Data Protector stores the Oracle integration parameters in two files on the Cell Manager:

- For every configured Oracle *instance* in the:
/etc/opt/omni/server/integ/config/Oracle8/<client_name>%<ORACLE_SID> file (HP-UX and Solaris systems), or in the
<Data_Protector_home>\Config\Server\Integ\Config\Oracle
8\<client_name>%<ORACLE_SID> file (Windows systems).

The parameters stored in the [REDACTED] are:

- Oracle home directory,
- Oracle version,
- encoded connection strings to the target database and recovery catalog and
- the variables which need to be exported prior to starting backup, and which affect the Oracle instance.

- Oracle *global* integration parameters in the:
/etc/opt/omni/server/integ/config/Oracle8/<client_name>%_OB2_GLOBAL file (HP-UX and Solaris systems), or in the
<Data_Protector_home>\Config\server\integ\config\oracle
8\<client_name>%_OB2_GLOBAL file (Windows systems).

The parameters stored in the [REDACTED] are:

- instance list (all Oracle instances on the Oracle server) and
- the variables that need to be exported prior to starting a backup, and which affect every Oracle instance on the Oracle server.

The configuration parameters are written to the Data Protector Oracle configuration files during configuration of the integration.

Global Configuration File Syntax

The syntax of the file is as follows:

IMPORTANT

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

```
INSTANCE_LIST=('<ORACLE_SID1>' [, '<ORACLE_SID2>' , '<ORACLE_SID3>' .  
..]);  
Environment={  
    [<ENV var1>=<value1>;]  
    [<ENV var2>=<value2>;]  
    ...]  
}
```

Example of Global Configuration File

This is an example of the Data Protector Oracle global configuration file:

```
INSTANCE_LIST=('ICE', 'BUREK');  
Environment={  
    OB2OPTS='-debug 1-200 GLOB.txt';  
}
```

IMPORTANT

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

Instance Configuration File Syntax

The syntax of the file is as follows:

```
TGTLogin=<encoded connection string to target database>;  
RCVLogin=<encoded connection string to recover catalog  
database>;  
ORACLE_HOME=<Oracle's instance home directory>;  
[ORACLE_VERSION=<Version of Oracle software>;]  
ORACLE_DBID=<Oracle database ID>;  
Environment={  
    [<ENV var1>=<value>;]
```

```
[<ENV var2>=<value>;  
...]  
}
```

Example of Instance Configuration File

This is an example of the Data Protector Oracle instance configuration file:

```
TGTLogin='EIBBKIB...BDGBBGFBBMFBB';  
RCVLogin='DIBBOH...BEFBBCFBBFGBB';  
ORACLE_HOME='/opt/oracle/product/8.1.6';  
ORACLE_VERSION='8.1.6 64bit';  
ORACLE_DBID='369153078';  
Environment={  
NLS_LANG='AMERICAN_AMERICA.US7ASCII';  
}
```

NOTE

The NLS_LANG variable is to be set only if the recovery catalog is used and if it uses a character set that is not the same as the character set used by the Oracle database to be backed up or restored. The value of the NLS_LANG must be set as the name of the character set used by the recovery catalog.

Setting, Retrieving, and Listing Data Protector Oracle Configuration Files Parameters Using the CLI

The Data Protector Oracle configuration files parameters are normally written to the Data Protector Oracle configuration files after the completed configuration of the Oracle instance in Data Protector.

On UNIX, the variable definitions that are command or shell-based must be entered in the following file, since such variable definitions are not possible in the Data Protector Oracle configuration files:

```
/etc/opt/omni/server/integ/config/oracle8/<ORACLE_SID>/.profile  
file (HP-UX and Solaris systems) or
```

```
/usr/omni/config/oracle8/<ORACLE_SID>/.profile (other UNIX  
systems)
```

The `util_cmd` Command

You can set, retrieve, or list the Data Protector Oracle configuration files' 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 Oracle client.

You can upload a configuration file to or download it from a Cell Manager using the `-putfile` or `-getfile` options.

The command resides in the `/opt/omni/sbin` (HP-UX and Solaris systems), `/usr/omni/bin` (other UNIX systems) or `<Data_Protector_home>\bin` (Windows systems) directory.

Cluster-Aware

[REDACTED]. 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] Oracle8 <Oracle8_instance> [-local \  
<filename>]
```

```
util_cmd -getopt[ion] [Oracle8 <Oracle8_instance>] \  
<option_name> [-sub[list] <sublist_name>] [-local \  
<filename>]
```

```
util_cmd -putopt[ion] [Oracle8 <Oracle8_instance>] \  
<option_name> [<option_value>] [-sub[list] <sublist_name>] \  
[-local <filename>]
```

```
util_cmd -getfile [Oracle8 <filename>] [-local <filename>]
```

```
util_cmd -putfile [Oracle8 <filename>] [-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.

`<filename>` is the name of the configuration file that is downloaded from or uploaded to `/etc/opt/omni/server/config/Oracle8` (on UNIX) or `<Data_Protector_home>\Config\server\integ\config\Oracle8` (on Windows)

`[-local <filename>]` specifies one of the following:

- When used with the `-getconf[ig]` option, it specifies the filename of the file that the command output is written to. If the `-local` option is not specified, the output is written to the standard output.
- When with the `-getopt[ion]`, it specifies the filename of the file from which the parameter and its value are to be taken from and then written to the standard output. If the `-local` option is not specified, the parameter and its value are taken from one of the Data Protector Oracle configuration files and then written to the standard output.
- When 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 one of the Data Protector Oracle configuration files.
- When used with the `-getfile` or `-putfile` option, it specifies the local directory to/from where the file will be copied.

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, one of the Data Protector Oracle configuration files is missing on the Cell Manager, etc.

Setting Parameters

To set the OB2OPTS and NLS_LANG parameters for the Oracle instance ICE, use the following commands:

NOTE

The NLS_LANG variable is to be set only if the recovery catalog is used and if it uses a character set that is not the same as the character set used by the Oracle database to be backed up or restored. The value of the NLS_LANG must be set as the name of the character set used by the recovery catalog.

HP-UX and Solaris

```
/opt/omni/sbin/util_cmd -putopt Oracle8 ICE OB2OPTS \  
'-debug 1-200 INSTANCE.txt' -sublist Environment  
  
/opt/omni/sbin/util_cmd -putopt Oracle8 ICE NLS_LANG \  
AMERICAN_AMERICA.US7ASCII -sublist Environment
```

Other UNIX

```
/usr/omni/bin/util_cmd -putopt Oracle8 ICE OB2OPTS '-debug \  
1-200 INSTANCE.txt' -sublist Environment  
  
/usr/omni/bin/util_cmd -putopt Oracle8 ICE NLS_LANG \  
AMERICAN_AMERICA.US7ASCII -sublist Environment
```

Windows

```
<Data_Protector_home>\bin\util_cmd -putopt Oracle8 ICE  
OB2OPTS "-debug 1-200 INSTANCE.txt" -sublist Environment  
  
<Data_Protector_home>\bin\util_cmd -putopt Oracle8 ICE  
NLS_LANG AMERICAN_AMERICA.US7ASCII -sublist Environment
```

Retrieving Parameters

To retrieve the value of the OB2OPTS parameter for the instance ICE, use the following command:

- On HP-UX and Solaris: `/opt/omni/sbin/util_cmd -getopt Oracle8 ICE OB2OPTS \ -sublist Environment`
- On other UNIX: `/usr/omni/bin/util_cmd -getopt Oracle8 ICE OB2OPTS \ -sublist Environment`
- On Windows: `<Data_Protector_home>\bin\util_cmd -getopt Oracle8 ICE OB2OPTS -sublist Environment`

Listing Parameters To list the Data Protector configuration files' parameters for the instance ICE, use the following command:

- On HP-UX and Solaris: `/opt/omni/sbin/util_cmd -getconf Oracle8 ICE`
- On other UNIX: `/usr/omni/bin/util_cmd -getconf Oracle8 ICE`
- On Windows: `<Data_Protector_home>\bin\util_cmd -getconf Oracle8 ICE`

Configuring the Integration

It is assumed that the installation of Data Protector software components on the Oracle Server system was successful.

Prerequisites

The following requirements must be met before you start configuring the Oracle integration:

- The Oracle Server software must be installed and the Oracle Target Database must be online.

- [REDACTED]

Refer to the *Oracle8i Recovery Manager User's Guide and References* for more information about different connection options.

- [REDACTED] Refer to the *Oracle8i Recovery Manager User's Guide and References* for further information about how to create a recovery catalog.

- [REDACTED]

Refer to “Setting, Retrieving, and Listing Data Protector Oracle Configuration Files Parameters Using the CLI” on page 16 for more information on how to do that.

- [REDACTED] for the Oracle Target Database and the recovery catalog, if used.

Refer to the Oracle documentation for further information.

Refer to “Troubleshooting” on page 113 for details about how to check all the prerequisites listed above.

If you are using a shared server, create a new service name that will use a dedicated server process instead of the dispatcher for connecting. Add the (SERVER=DEDICATED) parameter to the CONNECT_DATA part of the connection string.

- [REDACTED] properly. Refer to the Oracle documentation for more information.

-  Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for information on how to change the Data Protector Inet service account.

Configuration Overview

The following list gives an overview of the global tasks for configuring the Data Protector Oracle integration:

1. “Linking Oracle with the Data Protector Database Library on UNIX Systems” on page 22.
2. “Configuring an Oracle User in Data Protector on UNIX Systems” on page 29.
3. “Configuring the Oracle Server” on page 30.
4. “Configuring an Oracle Backup” on page 43.

Linking Oracle with the Data Protector Database Library on UNIX Systems

In order to use the Data Protector Oracle integration, you need to manually link the Oracle server software and the Data Protector Database Library on the Data Protector Oracle Server system.

The Data Protector Database Library 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 the Data Protector Database Library. You must rebuild (Oracle8) or re-link (Oracle8i/9i) 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 Systems” on page 108 for more information on removing the integration link.

MC/ServiceGuard Clients

When linking Oracle with the Data Protector database library, link it on all nodes.



Linking Oracle with the Data Protector Database Library on HP-UX systems

On Oracle Server systems running on HP-UX, the Data Protector Database Library `libob2oracle8.sl` (`libob2oracle8_64bit.sl`) is located in the `/opt/omni/lib` directory.

NOTE

On the IA-64 architecture, the extension for shared objects is `.so`. The filename for the Data Protector Database Library for the 32-bit version of Oracle is `libob2oracle8.so`, whereas for the 64-bit version of Oracle it is `libob2oracle8_64bit.so`.

Proceed as follows:

1. On the Oracle Server system, connect to the Oracle database as an Oracle operating system user and shut down all Oracle instances.

Oracle8

2. If you have Oracle8 installed, perform the following:
 - a. Change to the `<ORACLE_HOME>/rdbms/lib` directory:

```
cd <ORACLE_HOME>/rdbms/lib
```
 - b. If you have a 32-bit integration, execute the following command:

```
make -f ins_rdbms.mk LLIBMM= LLIBOBK="-L/opt/omni/lib  
/opt/omni/lib/libob2oracle8.sl" ioracle,
```

while for a 64-bit integration, the respective command is:

```
make -f ins_rdbms.mk LLIBMM= LLIBOBK="-L/opt/omni/lib  
/opt/omni/lib/libob2oracle8_64bit.sl" ioracle
```

Oracle8i/9i

3. If you have Oracle8i/9i installed, perform the following:
 - a. Change to the `<ORACLE_HOME>/lib` directory:

```
cd <ORACLE_HOME>/lib (32-bit Oracle),  
cd <ORACLE_HOME>/lib64 (64-bit Oracle8i) or  
cd <ORACLE_HOME>/lib (64-bit Oracle9i).
```
 - b. Perform this step only if the `libobk.sl` file is already created in the `<ORACLE_HOME>/lib` directory. If this is not the case, skip this step.

Execute the following command:

```
mv libobk.sl libobk.sl.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 the `libobk.sl.orig` file.

c. `ln -s /opt/omni/lib/libob2oracle8.sl libobk.sl` (32-bit Oracle) or

`ln -s /opt/omni/lib/libob2oracle8_64bit.sl libobk.sl` (64-bit Oracle).

4. Start all Oracle instances.

Oracle8

5. If you have Oracle8 installed, check if the `libob2oracle8.sl` (`libob2oracle8_64bit.sl`) file is linked with the Oracle8 executable using the following command:

```
/usr/bin/chatr <ORACLE_HOME>/bin/oracle (32-bit Oracle8)
```

```
/usr/ccs/bin/ldd <ORACLE_HOME>/bin/oracle (64-bit Oracle8)
```

The `/opt/omni/lib/libob2oracle8.sl` (`libob2oracle8_64bit.sl`) library must be listed in the shared library list.

The `chatr` command also verifies whether the shared library dynamic path `SHLIB_PATH` is enabled.

Example

The following is an example extract of the command output:

```
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
```

```
dynamic /usr/lib/librt.2
```

```
dynamic /usr/lib/libnss_dns.1
dynamic /usr/lib/libdld.2
shared library binding:
    deferred
static branch prediction disabled
kernel assisted branch prediction enabled
lazy swap allocation disabled
text segment locking disabled
data segment locking disabled
data page size: D (default)
instruction page size: D (default)
```

The line starting with the `SHLIB_PATH` entry should be as given in the example above. If this line is different, then enable the shared library dynamic path as follows:

- a. Stop all Oracle8 instances.
- b. `$ chatr +s enable $ORACLE_HOME/bin/oracle`
- c. Start all Oracle8 instances.

Linking Oracle with the Data Protector Database Library on Solaris Systems

On Oracle Server systems running on Solaris systems, the Data Protector Database Library `libob2oracle8.so` (`libob2oracle8_64bit.so`) is located in the `/opt/omni/lib` directory.

Proceed as follows:

1. On the Oracle Server system, connect to the database as an Oracle operating system user and shut down all Oracle instances.

Oracle8

2. If you have Oracle8 installed, proceed as follows:
 - a. Change to the `<ORACLE_HOME>/rdbms/lib` directory:
`cd <ORACLE_HOME>/rdbms/lib`

- b. Execute the following command:

```
make -e -f ins_rdbms.mk LLIBMM=  
LLIBOBK="-L/opt/omni/lib  
/opt/omni/lib/libob2oracle8.so" ioracle (32-bit Oracle). or  
  
make -e -f ins_rdbms.mk LLIBMM=  
LLIBOBK="-L/opt/omni/lib  
/opt/omni/lib/libob2oracle8_64bit.so" ioracle (64-bit  
Oracle).
```

Oracle8i/9i

3. If you have Oracle8i/9i installed, proceed as follows:

- a. Change to the `<ORACLE_HOME>/lib` directory:

```
cd <ORACLE_HOME>/lib (32-bit Oracle),  
cd <ORACLE_HOME>/lib64 (64-bit Oracle8i) or  
cd <ORACLE_HOME>/lib (64-bit Oracle9i).
```

- b. Perform this step only if the `libobk.so` file is already created in the `<ORACLE_HOME>/lib` directory. If this is not the case, skip this step.

Execute the following command:

```
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 the `libobk.so.orig` file.

- c. `ln -s /opt/omni/lib/libob2oracle8.so libobk.so` (32-bit Oracle). or

```
ln -s /opt/omni/lib /libob2oracle8_64bit.so libobk.so  
(64-bit Oracle).
```

4. Start all Oracle instances.

Oracle8

5. If you have Oracle8 installed, check if the `libob2oracle8.so` file is linked with the Oracle8 executable using the following command:

```
/opt/bin/ldd <ORACLE_HOME>/bin/oracle
```

The `/opt/omni/lib/libob2oracle8.so` library has to be listed as required by the Oracle8 executable, and the `LD_LIBRARY_PATH` must be enabled. If the `/opt/omni/lib/libob2oracle8.so` library is not listed, then you need to re-build the Oracle binary as follows:

```
make -f ins_rdbms.mk LLIBMM= LLIBOBK="-L/opt/omni/lib  
/opt/omni/lib/libob2oracle8.so" ioracle (32-bit Oracle). or  
make -f ins_rdbms.mk LLIBMM= LLIBOBK="-L/opt/omni/lib  
/opt/omni/lib/libob2oracle8_64bit.so" ioracle (64-bit  
Oracle).
```

Linking Oracle with the Data Protector Database Library on Other UNIX Systems

On Oracle Server systems running on other UNIX systems, the Data Protector Database Library `libob2oracle8.so` (`libob2oracle8_64bit.so`) is located in the `/usr/omni/lib` directory.

NOTE

On AIX systems, the extensions for shared objects is `.a`. The filename for the Data Protector Database Library for the 32-bit version of Oracle is `libob2oracle8.a`, whereas for the 64-bit version it is `libob2oracle8_64bit.a`.

Proceed as follows:

1. On the Oracle Server system, connect to the database as an Oracle operating system user and shut down all Oracle instances.

Oracle8

2. If you have Oracle8 installed, proceed as follows:

- a. Change to the `<ORACLE_HOME>/rdbms/lib` directory:

```
cd <ORACLE_HOME>/rdbms/lib
```

- b. Execute the following command:

```
make -f ins_rdbms.mk LLIBMM= LLIBOBK="-L/usr/omni/lib  
/opt/omni/lib/libob2oracle8.so" ioracle (32-bit Oracle) or  
make -f ins_rdbms.mk LLIBMM= LLIBOBK="-L/usr/omni/lib  
/opt/omni/lib/libob2oracle8_64bit.so" ioracle (64-bit  
Oracle).
```

Oracle8i/9i

3. If you have Oracle8i/9i installed, proceed as follows:

a. Change to the `<ORACLE_HOME>/lib` directory:

```
cd <ORACLE_HOME>/lib (32bit Oracle),
```

```
cd <ORACLE_HOME>/lib64 (64bit Oracle8i) or
```

```
cd <ORACLE_HOME>/lib (64bit Oracle9i).
```

b. Perform this step only if the `libobk.so` file is already created in the `<ORACLE_HOME>/lib` directory. If this is not the case, skip this step.

Execute the following command:

```
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 the `libobk.so.orig` file.

c. `ln -s /usr/omni/lib/libob2oracle8.so libobk.so` (32-bit Oracle) or

```
ln -s /usr/omni/lib /libob2oracle8_64bit.so libobk.so  
(64-bit Oracle).
```

4. Start all Oracle instances.

Oracle8

5. If you have Oracle8 installed, check if the `libob2oracle8.so` file is linked with the Oracle8 executable using the following command:

```
/usr/bin/dump -H <ORACLE_HOME>/bin/oracle (AIX systems) or
```

```
/usr/bin/ldd -s <ORACLE_HOME>/bin/oracle (other UNIX  
systems).
```

The `/usr/omni/lib/libob2oracle8.so` library has to be listed as required by the Oracle8 executable and the `LD_LIBRARY_PATH` must be enabled. If the `/usr/omni/lib/libob2oracle8.so` library is not listed, then you need to re-build the Oracle binary as follows:

```
make -e -f ins_rdbms.mk LLIBMM= LLIBOBK="-L/usr/omni/lib  
/usr/omni/lib/libob2oracle8.so" ioracle (32-bit Oracle) or
```



```
make -e -f ins_rdbms.mk LLIBMM= LLIBOBK="-L/usr/omni/lib  
/usr/omni/lib/libob2oracle8_64bit.so" ioracle (64-bit  
Oracle).
```

Configuring an Oracle User in Data Protector on UNIX Systems

[REDACTED]

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

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

If properly configured, this user is allowed to back up or restore an Oracle database.

[REDACTED]

As the owner of the backup specification,

[REDACTED]

[REDACTED]

[REDACTED]

or

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

Figure 1-2

Finding the Oracle User

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


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 in order to be able to back up the Oracle database using Data Protector.

IMPORTANT

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

After the two users are added to either 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.

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

MC/ServiceGuard

in the cluster.

Refer to chapter “Configuring Users and User Groups” in the *HP OpenView Storage Data Protector Administrator’s Guide* for detailed information on Data Protector user rights and how to add a user to a user group.

Configuring the Oracle Server

Before You Begin

It is recommended that you configure and run a Data Protector filesystem backup of the Oracle Server system. A filesystem backup can be performed only if you have installed the Disk Agent on the Oracle Server system.

In case of problems, this type of backup is much easier to troubleshoot than an integration of the Oracle Server system with Data Protector.

Any device can be used for this test. Configure a standard filesystem backup, which can include one directory only. The test should include a partial restore to the Oracle Server system as well.

Configuring an Oracle Server system 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 online during the configuration procedure. The configuration must be done for each Oracle Server instance.

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.

The configuration is performed using the Data Protector User Interface.

Configuring the Oracle Server Using the Data Protector GUI

The procedure below describes the configuration of the Oracle Server system using the Data Protector GUI.

HP-UX with RAC

If you want to use virtual hostname and the global database name, [REDACTED] containing only the virtual IP and the virtual hostname parameters and distribute it among the RAC nodes.

RAC

[REDACTED] However, if you use an virtual hostname, then perform the configuration only on the node that will be used for backup/restore.

The configuration is performed at the same time that a new backup specification is created:

1. In the Data Protector Manager, switch to the Backup context.
2. In the Scoping Pane, expand Backup, and then Backup Specifications. Right-click Oracle Server.
3. Click Add Backup. In the Create New Backup dialog box, double-click any of the predefined backup templates.
4. In the Results Area, enter the following information:
 - Name of the Oracle Server system you want to configure. If the application is cluster-aware, select the virtual server of the Oracle resource group.



- Name of the Oracle Server instance.

or the name of the Oracle node instance if you selected the virtual server of the Oracle resource group.

- as described in “Configuring an Oracle User in Data Protector on UNIX Systems” on page 29. for detailed information on how to identify that user.

Figure 1-3 Specifying an Oracle Server System on UNIX

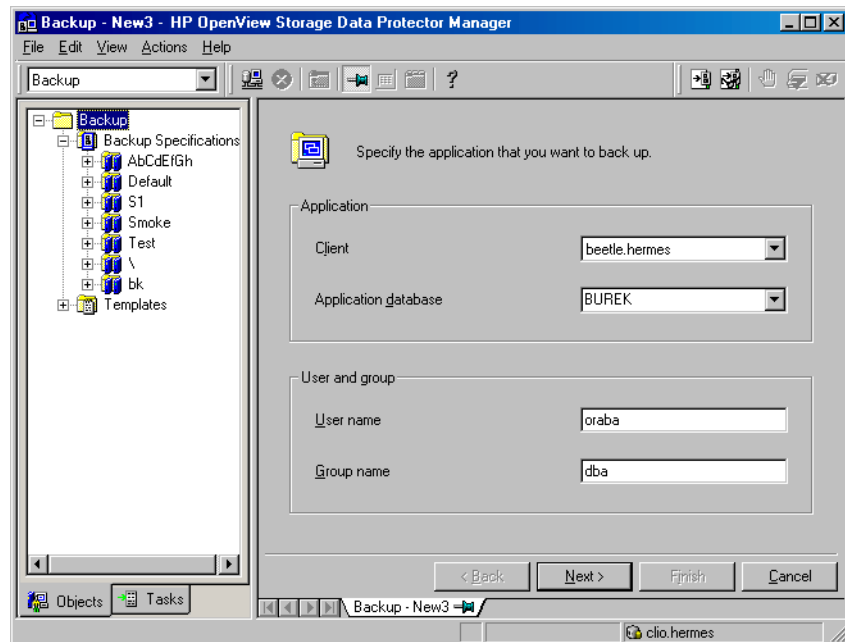
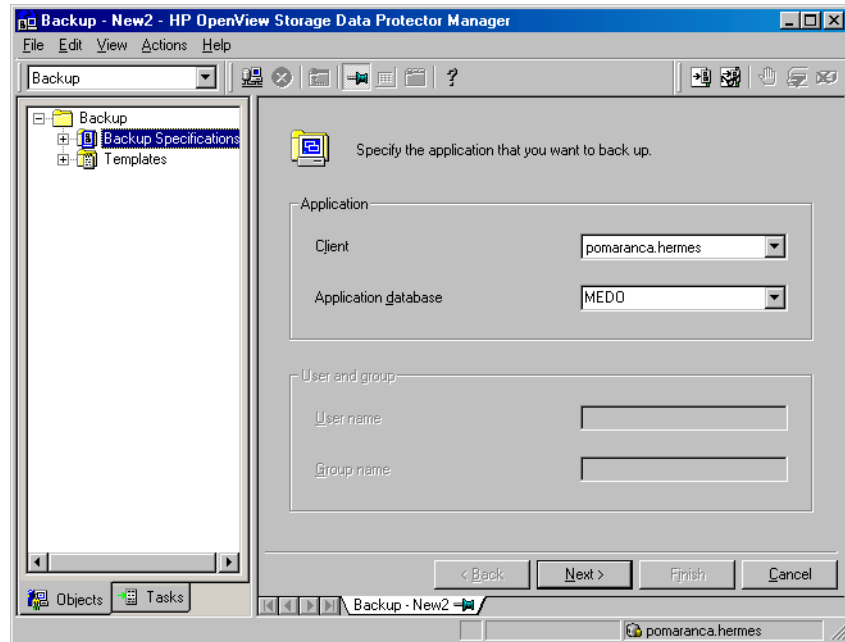


Figure 1-4 Specifying an Oracle Server System on Windows



Click Next. If the Oracle instance has not yet been configured, the Configure Oracle dialog window is displayed.

5. In the General property page, enter the Oracle Server home directory.

Figure 1-5 **Configuring Oracle on UNIX - General**

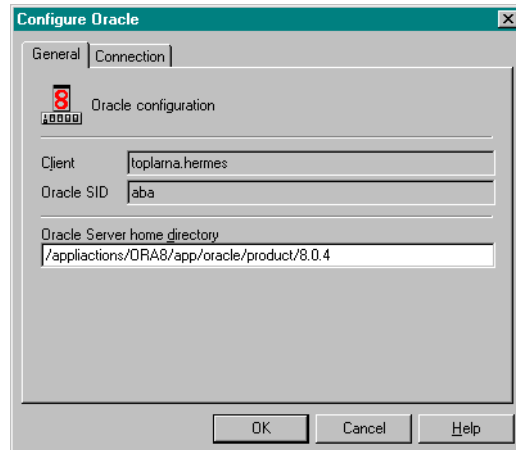
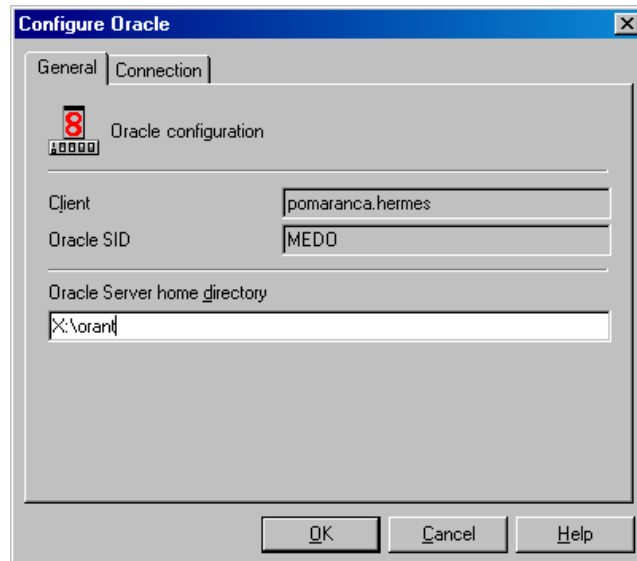


Figure 1-6 **Configuring Oracle on Windows - General**



In the Connection property page, enter the login information for the Oracle Target Database. Note that the username entered must have the SYSDBA privilege granted.

To grant the sysdba privilege, start the svrmgr1 (for Oracle8/8i) or sqlplus (for Oracle9i/10g) prompt and enter the following:

```
connect <user>/<password>@<service>;  
grant sysdba to <user>;
```

Refer to the Oracle documentation for more information on user privileges in Oracle.



Figure 1-7 **Configure Oracle on UNIX - Connection**

The screenshot shows the 'Configure Oracle' dialog box with the 'Connection' tab selected. It contains two sections for login information. The first section, 'Oracle login information to target database', has fields for 'User name' (sys), 'Password' (masked with asterisks), and 'Service' (aba). The second section, 'Oracle login information to recovery database', has fields for 'User name' (rman), 'Password' (masked with asterisks), and 'Service' (TCLAT). Between these sections are two radio buttons: 'Use target database control file instead of recovery catalog' (unselected) and 'Use recovery catalog' (selected). At the bottom are 'OK', 'Cancel', and 'Help' buttons.

Figure 1-8 **Configuring Oracle on Windows - Connection**

The screenshot shows the 'Configure Oracle' dialog box with the 'Connection' tab selected. It contains two sections for login information. The first section, 'Oracle login information to target database', has fields for 'User name' (sys), 'Password' (masked with asterisks), and 'Service' (MEDO). The second section, 'Oracle login information to recovery database', has fields for 'User name' (rman), 'Password' (masked with asterisks), and 'Service' (TCLAT). Between these sections are two radio buttons: 'Use target database controlfile instead of recovery catalog' (unselected) and 'Use recovery catalog' (selected). At the bottom are 'OK', 'Cancel', and 'Help' buttons.

What Happens?

The following takes place after clicking OK in the Configure Oracle window.

1. The `util_oracle8.exe` executable is started on the Oracle Server system, and saves the configuration parameters in the Data Protector Oracle configuration files. See “Data Protector Oracle Configuration Files” on page 14.
2. If the Use recovery catalog button was selected, the `util_oracle8.exe` executable starts the Oracle RMAN command, which registers the Oracle Target Database in the recovery catalog.


Information about the Oracle Target Database’s structure is transferred to the recovery catalog from the Oracle Target Database’s control files.

Checking the Configuration

After clicking OK in the Configure Oracle window, the next page in the wizard allows you to check the configuration.

Data Protector will verify the configuration by trying to connect to the Oracle Server system using the information specified and saved during the configuration procedure.

Figure 1-9 **Checking the Configuration**



Proceed as follows to check the configuration:

1. Right-click the Oracle Server system.
2. Click Check Configuration.

If the configuration is successful, you will receive a message confirming that the integration was properly configured. If it was not, you will receive a message explaining the reasons for the unsuccessful configuration.


IMPORTANT

On UNIX, when checking the Data Protector Oracle integration configuration using the Data Protector GUI, it is possible that although the GUI check returns a successful result, you may still receive the 12:8300 error when trying to start a backup session. Such a backup session will not start. For more information, see “Troubleshooting” on page 113.

The configuration can also be checked if you have already created and saved a backup specification for backing up a particular Oracle Server:

1. In the Data Protector Manager, switch to the Backup context. In the Scoping Pane, expand Backup, Backup Specification, then Oracle Server.
2. In the Results Area, right-click the backup specification.
3. In the Source property page, right-click the name of the client system, then click Check Configuration.

Configuring the Oracle Server Using the CLI

 the user root or as the Oracle user that is identified as described in “Configuring an Oracle User in Data Protector on UNIX Systems” on page 29. Note that the identified Oracle user and the user root must also be added to Data Protector admin or operator group.

Cluster-Aware Clients

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 Windows: set OB2BARHOSTNAME=<virtual_hostname>
- On UNIX: export OB2BARHOSTNAME=<virtual_hostname>

**Configuration
Command Syntax**

Run the following command to configure an Oracle Server system using the Data Protector CLI:

- On HP-UX systems without RAC or on Solaris systems:

```
/opt/omni/sbin/util_oracle8.exe -CONFIG <ORACLE_SID> \  
<ORACLE_HOME> <Target_Database_Login> [<Recovery_Catalog \  
_Login>]
```

- On HP-UX systems *with* RAC, execute the command on every RAC node you want to backup or use for restore:

— Using RAC node names:

```
/opt/omni/sbin/util_oracle8.exe -CONFIG \  
<ORACLE_GLOBAL_INSTANCE_SID> <ORACLE_HOME> \  
<Target_Database_Login> [<Recovery_Catalog_Login>]
```

— Using the virtual hostname:

```
/opt/omni/sbin/util_oracle8.exe -CONFIG <ORACLE_SID> \  
<ORACLE_HOME> <Target_Database_Login> \  
[<Recovery_Catalog_Login>]
```

- On other UNIX systems *without* RAC:

```
/usr/omni/bin/util_oracle8.exe -CONFIG <ORACLE_SID> \  
<ORACLE_HOME> <Target_Database_Login> \  
[<Recovery_Catalog_Login>]
```

- On other UNIX systems *with* RAC:

```
/usr/omni/bin/util_oracle8.exe -CONFIG \  
<ORACLE_GLOBAL_INSTANCE_SID> <ORACLE_HOME> \  
<Target_Database_Login> [<Recovery_Catalog_Login>]
```

The last connection string is required only if you are using the recovery catalog, otherwise the database control files are used instead.

The variables are defined as follows:

<ORACLE_SID>	The name of the Oracle Server instance.
<ORACLE_GLOBAL_INSTANCE_SID>	The name of the RAC global instance.
<ORACLE_HOME>	The home directory of the Oracle Server instance.

<code><Target_Database_Login></code>	The format of the login information is <code><user_name>/<password>@<service></code> . It is used to connect to the Target Database via the Recovery Manager.
<code><Recovery_Catalog_Login></code>	The format of the login information is <code><user_name>/<password>@<service></code> . It is used to connect to the recovery catalog database via the Recovery Manager. This parameter is optional and is to be used only if you are using the recovery catalog database.

Example Configurations

In the examples below, the Oracle service name and the ORACLE_SID of the Target database is aba, and the user name sys is identified by the password manager.

The recovery catalog owner is rman, the corresponding Oracle service name and the ORACLE_SID is TCAT, and the user name rman is identified by the password rman.

HP-UX and Solaris Example

```
/opt/omni/sbin/util_oracle8.exe -CONFIG aba \  
/applications/ORA8/app/oracle/product/8.0.4 \  
sys/manager@aba rman/rman@TCAT
```

Other UNIX Example

```
/usr/omni/bin/util_oracle8.exe -CONFIG aba \  
/applications/ORA8/app/oracle/product/8.0.4 \  
sys/manager@aba rman/rman@TCAT
```

Windows Example

```
<Data_Protector_home>\bin\util_oracle8.exe -CONFIG aba  
x:\orant\ sys/manager@aba rman/rman@TCAT
```

What Happens?

The following happens after typing the command at the command line:

1. The `util_oracle8.exe` executable is started on the Oracle server system, which saves the configuration parameters in the Data Protector Oracle configuration files. See “Data Protector Oracle Configuration Files” on page 14.
2. If the `<Recovery_Catalog_Login>` option has been specified in the command line, the `util_oracle8.exe` executable starts the Oracle RMAN command line interface, which registers the Oracle Target Database to the recovery catalog.

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

If you need to export some variables before starting the Oracle8 Server Manager or Oracle9i/10g SQL Plus, TNS listener, or Recovery Manager, these variables must be defined in the `Environment` section of the Data Protector Oracle global configuration file. See "Data Protector Oracle Configuration Files" on page 14 for information on how to define such variables.

Checking the Configuration

On UNIX, 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 an Oracle User in Data Protector on UNIX Systems" on page 29. The identified Oracle user and the user `root` must also be added to Data Protector `admin` or `operator` group.

Run the following command on the Oracle Server system:

- On HP-UX and Solaris: `/opt/omni/sbin/util_oracle8.exe -CHKCONF <ORACLE_SID>`
- On other UNIX: `/usr/omni/bin/util_oracle8.exe -CHKCONF <ORACLE_SID>`
- On Windows: `<Data_Protector_home>\bin\util_oracle8.exe -CHKCONF <ORACLE_SID>`.

Data Protector attempts to connect to the Oracle 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>`.

On UNIX, to obtain an error description, run the following command on the Oracle Server system:

- On HP-UX and Solaris: `/opt/omni/sbin/omnigetmsg 12 <Error_number>`
- On other UNIX: `/usr/omni/bin/omnigetmsg 12 <Error_number>`

`/opt/omni/sbin/omnigetmsg 12`

Configuring an Oracle Backup

To configure an Oracle backup, perform the following steps:

1. Configure the devices you plan to use for a backup. Refer to 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 Oracle backup specification. See “Creating a Data Protector Oracle Backup Specification” on page 44.

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, refer to “Creating a Data Protector Oracle Backup Specification” on page 44 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 Data Protector Manager, switch to the Backup context.
2. In the Scoping Pane, expand Backup, and then Backup Specifications. Right-click Oracle Server and then click Add Backup. The Create New Backup dialog box is displayed.
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 Results Area, provide the following information:

- In the Client drop-down list, select the Oracle Server that you want to back up. If the application is cluster-aware, select the virtual server of the Oracle resource group.
- In the Application database drop-down list, select the name of the Oracle application instance that you want to back up.
- On UNIX, enter also the Oracle user name and its group name as described in “Configuring an Oracle User in Data Protector on UNIX Systems” on page 29.

Click Next.

5. If the Oracle Server is already configured, the Source property page is displayed. Otherwise, you are prompted to configure it. See “Configuring the Oracle Server” on page 30 for details.
6. In the Source property page, select the database objects you want to back up.

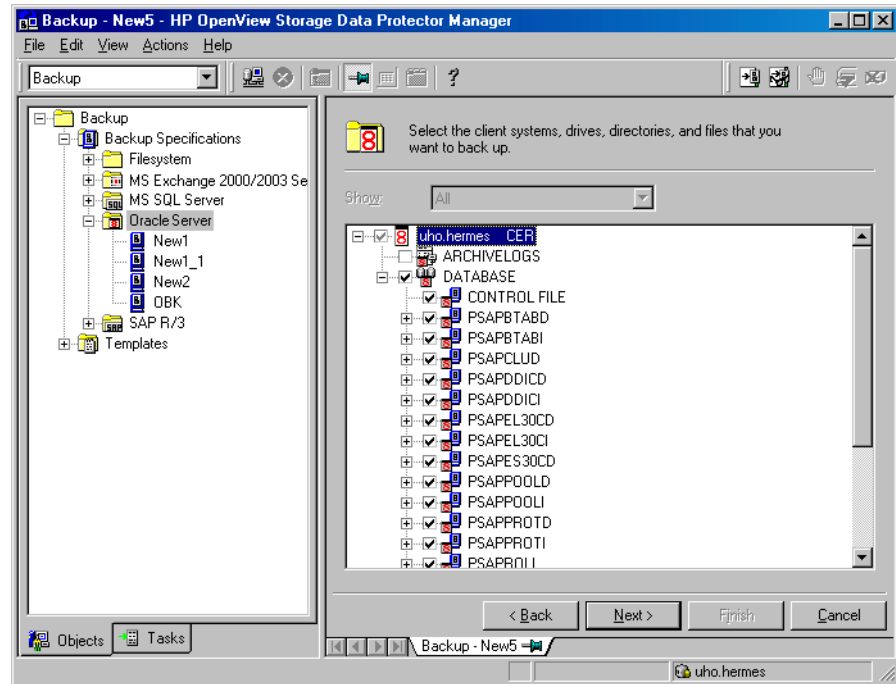
For example, a single tablespace can be separately selected for the backup, but for a complete and consistent online backup of the Oracle Target Database, the ARCHIVELOGS must be selected.

NOTE

If your Oracle Target Database uses a recovery catalog, this is backed up by default after each Oracle Target Database backup. This is not the case if control files are used instead of the recovery catalog. The control files must be backed up separately.

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, refer to the *HP OpenView Storage Data Protector Administrator's Guide*.

Click *Next* to proceed.

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 Options (Oracle specific backup options), see “Oracle Specific Backup Options” on page 48 or online Help.

Click Next.

9. Optionally, schedule the backup. For more details, see “Scheduling a Backup” on page 62.


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 channel names. 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 1-11 **Saving the Backup Specification**



Click OK.

To start the backup, see “Backing Up an Oracle Database” on page 58.

11. On UNIX, after the backup specification is saved, verify that the owner of the backup specification is the specified Oracle user. See “Configuring an Oracle User in Data Protector on UNIX Systems” on page 29 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.

- On UNIX:
`/etc/opt/omni/server/barlists/oracle8/<Backup_Spec_Name>`
- On Windows:
`<Data_Protector_home>\Config\server\Barlists\Oracle8\
<Backup_Specification_Name>`

13. It is recommended to test the backup specification. See “Testing the Integration” on page 55 for details.

Oracle Specific Backup Options

The window can be accessed from the Options property page of an Oracle backup specification by clicking the Advanced tab.

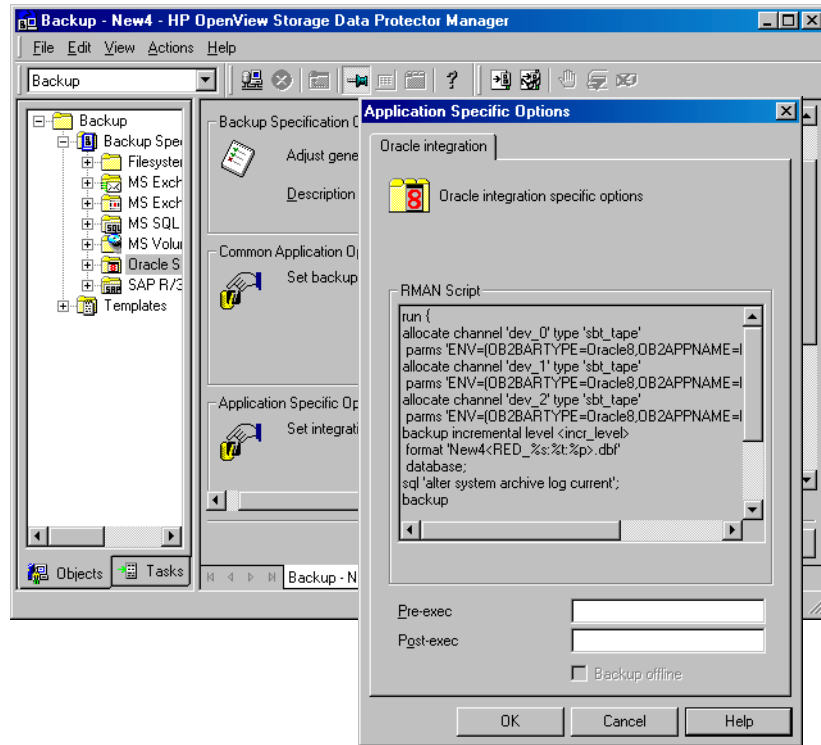
- The RMAN script section is created by Data Protector during the creation of a backup specification and reflects the backup specification’s selections and settings. For information on how to edit the Oracle RMAN script, see “Editing the Oracle RMAN Script” on page 50.
- ” on page 49.

TIP

If the Application Specific Options window is not displayed when clicking the Advanced button, this means that you have deselected all backup objects that were present in the backup specification, either by using the Data Protector GUI or by editing the RMAN script section.

Using the Data Protector GUI, select backup objects to display the Application Specific Options window when clicking the Advanced button.

Figure 1-12 Accessing the Application Specific Options Window



Specifying the Pre- and Post-exec commands

To specify Pre- or Post-exec commands, enter the command in the box next to the following two options:

- **Pre-exec**

Specifies a command/script that is started on the Oracle server system before backup. The command is started by the `ob2rman.exe` scrip. The full pathname of the command must be provided in the backup specification.

On UNIX, it runs under the account of the backup specification owner. See “Configuring an Oracle User in Data Protector on UNIX Systems” on page 29 for details about this user.

- **Post-exec**

Specifies a command/script that is started on the Oracle server system after backup. The command is started by the `ob2rman.exe` script. The full pathname of the command must be provided in the backup specification.

On UNIX, it runs under the account of the backup specification owner. See “Configuring an Oracle User in Data Protector on UNIX Systems” on page 29 for details about this user.

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 file but is kept in memory until the backup specification is either saved or manually edited by clicking the `Edit` button.

IMPORTANT



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 Oracle9i).
- 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 data files (the first backup command)
 - Archive logs (the second backup command)
 - Control files (the third or last backup command)

The RMAN scripts with all combinations of the three 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.

IMPORTANT

When editing the RMAN script section of the Data Protector Oracle backup specification, make sure that all manually-entered RMAN commands are *backup-related*. The RMAN script section of the Data Protector Oracle backup specification is not meant for any other Oracle-related tasks, such as maintenance, configuration, registration, etc.

To edit an Oracle RMAN script, click **Edit** in the Application Specific Options window (see Figure 1-12 on page 49), edit the script, and then click **Save** to save the changes to the script.

Refer to the *Oracle8i 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 (see also “Example of the RMAN Script” on page 54):

- **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 16 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=<ORACLE_SID>,  
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 or the datafile**. The backup statement consists of the following:

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

```
format  
'<Backup_Specification_Name><<ORACLE_SID>_%s:%t:%p>.db  
f'
```

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 <ORACLE_SID>, 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><<ORACLE_SID>_%s:%t:%p>.db  
f'
```

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 <ORACLE_SID>.

- The RMAN `archive log 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:

```
archive log 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><<ORACLE_SID>_%s:%t:%p>.db  
f'
```

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 <ORACLE_SID>, 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

Oracle8i/9i/10g Duplex Mode

Oracle8i/9i/10g 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, [REDACTED]

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.

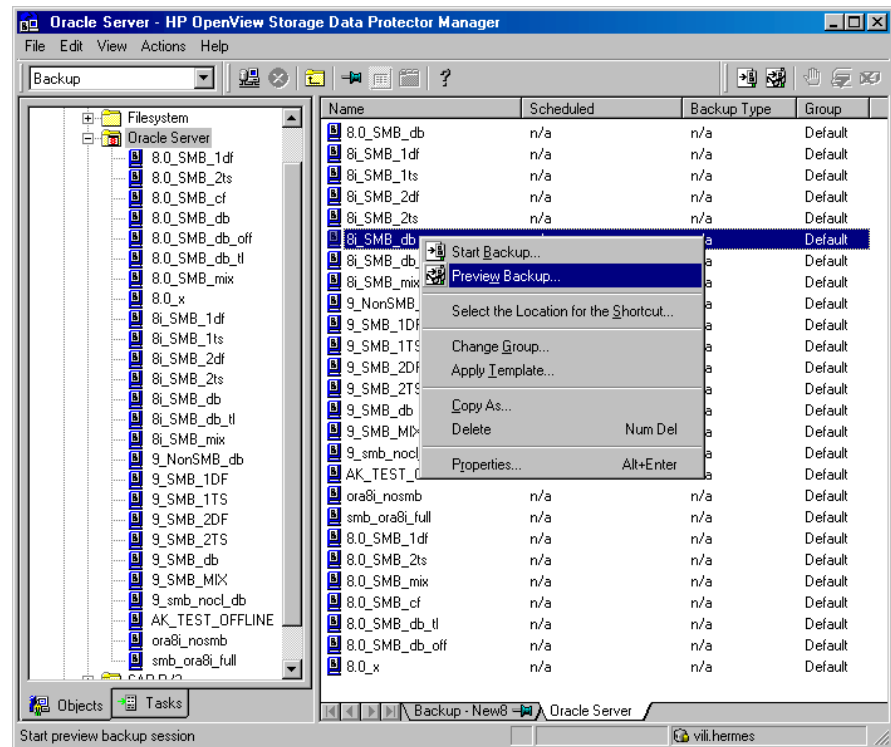
Testing Using the Data Protector GUI

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.

You have to run the `omnib` command with the `-test_bar` option. Run the following command:

- 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 Windows: `<Data_Protector_home>\bin\omnib
-oracle8_list <backup_specification_name> -test_bar`

What Happens?

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

1. The `ob2rman.exe` script is started, which then starts the Oracle RMAN backup command.

The Oracle Target Database is backed up to the following file on the Oracle server system without using Data Protector for data transfer:

- On UNIX: `/dev/null/`
- On Windows:
`<Data_Protector_home>\tmp\<Backup_Specification_Name>
_TEST_FILE`

This is how the Oracle side of the integration is checked.

2. The Data Protector `testbar2` command is started by the `ob2rman.exe` script, which checks:
 - The communication within Data Protector
 - The syntax of the Oracle backup specification
 - If the devices are correctly specified
 - If the required media reside in the devices

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 opened 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




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. 
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 opened database is inconsistent, because portions of the database are being modified and written to disk while the backup is progressing.

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. Restore the Archived Redo Logs to disk.
3. 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.

In order to be able to restore to a consistent state,

- 1.

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 the Oracle8 Server Manager or Oracle9i/10g SQL Plus and issuing the following command:

If the Oracle Target Database is not operating in the ARCHIVELOG mode, shut it down first and mount it again. Issue the following command at the Oracle8 Server Manager or Oracle9i/10g SQL Plus prompt:

```
alter database archivelog;  
  
archive log start;  
  
alter database open;
```

IMPORTANT

On Windows, do not use Oracle reserved names, such as ALLOCATE, BACKUP, CATALOG, FORMAT, INCREMENTAL, RUN, TABLESPACE and others for tablespace names. Refer to the Oracle documentation for a complete list. Oracle reserved names used as names for tablespaces may cause the backup to fail.

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 62.
- 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 63.
- Start a backup on the Oracle server using either Oracle Recovery Manager or Oracle Enterprise Manager. See “Starting Oracle Backup Using the Oracle Recovery Manager (RMAN)” on page 65.

Backup Procedure

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

1. Data Protector executes `ob2rman.exe` on the client. This command starts the Recovery Manager (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 Database Library 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.

Automatic

IMPORTANT

When backing up or restoring the catalog database, the value of the NLS_LANG variable, which is set for the target database, is used as the default NLS_LANG value for the catalog database.

If you want to use a different NLS_LANG value for importing or exporting the catalog database, you can specify this in the Oracle instance configuration file in the section Environment. Refer to the section “Setting, Retrieving, and Listing Data Protector Oracle Configuration Files Parameters Using the CLI” on page 16.

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. Refer to the *Oracle8i Recovery Manager User's Guide and References* for detailed information about deleting data from the recovery catalog.

NOTE

You can obtain the primary keys of the records to be deleted by issuing the `list` command, as shown in the following example:

If the (data) file is stored on a disk, issue the following command:

```
list backup of tablespace temp;
```

Primary keys in the `temp` tablespace are listed. You need these keys to delete the records, as shown in the following example:

```
allocate channel for delete type 'sbt_tape';
```

```
change backupset <primary_key> delete;  
release channel;
```

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 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:

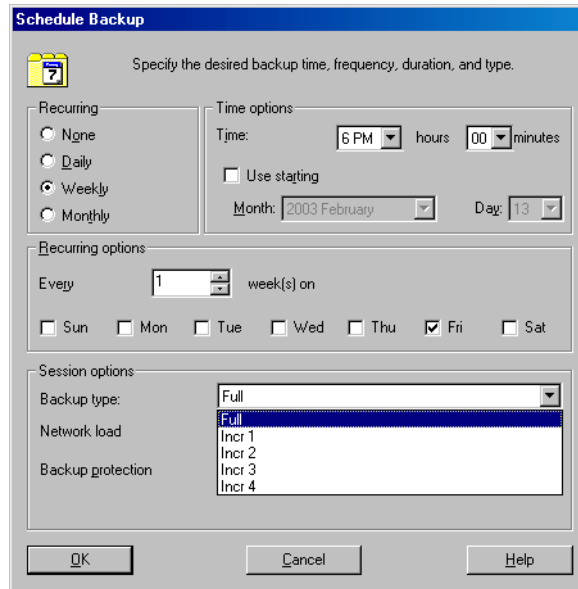
- Weekly full backup
- Daily incremental backup
- Archived Log backups as needed

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 63. Refer to 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.

Starting a Backup Using the GUI

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

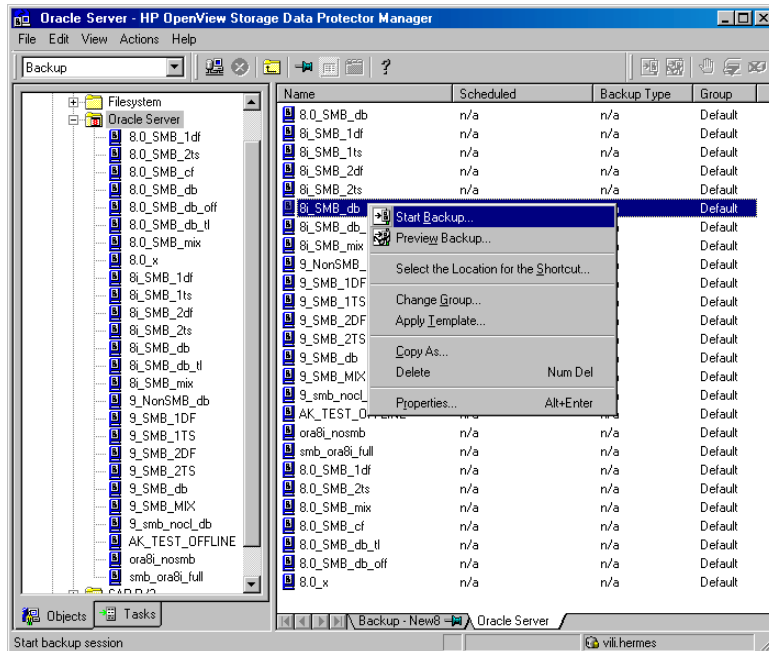
1. In the HP OpenView Storage Data Protector Manager, switch to the Backup context.
2. In the Scoping Pane, expand Backup, Backup Specifications, and then Oracle Server.
3. Right-click the backup specification and select Start Backup.

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 63. Refer to the RMAN documentation for details on incremental backup levels.

Click OK.

Figure 1-15 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>`

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 Oracle Server system and run the following command:

```
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}
```

Refer to 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 the Oracle Recovery Manager (RMAN)

The Oracle Recovery Manager (RMAN) utility is an Oracle CLI that allows you to perform a backup, restore, or recovery of Oracle database objects. To start an Oracle backup using RMAN, an Oracle backup specification must be created.

See “Configuring an Oracle Backup” on page 43 for information on how to create an Oracle backup specification. To start an Oracle backup using RMAN, follow these steps:

1. Connect to the Oracle Target Database specified in the backup specification.

To connect to an Oracle Target Database that uses the recovery catalog, run the following command:

- On UNIX: `<ORACLE_HOME>/bin/rman target <Target_Database_Login> rcvcat <Recovery_Catalog_Login>`

- On Windows:

For Oracle8:

```
<ORACLE_HOME>\bin\rman80 target  
<Target_Database_Login> rcvcat  
<Recovery_Catalog_Login>
```

For Oracle8i/9i/10g:

```
<ORACLE_HOME>\bin\rman target  
<Target_Database_Login> rcvcat  
<Recovery_Catalog_Login>
```

To connect to an Oracle Target Database without using the recovery catalog, run the following command:

- On UNIX: `<ORACLE_HOME>/bin/rman target
<Target_Database_Login> nocatalog`

- On Windows:

For Oracle8:

```
<ORACLE_HOME>\bin\rman80 target  
<Target_Database_Login> nocatalog
```

For Oracle8i/9i/10g:

```
<ORACLE_HOME>\bin\rman target <Target_Database_Login>  
nocatalog
```

See the “Glossary” on page G- 1 for details on the login information syntax.

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 Oracle is linked to Data Protector, 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 systems, a maximum of 16 channels can be allocated.

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

3. You must specify the `parms` operand in the following form:

```
parms 'ENV(OB2BARTYPE=Oracle8,  
OB2APPNAME=<ORACLE_SID>,OB2BARLIST=<backup_  
specification_name>)' ;
```

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

4. Optionally, you may specify `format` as follows:

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

Note that `%s:%t:%p` and the Oracle SID 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'
```

Refer to the *Oracle8i 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, you may 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.

In order to run a backup using RMAN, start the Recovery Manager by running the following command from the `<ORACLE_HOME>` directory (if you are using recovery catalog):

- On UNIX: `bin/rman target <Target_Database_Login> rcvcat <Recovery_Catalog_Login>`

- On Windows:

For Oracle8:

```
bin\rman80 target <Target _Database_Login> rcvcat  
<Recovery_Catalog_Login>
```

For Oracle8i/9i/10g:

```
bin\rman target <Target _Database_Login> rcvcat  
<Recovery_Catalog_Login>
```

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 `credo`, using a backup specification named `alma`, enter the following command sequence:

```
run {  
  allocate channel 'dev_0' type 'sbt_tape'  
  parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=credo,OB2BARLIST=alma)';  
  backup  
  incremental level 0  
  format 'alma<credo_%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=credo,OB2BARLIST=alma)';  
  allocate channel 'dev_1' type 'sbt_tape'  
  parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=credo,OB2BARLIST=alma)';  
  allocate channel 'dev_2' type 'sbt_tape'  
  parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=credo,OB2BARLIST=alma)';
```



```
backup  
incremental level 0  
format 'alma<credo_%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 alma, the RMAN script should look like this:

```
run {  
  allocate channel 'dev_0' type 'sbt_tape'  
  parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=credo,OB2BARLIST=alma)';  
  allocate channel 'dev_1' type 'sbt_tape'  
  parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=credo,OB2BARLIST=alma)';  
  allocate channel 'dev_2' type 'sbt_tape'  
  parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=credo,OB2BARLIST=alma)';  
  backup  
  incremental level 0  
  format 'alma<credo_%s:%t>.dbf'  
  tablespace SYSTEM, RONA  
  sql 'alter system archive log current'  
  format 'alma<credo_%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 alma is complete, run the following script:

```
run {  
  allocate channel 'dev_0' type 'sbt_tape'  
  parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=credo,OB2BARLIST=alma)';  
  allocate channel 'dev_1' type 'sbt_tape'  
  parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=credo,OB2BARLIST=alma)';  
  allocate channel 'dev_2' type 'sbt_tape'  
  parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=credo,OB2BARLIST=alma)';  
  backup
```

```
incremental level 0
(archivelog low logseq 5 high logseq 105 thread 1 all delete input
format 'alma<credo_%s:%t:%p>.dbf');
}
```

If the backup fails, the logs are not deleted.

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=credo,OB2BARLIST=alma)';
allocate channel 'dev_1' type 'sbt_tape'
'parms'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=credo,OB2BARLIST=alma)';
allocate channel 'dev_2' type 'sbt_tape'
parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=credo,OB2BARLIST=alma)';
backup
incremental level 0
(tablespace COSTS current controlfile
format 'alma<credo_%s:%t>.dbf');
;}
```

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 alma 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 {
allocate channel 'dev_0' type 'sbt_tape'
parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=credo,OB2BARLIST=alma)';
allocate channel 'dev_1' type 'sbt_tape'
parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=credo,OB2BARLIST=alma)';
```

```
allocate channel 'dev_2' type 'sbt_tape'
parms'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=credo,OB2BARLIST=alma)';
backup
incremental level 0
format 'alma<credo_%s:%t>.dbf'
database;
set maxcorrupt for datafile
'/oracle/data1.dbs' to 10;}
```

On Windows

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

Restoring an Oracle Database

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

Restore Methods

- Data Protector GUI. See “Restoring Oracle Using the Data Protector GUI” on page 74.
- Oracle Recovery Manager (RMAN). See “Restoring Oracle Using RMAN” on page 91.

Restorable Items

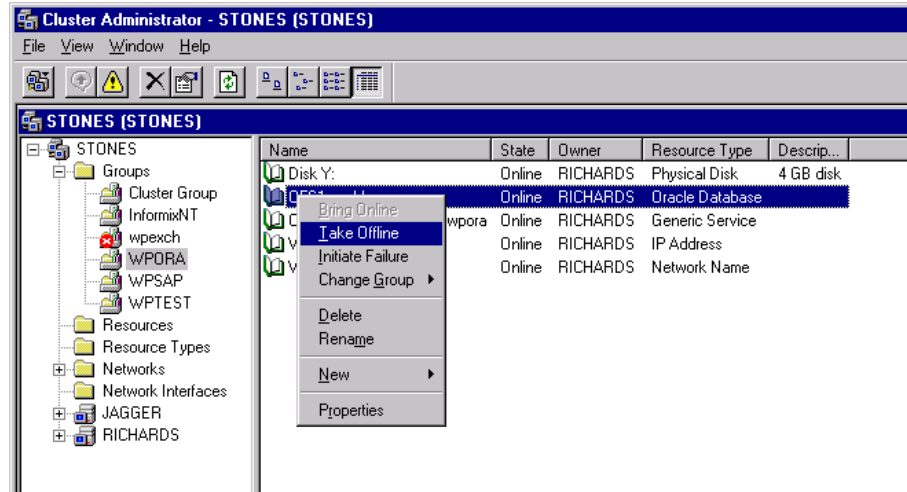
It is possible to restore the following database objects using both the Data Protector GUI or RMAN.

- Control files
- Datafiles
- Tablespaces
- Databases

MS Cluster Server Clients


Before you start restoring a cluster-aware Oracle server, [REDACTED] 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 <Oracle_SID>.world resource, which is an Oracle Database resource.

Figure 1-17 Checking Properties



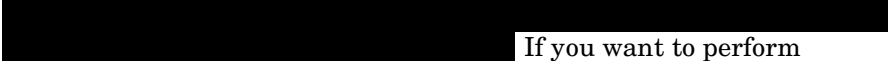
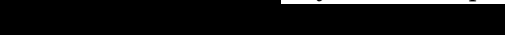
MC/ServiceGuard Clients

When restoring the database from a backup performed on a virtual host, you should [REDACTED] 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;  
}
```

Restoring Oracle Using the Data Protector GUI

 If you want to perform additional actions,  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:
 - Oracle8/8i:

Refer to “The Recovery Catalog was lost and the control file cannot be restored from Data Protector managed backup.” on page 138.
 - Oracle9i/10g:
 1. Restore the control file from automatic backup.

If no automatic backup of the control file is available, refer to “The Recovery Catalog was lost and the control file cannot be restored from Data Protector managed backup.” on page 138
 2. Restore the database or data items.

Changing The Database State

Before you restore any database item you need to ensure that the database is in the correct state. The following table tells you which state the database needs to be in to restore a particular type of database item.

Table 1-1

Required Database States

Item to restore	Database state
Control file	NoMount (started)
All other items	Mount

To put the database into the correct state carry out the following procedure:

Open a command window and enter the following:

```
sqlplus /nolog
```

In the SQL> prompt, enter:

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

```
SQL>shutdown immediate;
```


If you are restoring a control file put the database into NoMount state.

```
SQL>startup nomount;
```

If you are restoring any other database item put the database in Mount state.

```
SQL>startup mount;
```

Restoring the Recovery Catalog Database



Carry out the following procedure to restore the recovery catalog database:

NOTE

If the recovery catalog uses a character set that is not the same as the character set used by the Oracle database that was backed up, the `NLS_LANG` variable must be set in the Oracle instance configuration file. See “Setting, Retrieving, and Listing Data Protector Oracle Configuration Files Parameters Using the CLI” on page 16 for more information on how to do that.

1. Ensure the recovery catalog database exists and is empty.

```
# ./util_cmd -getconf Oracle8 <SID>
```

If the recovery catalog database was selected as one of the Data Protector backup options when the original database backup was configured, this command returns the following output:

```
ORACLE_HOME=/app/oracle9i/product/9.2.0.1.0';
TGTLLogin='EIBBKIBBEIBBQDBBOHBBCHBBPHBBBIBBCHBBEIBBB
FBBFGBBFFBBDFFBB';
RCVLogin='DIBBOHBBCHBBPHBBQDBBDIBBOHBBCHBBPH
FBBFGBBFFBBDFFBB';
ORACLE_VERSION='9.2.0';
Configuration read/write operation successful.
```

2. [REDACTED] using the Data Protector GUI.
3. Ensure that the recovery catalog database is in the [REDACTED] state. Run the following command:

```
sqlplus /nolog
```


In the SQL> prompt, enter

```
SQL>connect <user>/<password>@<CATALOG_SERVICE_NAME> as  
sysdba;
```

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

If the database is not open, perform one of the following steps:

- If the database is shut down, start it:

```
SQL>startup;
```

- If the database is in nomount state, mount the database and then open it:

```
SQL>alter database mount;
```

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

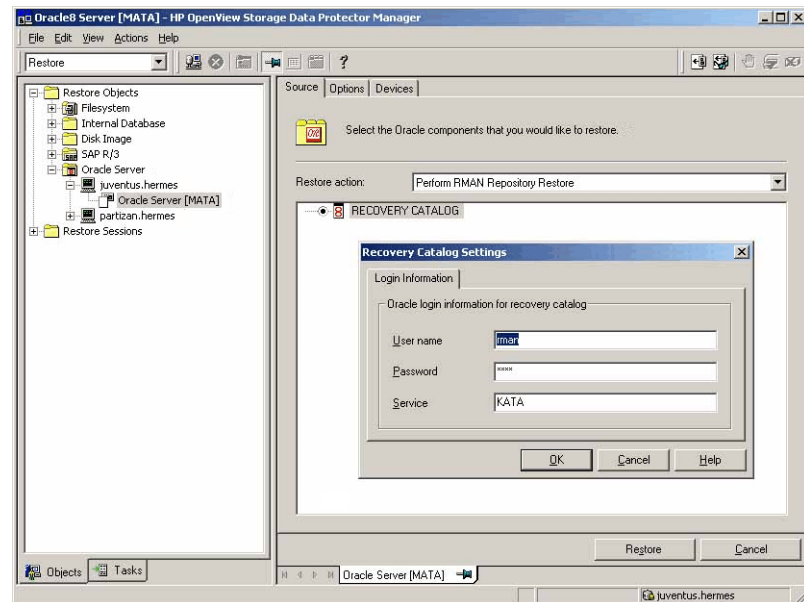
- If the database is mounted, open it:

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

4. In the Data Protector GUI, switch to the Restore context and select the Oracle server and Oracle instance for which you need to restore the recovery catalog database.
5. In the Results Area, select Recovery Catalog. If you want to change or enter the recovery catalog login strings, right-click on the Recovery Catalog and select Change Settings.

Figure 1-18

Recovery Catalog and Change Recovery Catalog Settings Dialogs



6. Enter the recovery catalog database login strings in the Recovery Catalog Settings dialog.
7. Select the Options pane on the Data Protector GUI. Select the Session ID from the Session ID drop-down list. For further information, see “Restore and Recovery Options” on page 86.
8. Enter the user name and password to the recovery catalog database in the User name and User group fields on the Options pane.
9. Click the Restore button. At this point the recovery catalog database is restored. You can now proceed to restore your 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 can restore any other part of the database.

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.exe` at the end of a backup session. The recovery catalog is *not* required for this option.

The control files (`ctrl<ORACLE_SID>.dbf`) are restored to `/var/opt/omni/tmp` (HP-UX and Solaris systems), `/usr/opt/omni/tmp` (other UNIX systems), or `<Oracle_home>\tmp` (Windows systems). After the restore, run the following script:

```
run {
  allocate channel 'dev0' type disk;
  replicate 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 Oracle8/8i.

The control file was automatically backed up by RMAN and the recovery catalog is *not* available. The database must be in the `nomount` state.

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. Refer to the Oracle9i/10g documentation on how to set up RMAN AUTOBACKUP.

- Restoring from RMAN backup set (`CONTROLFILE FROM RMAN BACKUPSET`)

This type of restore is available if you selected the control file in the backup specification.

The recovery catalog *is* required and the database must be in `nomount` state.

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

The procedure for restoring the control file is as follows:

1. Open an `sqlplus` (for Oracle9i/10g) or `svrmgrl` (for Oracle8/8i) window and put the database in the `nomount` state.
2. In the Data Protector GUI switch to the `Restore` context. Expand `Restore` and select the Oracle Server and the Oracle instance for which you need to restore the control file.

Restoring to Another Client

Optionally, to restore the control file to another client, click on the `Options` tab and select the name of the client from the `Restore to client` drop-down list. If the target instance requires different connection parameters, click on `Settings` and enter the username, password and service in the dialog box.

3. In the `Restore Action` drop-down list, select `Perform RMAN Repository Restore`.

In the `Results` area select the control file.

Set the required options. Refer to “`Restore and Recovery Options`” on page 86 for a more detailed description of available options.

4. Click `Restore`.

You must now proceed to recover the Oracle database or items within the database.

Restoring Oracle Database Objects

Before you restore Oracle database objects you need to ensure that you have an up-to-date version of the recovery catalog database and the control file. These contain the database structure information. If you do not already have up-to-date versions of these files, you need to restore them. See “`Restoring the Recovery Catalog Database`” on page 75 and “`Restoring the Control File`” on page 78 for more information.

To restore Oracle database items, proceed as follows:

1. First ensure that the database is in the correct state before you begin to restore database items. Details of this can be found in “`Changing The Database State`” on page 75.
2. In the `Context List`, select `Restore`.

In the Data Protector GUI switch to the `Restore` context. Expand `Restore` and select the Oracle Server and the Oracle instance for which you need to restore the control file.

Restoring to Another Client

Optionally, to restore the database objects to another client, click on the **Options** tab and select the name of the client from the **Restore to client** drop-down list. If the target instance requires different connection parameters, click on **Settings** and enter the username, password and service in the dialog box.

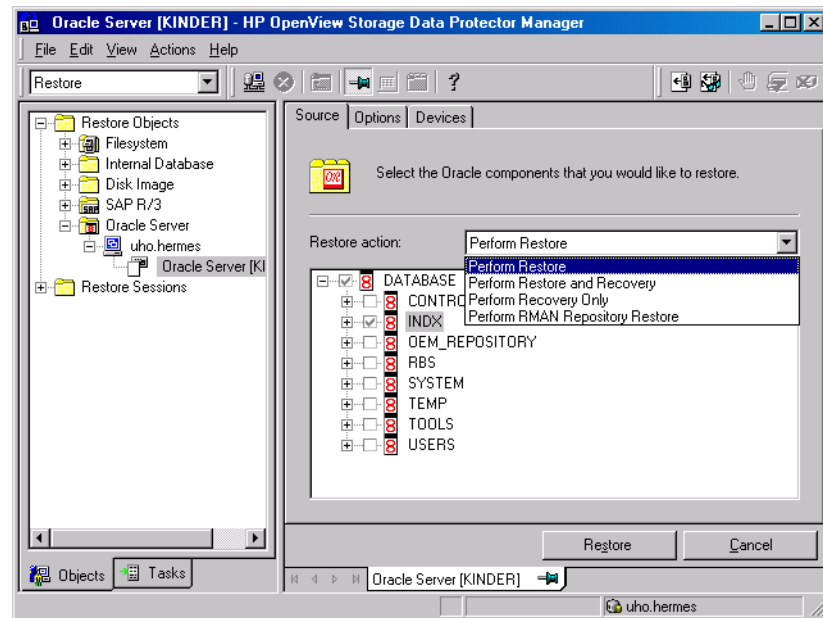
3. In the **Restore action** drop-down list, select the type of restore action you wish to perform. For more details on these options, see “Restore and Recovery Options” on page 86.

IMPORTANT

If you do not select **Perform Restore & Recovery** or **Perform Recovery Only** in the **Restore action** drop-down list you will have to recover database items manually from the command line. For more information about using RMAN refer to “Restoring Oracle Using RMAN” on page 91.

Figure 1-19

Source Pane

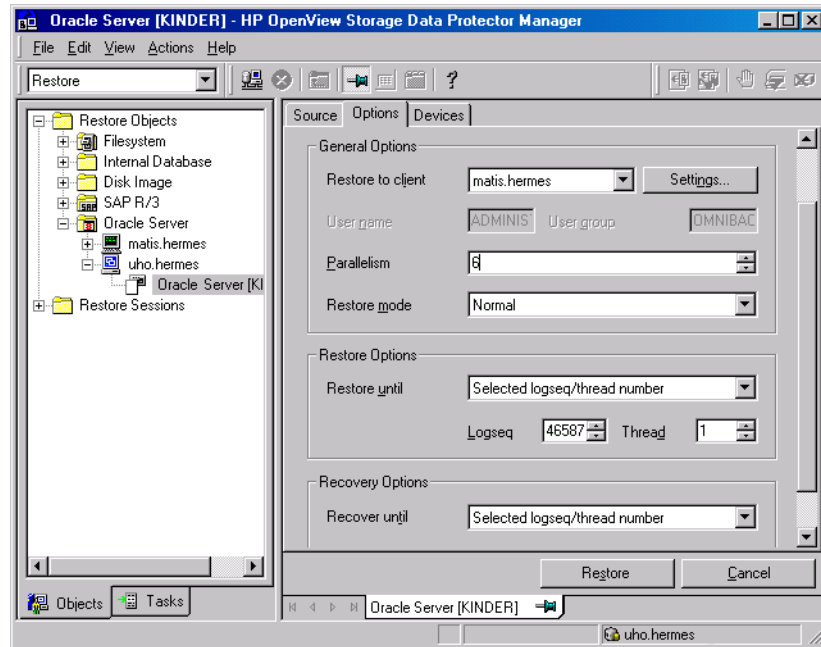


4. In the Results Area, click on the items you would like to restore. If you are restoring datafiles, it is possible to restore the files to a new location by right-clicking on the filename. A dialog will appear into which you can enter the new datafile location.

If you select to restore the datafile to another location or with a different name, the datafile will be restored to the selected location. If you want Oracle to use the datafile under the new name, you must issue a switch statement afterwards using Oracle tools. For more details refer to *"Recovery Manager User's Guide and Reference"*.

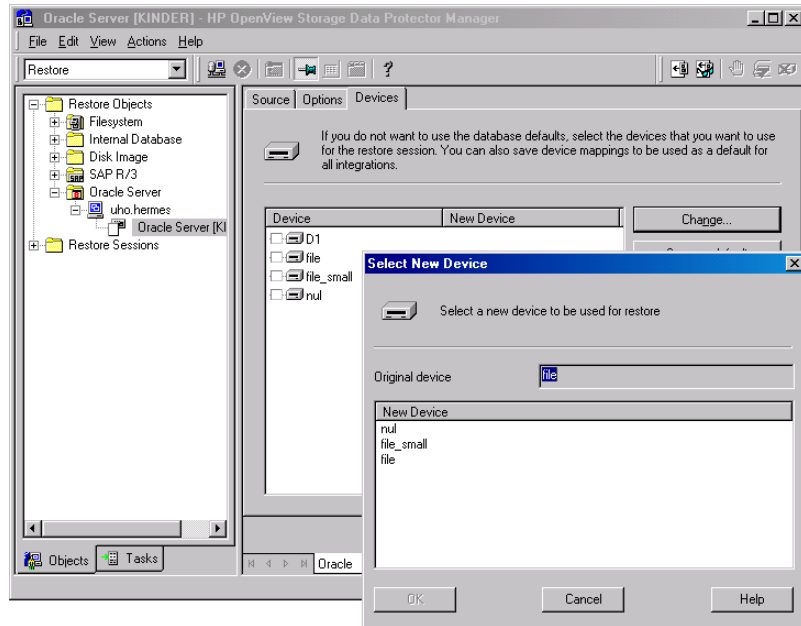
5. In the Options pane, enter the restore or recovery details. See "Restore and Recovery Options" on page 86 for more information about the fields in this pane.

Figure 1-20 Options Pane



6. In the **Devices** pane, select the devices to be used for the restore. You can select devices from which to restore the database other than those which were used for the original backup, although Data Protector defaults to the original device on which the backup was made. For more information on the **Devices** pane, press **F1**.

Figure 1-21 **Devices Pane**



7. If you would like to change the device from which an item is restored, select your desired device and click **Change**.
8. After selecting all the devices, click **Restore**. The restore procedure starts.

After Recovering the Database

Once you have restored the lost database objects you need to ensure that the database is in the correct state.

If you used one of the options on the **Source** pane containing the word “Recovery” then the database will automatically be put into Open state by Data Protector.

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

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

```
rman target <Target_Database_Login> rcvcat  
<Recovery_Catalog_Login>RMAN> RESET DATABASE;
```



```
RMAN> exit
```

If you did not choose to use Data Protector to Recover database items, you need to carry out the following procedure after the database has been restored:

Open a command line window and enter the following commands:

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

Restoring Tablespaces and Datafiles

You can restore tablespaces and datafiles as follows:

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;
```

Restore and Recovery Options

Source Pane

The following describes each of the options on the `Source` pane. This pane is used to define the combination of Restore and Recovery you would like to perform with 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.

The options below are available on the `Restore` action drop-down list.

Perform Restore

This option specifies that the user will restore the database objects using Data Protector and will then perform the recovery manually using RMAN. For information on how to recover database items using RMAN see “Restoring Oracle Using RMAN” on page 91.

Perform Restore and Recovery

This option means that the user will perform both the Restore and the Recovery from the GUI.

Perform Recovery Only

The option specifies that the user will perform the recovery only using the GUI.

This option is meant to be used after Instant Recovery has been successfully completed.

Perform RMAN Repository Restore

This option is used to restore the recovery catalog or the control file if the database objects are not available in the `Source` property page.

The `Source` window allows you to specify another location to which to restore an Oracle database file. This is done using the `Restore As` dialog which appears when you right-click the mouse on a database item listed in the `Results` Area.

Options Pane

The following describe each of the fields on the `Options` pane.

Restore to client

This option specifies the name of the Oracle server to which the user wants to restore the database item. This defaults to the original backup server.

NOTE

When restoring to another client, you should select the database SID on the `Source` pane and then select the client to which to restore the database on the `Options` pane. At this point Data Protector automatically restores the control file and then all the objects in the database.

User name (UNIX systems only)

Use this field to enter the Oracle User Name. The user needs to be a member of the 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 an Oracle User in Data Protector on UNIX Systems” on page 29 for more information on this user and on how to identify it.

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.

To optimize restore performance, specify the same number of data streams as were used during the backup. If, for example you set the backup concurrency to 3 then set the number of parallel data streams to 3

as well. Note that if a very high number of parallel data streams are specified this may result in a resource problem because too much memory is being used.

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 either a non-ZDB or ZDB using the backup set method was made of the database, such as ZDB using a version of Data Protector which is older than 5.0.

- Proxy copy

This option should be used when the original Oracle backup was made using the Oracle proxy-copy method, such as ZDB of Oracle8i/9i using Data Protector version 5.1.

This option is disabled when you are performing a restore after Instant Recovery.

Restore until

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

This option is disabled when you are performing a restore after Instant Recovery.

- Now

This is the default option. Data Protector restores the database to the most recent time possible.

- Selected time

Using this option you can specify an exact time to which Data Protector restores database objects.

- Selected logseq/thread number

A logseq number is a redo log sequence number. You can enter a particular redo log sequence and a thread number which will act as an upper limit of redo logs to restore.

- Selected SCN number

This option allows you to enter the SCN number to which to perform the restore.

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

This is the default option. Data Protector starts RMAN to restore the database to the most recent time possible by applying all archive redo logs.

- Selected time

Using this option you can specify an exact time to which the archive logs are applied.

- Selected logseq/thread number

A logseq number is a redo log sequence number. You can enter a particular redo log sequence and a thread number which will act as an upper limit of redo logs to recovery.

- Selected SCN number

This option allows you to enter the SCN number to which to 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 the following:

```
rman target <Target_Database_Login> rcvcat  
<Recovery_Catalog_Login>
```

```
RMAN> RESET DATABASE;
```

```
RMAN> exit
```

Open database after recovery

Opens the database after a recovery is performed.

Reset logs

Resets the archive logs after the database is opened.

The following are recommendations on when to reset the logs. *Always* reset the logs:

- After an incomplete recovery, that is if not all archive redo logs will be applied.
- If a backup of a control file is used in recovery.

Do not reset the logs:

- After a complete recovery where the backup of a control file is not used in recovery.
- If the archive logs are used for a standby database. If you must reset the archive logs, then you have to recreate the standby database.

If you select to 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.

Target DB login

This option lets you change the target database login information, i.e. the username and password of the user who has SYSDBA privileges and the service name to which Data Protector should connect.

Restoring Oracle Using RMAN

Data Protector acts as a media management utility for the Oracle system, therefore the Oracle Recovery Manager (RMAN) can be used for a restore.

Refer to the *Oracle8i Recovery Manager User's Guide and References* for detailed information on how to perform database, tablespace, control file, and datafile restore and recovery.

IMPORTANT

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. For additional information on how you can restore an Oracle database using the RMAN utility, refer to the Oracle documentation.

This section provides some examples of how you can restore an Oracle database. The following examples of restore are given:

- “Example of Full Database Restore” on page 94
- “Example of Point-in-Time Restore” on page 95
- “Example of Tablespace Restore” on page 97
- “Example of Datafile Restore” on page 99
- “Example of Archive Log Restore” on page 101

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 on the version of the Oracle database you are using. For detailed steps to perform the restore of control files, please refer to the *Recovery Manager User's Guide and Reference*.

Preparing the Oracle Database for Restore

The restore of an Oracle database can be performed when the database is in a mount (quiescent) 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 are using a 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 database. Refer to “Restoring an Oracle Database” on page 72 for details on how to restore the recovery catalog database.
- If you are using control files, they must be accessible in order to perform the restore. If the control files are not available, you will probably need to restore them. Refer to the *Recovery Manager User's Guide and Reference* for more details.

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

When you are sure that the recovery catalog database or control files are in place, start the recovery catalog database and the listener.

- Make sure that the following environment variables are set:

- ✓ ORACLE_BASE
- ✓ ORACLE_HOME
- ✓ ORACLE_TERM
- ✓ ORACLE_SID
- ✓ PATH
- ✓ NLS_LANG
- ✓ NLS_DATE_FORMAT

Example of Environment Variables on UNIX

```
ORACLE_BASE=/opt/oracle
ORACLE_HOME=/opt/oracle/product/8.1.6
ORACLE_TERM=hp
ORACLE_SID=PROD
PATH=$PATH:/opt/oracle/product/8.1.6/bin
NLS_LANG=american
NLS_DATE_FORMAT='Mon DD YYYY HH24:MI:SS'
```


Example of Environment Variables on Windows

```
ORACLE_BASE=<Oracle_home>
ORACLE_HOME=<Oracle_home>\product\8.1.6
ORACLE_TERM=hp
ORACLE_SID=PROD
PATH=$PATH:<Oracle_home>\product\8.1.6\bin
NLS_LANG=american
NLS_DATE_FORMAT='Mon DD YYYY HH24:MI:SS'
```

- Check that the `/etc/oratab` file has the following line:
 - On UNIX: `PROD:/opt/oracle/product/8.1.6:N`
 - On Windows: `PROD:<Oracle_home>\product\8.1.6:N`The last letter determines whether the database will automatically start upon bootup (Y) or not (N).

Connection Strings Used in the Examples

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

- Target connection string for target database:
`sys/manager@PROD`
where `sys` is the username, `manager` is the password and `PROD` is the name of the Oracle database.
- Recovery catalog connection string for recovery catalog database:
`rman/rman@CATAL`
where `rman` is the username and password and `CATAL` is the name of the Oracle database.

Example of Full Database Restore

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

1. Log in to the Oracle RMAN:

If you are using the recovery catalog database, run the following command:

- On UNIX: `<ORACLE_HOME>/bin/rman target sys/manager@PROD rcvcat rman/rman@CATAL`
- On Windows: `<ORACLE_HOME>\bin\rman target sys/manager@PROD rcvcat rman/rman@CATAL`

If you are using control files instead of the recovery catalog database, run the following command:

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

2. Start the full database restore:

```
run{
allocate channel 'dev1' type 'sbt_tape' parms
'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=<ORACLE_SID>)';
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 are using the recovery catalog database, run the following command:

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

If you are using control files, run the following command:

- On UNIX: `<ORACLE_HOME>/bin/rman target sys/manager@PROD
nocatalog cmdfile=/var/opt/omni/tmp/restore_datafile`
- On Windows: `<ORACLE_HOME>\bin\rman target
sys/manager@PROD nocatalog
cmdfile=<Data_Protector_home>\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, follow the steps below:

1. Log in to the Oracle RMAN:

If you are using the recovery catalog database, run the following command:

- On UNIX: `<ORACLE_HOME>/bin/rman target sys/manager@PROD
rcvcat rman/rman@CATAL`
- On Windows: `<ORACLE_HOME>\bin\rman target
sys/manager@PROD rcvcat rman/rman@CATAL`

If you are using control files, run the following command:

- On UNIX: `<ORACLE_HOME>/bin/rman target sys/manager@PROD
nocatalog`
- On Windows: `<ORACLE_HOME>\bin\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=<ORACLE_SID>)';

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. Follow the steps below:

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 are using the recovery catalog database, run the following command:

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

If you are using control files, run the following command:

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

Example of Tablespace Restore

If a table is missing or corrupted, you need to perform a restore 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. Follow the steps below:

1. Log in to the Oracle RMAN:

If you are using the recovery catalog database, run the following command:

- On UNIX: `<ORACLE_HOME>/bin/rman target sys/manager@PROD rcvcat rman/rman@CATAL`
- On Windows: `<ORACLE_HOME>\bin\rman target sys/manager@PROD rcvcat rman/rman@CATAL`

If you are using control files, run the following command:

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

2. Start the tablespace restore.

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

```
run{
allocate channel <dev1> type 'sbt_tape' parms
'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=<ORACLE_SID>)' ;
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 the tablespace should have the following format:

```
run{
    allocate channel <dev1> type 'sbt_tape' parms
    'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=<ORACLE_SID>)' ;
    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. The procedure in such cases is as follows:

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 are using the recovery catalog database, run the following command:

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

If you are using control files, run the following command:

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

Example of Datafile Restore

To restore a datafile, you may take only a part of the database offline. To perform a datafile restore, follow the steps below:

1. Log in to the Oracle RMAN:

If you are using the recovery catalog database, run the following command:

- On UNIX: `<ORACLE_HOME>/bin/rman target sys/manager@PROD rcvcat rman/rman@CATAL`
- On Windows: `<ORACLE_HOME>\bin\rman target sys/manager@PROD rcvcat rman/rman@CATAL`

If you are using control files, run the following command:

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

2. Start the datafile restore:

- 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=<ORACLE_SID>)';
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=<ORACLE_SID>)';

sql "alter database datafile
'<Oracle_home>\data\oradata\DATA\temp01.dbf'
offline";

restore datafile
'<Oracle_home>\data\oradata\DATA\temp01.dbf';

recover datafile
'<Oracle_home>\data\oradata\DATA\temp01.dbf';

sql "alter database datafile
'<Oracle_home>\data\oradata\DATA\temp01.dbf'
online";

release channel dev1;

}
```

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

UNIX

```
run{
allocate channel dev1 type 'sbt_tape' parms
'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=<ORACLE_SID>)';

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=<ORACLE_SID>)';

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. The procedure in such cases is as follows:

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 are using the recovery catalog database, run the following command:

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

If you are using control files, run the following command:

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

Example of Archive Log Restore

To restore an archive log, follow the steps below:

1. Login to the Oracle RMAN:

If you are using the recovery catalog database, run the following command:

- On UNIX: `<ORACLE_HOME>/bin/rman target sys/manager@PROD rcvcat rman/rman@CATAL`
- On Windows: `<ORACLE_HOME>\bin\rman target sys/manager@PROD rcvcat rman/rman@CATAL`

If you are using control files, run the following command:

- On UNIX: `<ORACLE_HOME>/bin/rman target sys/manager@PROD nocatalog`

- On Windows: `<ORACLE_HOME>\bin\rman target sys/manager@PROD nocatalog`
2. Start the archive log restore:

```
run{  
  allocate channel dev1 type 'sbt_tape' parms  
  'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=<ORACLE_SID>)' ;  
  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. The procedure in such cases is as follows:

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 are using the recovery catalog database, run the following command:

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

If you are using control files, run the following command:

- On UNIX: `<ORACLE_HOME>/bin/rman target sys/manager@PROD nocatalog cmdfile=/var/opt/omni/tmp/restore_arch`
- On Windows: `<ORACLE_HOME>\bin\rman target sys/manager@PROD nocatalog cmdfile=<Data_Protector_home>\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. See “Data Protector Oracle Configuration Files” on page 14.

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

UNIX

- if you are using an `sh` - like shell, enter the following commands:

```
OB2APPNAME="<ORACLE_SID>"
```

```
export OB2APPNAME
```

- if you are using a `csh` - like shell, enter the following commands:

```
setenv OB2APPNAME "<ORACLE_SID>"
```

Windows

```
set OB2APPNAME=<ORACLE_SID>
```

3. Run the following command:

- On HP-UX and Solaris: `/opt/omni/sbin/ob2rman.exe -restore_catalog -session <session_ID> -apphost <application_hostname>`

Restoring an Oracle Database

- On other UNIX: `/usr/omni/bin/ob2rman.exe`
`-restore_catalog -session <session_ID> -apphost`
`<application_hostname>`
- On Windows: `<Data_Protector_home>\bin\ob2rman.exe`
`-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.

After issuing the above command, the `rcvcat.exp` file is restored by the Data Protector `obkrestore` utility, which is then read by the Oracle import utility. The Oracle Import utility then restores the file back to the Oracle database. On UNIX, this file resides in the `/var/opt/omni/tmp` directory.

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/restoreddev` (UNIX systems) or

`<Data_Protector_home>\Config\server\Cell\restoreddev` (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 `restoreddev` 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 Disaster Recovery chapter in the *HP OpenView Storage Data Protector Administrator's 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 74.
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 the `/var/opt/omni/log/oracle8.log` (HP-UX and Solaris systems), `usr/omni/log/oracle8.log` (other UNIX systems), or `<Data_Protector_home>\log\oracle8.log` (Windows systems).
- Oracle writes the logs in the `<Oracle user dump directory>\sbtio.log` file.

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 the Data Protector Database Library. You must rebuild (Oracle8) or re-link (Oracle8i/9i) 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 108 or “Removing the Data Protector Oracle Integration Link on Solaris and other UNIX Systems” on page 109.

Removing the Data Protector Oracle Integration Link on HP-UX Systems

1. On the Oracle Server system, connect to the Oracle database as an Oracle operating system user and shut down all Oracle instances.

Oracle8

2. For Oracle8, perform the following:
 - a. Change to the `<ORACLE_HOME>/rdbms/lib` directory:

```
cd <ORACLE_HOME>/rdbms/lib
```
 - b. Execute the following command:

```
make -f ins_rdbms.mk ioracle
```

IMPORTANT

The `make -f ins_rdbms.mk ioracle` command will work only if the `env_rdbms.mk` file was not changed.

Oracle8i/9i

3. For Oracle8i/9i, perform the following:
 - a. Change to the `<ORACLE_HOME>/lib` directory:

```
cd <ORACLE_HOME>/lib (32-bit Oracle),  
cd <ORACLE_HOME>/lib64 (64-bit Oracle8i) or  
cd <ORACLE_HOME>/lib (64-bit Oracle9i).
```
 - b. If the `libobk.sl.orig` file exists in the `<ORACLE_HOME>/lib` directory, execute the following command:

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

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

Removing the Data Protector Oracle Integration Link on Solaris and other UNIX Systems

1. On the Oracle Server system, connect to the Oracle database as an Oracle operating system user and shut down all Oracle instances.

Oracle8

2. For Oracle8, perform the following:
 - a. Change to the `<ORACLE_HOME>/rdbms/lib` directory:

```
cd <ORACLE_HOME>/rdbms/lib
```
 - b. Execute the following command:

```
make -f ins_rdbms.mk ioracle
```

IMPORTANT

The `make -f ins_rdbms.mk ioracle` command will work only if the `env_rdbms.mk` file was not changed.

Oracle8i/9i

3. For Oracle8i/9i, perform the following:
 - a. Change to the `<ORACLE_HOME>/lib` directory:

```
cd <ORACLE_HOME>/lib (32-bit Oracle),  
cd <ORACLE_HOME>/lib64 (64-bit Oracle8i) or
```

```
cd <ORACLE_HOME>/lib (64-bit Oracle9i).
```

- b. 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.

4. Start all Oracle instances.

Oracle8i/9i/10g RMAN Metadata and Data Protector Media Management Database Synchronization

This section describes how to synchronize Oracle8i/9i/10g Recovery Manager metadata with the Data Protector Media Management Database.

Recovery Manager 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. In order to have both catalogs synchronized, run the following command using RMAN:

```
allocate channel for maintenance type 'sbt_tape';  
crosscheck backup;  
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;
```

Refer to the *Oracle8i Recovery Manager User's Guide and References* for more details on recovery catalog maintenance.

TIP

[REDACTED]

[REDACTED]

[REDACTED]

Troubleshooting

This section lists the most common problems you may encounter while working with the integration.

The following sections provide some checking procedures you should perform before you call the Data Protector support. In this way, you may either resolve the problem yourself or identify the area where the difficulties are occurring.

Should you fail when performing a troubleshooting procedure, actions are proposed to help you work around the problem.

Before You Begin

- Ensure that the latest official Data Protector patches are installed. Refer to the “Verifying Which Data Protector Patches Are Installed” section in the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

Refer to the *HP OpenView Storage Data Protector Software Release Notes* or

http://www.openview.hp.com/products/datapro/spec_0001.html for an up-to-date list of supported versions, platforms, and other information.

- Refer to the *HP OpenView Storage Data Protector Software Release Notes* for a description of general Data Protector limitations, problems and workarounds, as well as a list of available Data Protector patches.

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 the Data Protector Database Library. You must rebuild (Oracle8) or re-link (Oracle8i/9i) 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 Systems” on page 108 for more information on how to make the Oracle server functional again.

Setting Up the Environment Variables

If you need to export some variables before starting the Oracle8 Server Manager or Oracle9i SQL Plus, TNS listener, or Recovery Manager, these variables must be exported as described in “Setting, Retrieving, and Listing Data Protector Oracle Configuration Files Parameters Using the CLI” on page 16.

General Troubleshooting

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

Problem

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

```
*RETVAL*8422
```

Action

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

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, refer to the Oracle documentation.

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) or SQL Plus (Oracle9i) from the `<ORACLE_HOME>` directory:

```
bin/svrmgrl (Oracle8/8i) or  
bin/sqlplus /nolog (Oracle9i).
```

At the SVRMGR (Oracle8) or SQL (Oracle9i) prompt, 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 `<ORACLE_SID>` as described on page 115.

Start the Server Manager (Oracle8) or SQL Plus (Oracle9i) from the `<ORACLE_HOME>` directory:

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

```
bin/sqlplus /nolog (Oracle9i).
```

At the SVRMGR (Oracle8) or SQL (Oracle9i) prompt 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 115.

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

```
bin/lsnrctl80 status <service> (for Oracle8) or
bin/lsnrctl status <service> (for Oracle8i/9i)
exit
```

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`).

Export `<ORACLE_HOME>` as described on page 115.

Start the Server Manager (Oracle8) or SQL Plus (Oracle9i) from the `<ORACLE_HOME>` directory:

```
bin/svrmgrl (Oracle8) or
bin/sqlplus /nolog (Oracle9i).
```

At the SVRMGR (Oracle8) or SQL (Oracle9i) prompt type:

```
connect <Target_Database_Login>
exit
and then
connect <Recovery_Catalog_Login>
exit
```


If this fails, refer to 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 115.

Start the Server Manager (Oracle8) or SQL Plus (Oracle9i) from the `<ORACLE_HOME>` directory:

`bin/svrmgrl` (Oracle8) or

`bin/sqlplus /nolog` (Oracle9i).

At the SVRMGR (Oracle8) or SQL (Oracle9i) prompt type:

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

```
exit
```

and

```
bin/svrmgrl connect <Recovery_Catalog_Login> as SYSDBA  
(Oracle8) or
```

```
bin/sqlplus connect <Recovery_Catalog_Login> as SYSDBA  
(Oracle9i)
```

```
exit
```

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

If this fails, refer to the Oracle documentation for instructions about how to set up the password file and any relevant parameters in the `init<ORACLE_SID>.ora` file.

5. **If you are using the recovery catalog database, verify that the Target Database is registered in the recovery catalog:**

Export `<ORACLE_HOME>` as described on page 115 and start the Oracle8 Server Manager or Oracle9i SQL Plus:

`bin/svrmgrl` (Oracle8) or

`bin/sqlplus /nolog` (Oracle9i).

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

```
connect <Recovery_Catalog_Login>;
```

```
select * from rc_database;
```

```
exit
```

If this fails, start the configuration using Data Protector or refer to 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 a Recovery Manager channel type disk.

If you are using the recovery catalog:

Export `<ORACLE_HOME>` as described on page 115 and start Recovery Manager:

```
bin/rman target <Target_Database_Login> rcvcat  
<Recovery_Catalog_Login> cmd_file=rman_script
```

If you are not using the recovery catalog:

Export `<ORACLE_HOME>` as described on page 115 and start Recovery Manager: `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, refer to the Oracle documentation for details about how to execute a backup and restore directly to disk using the Recovery Manager.

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, refer to the Oracle documentation.

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

Set `<ORACLE_HOME>` and `<ORACLE_SID>`:

Start the Server Manager (Oracle8) or SQL Plus (Oracle9i/10g) from the `<ORACLE_HOME>` directory:

`bin\svrmgrl` (Oracle8) or

`sqlplus /nolog` (Oracle9i/10g)

At the SVRMGR (Oracle8) or SQL (Oracle9i/10g) prompt, 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 `<ORACLE_SID>`.

Start the Server Manager (Oracle8) or SQL Plus (Oracle9i/10g) from the `<ORACLE_HOME>` directory:

`bin\svrmgrl` (Oracle8) or

`sqlplus /nolog` (Oracle9i/10g)

At the SVRMGR (Oracle8) or SQL (Oracle9i/10g) prompt, 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:

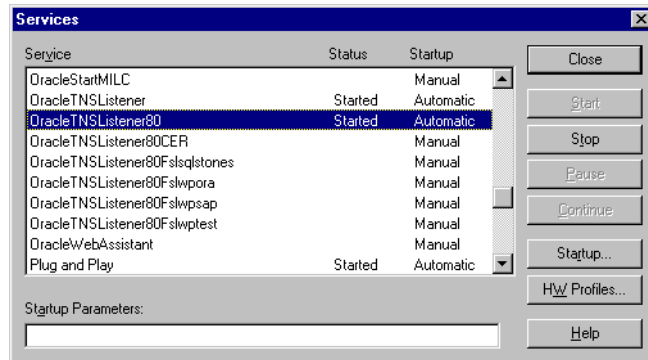
Note that with Oracle8i/9i/10g, the `lsnrctl` and not the `lsnrctl80` command is used.

```
bin\lsnrctl80 status <service>
```

If this fails, startup 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, click Administrative Tools, Services.

Figure 1-22 **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:

```
bin\svrmgrl (Oracle8) or
```

```
sqlplus /nolog (Oracle9i/10g)
```

At the SVRMGR (Oracle8) or SQL (Oracle9i/10g) prompt, type:

```
connect <Recovery_Catalog_Login>
```

```
exit
```

If this fails, refer to 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.

For Oracle8, start the Server Manager or SQL Plus from the `<ORACLE_HOME>` directory:

```
bin\svrmgrl (Oracle8) or
```

```
bin\sqlplus /nolog (Oracle9i/10g)
```

At the SVRMGR (Oracle8) or SQL (Oracle9i/10g) prompt, type:

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

```
exit
```

Connect to the recovery catalog:

For Oracle8:

```
bin\svrmgrl connect <Recovery_Catalog_Login> as SYSDBA
```

For Oracle9i/10g:

```
bin\sqlplus connect <Recovery_Catalog_Login> as SYSDBA
```

```
exit
```

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

If this fails, refer to the Oracle documentation for instructions about how to set up the password file and any relevant parameters in the `init<ORACLE_SID>.ora` file.

5. **If you are using the recovery catalog database, verify that the Target Database is registered in the recovery catalog:**

```
bin\svrmgrl (Oracle8) or
```

```
bin\sqlplus (Oracle9i/10g)
```

At the SVRMGR (Oracle8/8i) or SQL (Oracle9i/10g) prompt, type:

```
connect <Recovery_Catalog_Login>;
```

```
select * from rc_database;
```

```
exit
```

If this fails, start the configuration using Data Protector or refer to 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 a Recovery Manager channel type disk.**

If you are using the recovery catalog:

Set `<ORACLE_HOME>` and start Recovery Manager from the `<ORACLE_HOME>` directory:

For Oracle8:

```
bin\rman80 target <Target_Database_Login> rcvcat  
<Recovery_Catalog_Login> cmd_file=rman_script
```

For Oracle8i/9i/10g:

```
bin\rman target <Target_Database_Login> rcvcat  
<Recovery_Catalog_Login> cmd_file=rman_script
```

If you are not using the recovery catalog:

Set `<ORACLE_HOME>` and start Recovery Manager from the `<ORACLE_HOME>` directory:

For Oracle8:

```
bin\rman80 target <Target_Database_Login> nocatalog  
cmd_file=rman_script
```

For Oracle8i/9i/10g:

```
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, refer to the Oracle documentation for details about how to execute a backup and restore directly to disk using the Recovery Manager.

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

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.sl` (`libob2oracle8_64bit.sl`) file is linked with the Oracle8 executable. Note that on HP-UX IA-64 and Solaris systems, the extension for the Data Protector Database Library is `.so`, and on AIX, the extension is `.a`.

Export `<ORACLE_HOME>` and `<ORACLE_SID>` as described on page 115.

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 Data Protector library 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 the Data Protector Database Library dynamic path as follows:

```
/usr/bin/chatr +s enable <ORACLE_HOME>/bin/oracle
```

On Solaris and, 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 1-23

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.

Refer to the *HP OpenView Storage Data Protector Administrator's Guide* 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 an Oracle User in Data Protector on UNIX Systems” on page 29.

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**

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

2. **Verify that the Data Protector Database Library is loaded**

Once the setup has finished, you need to restart the Oracle services and verify that the

`<Drive_Label>:\<%SystemRoot%\system32\orasbt.dll` Data Protector Database Library 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.

Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for details about how to do a filesystem backup.

4. Verify the inet startup parameters:

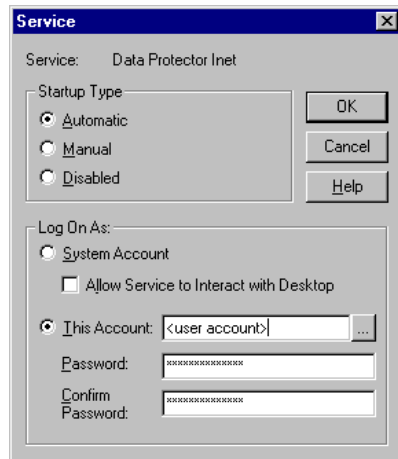
Check the Data Protector Inet service startup parameters on the Oracle Server system. Proceed as follows:

- 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 1-24

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 an Oracle User in Data Protector on UNIX Systems” on page 29. The identified Oracle user and the user `root` must also be added to Data Protector `admin` or `operator` group. Then run the following command:

```
/opt/omni/sbin/util_oracle8.exe -CHKCONF <ORACLE_SID>  
(HP-UX and Solaris systems) or
```

```
/usr/omni/bin/util_oracle8.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/sbin/omnigetmsg <set_number> <Error_number>  
(HP-UX and Solaris systems) or
```

```
/usr/omni/bin/omnigetmsg <set_number> <Error_number>  
(other UNIX systems).
```

The `*RETVAL*0` indicates successful configuration.

IMPORTANT

It is possible to receive a `*RETVAL*0` even though the backup still fails. This can happen due to the fact that the backup owner is not the Oracle user `root` or the Oracle user that is identified as described in “Configuring an Oracle User in Data Protector on UNIX Systems” on page 29.

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:<ORACLE_SID>  
-bar:<backup_specification_name> -perform:backup
```

On other UNIX systems:

```
/usr/omni/bin/testbar2 -type:Oracle8 -appname:<ORACLE_SID>  
-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. Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for instructions on troubleshooting devices.

Data Protector reports “Export of the Recovery Catalog Database Failed” when backing up Oracle9i

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.exe@machine "MAKI" Time: 10/01/01 16:07:53

Export of the Recovery Catalog Database failed.

Action

Log in to Oracle9i SQL Plus and grant the execute permission to the Oracle9i LT_EXPORT_PKG as follows (make sure that the user sys has the SYSDBA permission 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. Refer to the *HP OpenView Storage Data Protector Administrator's 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:

```
rman target <Target_Database_Login> rcvcat  
<Recovery_Catalog_Login>
```

```
RMAN> RESET DATABASE;
```

```
RMAN> exit
```

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.
- Replace the `<ORACLE_SID>` in the `OB2APPNAME` parameter with `<GLOBAL_DB_NAME>`.
- 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=<GLOBAL_DB  
NAME>,OB2BARLIST=RAC_arch)' connect  
<username>/<passwd>@<INSTANCE 1>;  
  allocate channel 'dev_2' type 'sbt_tape'  
    parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=<GLOBAL_DB  
NAME>,OB2BARLIST=RAC_arch)' connect  
<username>/<passwd>@<INSTANCE 2>;  
  backup  
    format 'RAC_arch<QU_%s:%t:%p>.dbf'  
    archive log 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, start the following command on the Oracle Server system:

```
<Data_Protector_home>\bin\util_oracle8.exe -CHKCONF  
<ORACLE_SID>
```

The *RETVAL*0 indicates successful configuration.

IMPORTANT

It is possible to receive a *RETVAL*0 even though the backup still fails. This can happen due to the fact that the backup owner is not the Oracle user root or the Oracle user that is identified as described in “Configuring an Oracle User in Data Protector on UNIX Systems” on page 29.

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 the following command:

```
<Data_Protector_home>\bin\testbar2 -type:Oracle8  
-appname:<ORACLE_SID> -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. Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for instructions on troubleshooting devices.

Data Protector reports “Export of the Recovery Catalog Database Failed” when backing up Oracle9i/10g

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.exe@machine "MAKI" Time: 10/01/01 16:07:53

Export of the Recovery Catalog Database failed.

Action

Log in to Oracle9i/10g SQL Plus and grant the execute permission to the Oracle9i/10g LT_EXPORT_PKG as follows (make sure that the user sys has the SYSDBA permission 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:

```
rman target <Target_Database_Login> rcvcat  
<Recovery_Catalog_Login>
```

```
RMAN> RESET DATABASE;
```

```
RMAN> exit
```

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, refer to 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:<ORACLE_SID>
                   -perform:restore
```

```
-object:<object_name>  
-version:<object_version>  
-bar:<backup_specification_name>
```

Other UNIX

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

Windows

```
<Data_Protector_home>\bin\testbar2 -type:Oracle8  
-appname:<ORACLE_SID>  
-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 91.

6. Try putting the database into the Open state manually after using the Data Protector GUI to recover and restore a Zero Downtime Backup or Standalone Backup session.

If you have used the Data Protector GUI to recover and restore a Zero Downtime Backup (ZDB) session or a standalone backup, 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;
```

“Binary util_orarest is missing” error message is displayed when browsing Oracle9i database for restore on Linux

Problem

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 <ORACLE_SID>
where <ORACLE_SID> is the Oracle SID of the target database.
```

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

“Binary util_orarest failed” error message is displayed when browsing Oracle9i database for restore on Linux or Tru64

Problem

The following error message is displayed when browsing *Oracle9i* database for restore on Linux or Tru64:

```
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 or Tru64):

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

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

Problem The Recovery Catalog was not used, the RMAN autobackup feature was not used (for Oracle9i/10g), and the control file cannot be restored from Data Protector managed backup. A valid control file backup exists on tape.

Action

- For Oracle8/8i, 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=<ORACLE_SID>,
    OB2BARHOSTNAME=<hostname>)' );
    dbms_backup_restore.restoresetdatafile;

    dbms_backup_restore.restorecontrolfileto('/tmp/tmp.cf');
    dbms_backup_restore.restorebackuppiece('<backup piece
handle>',done=>done);
END;
```

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

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

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

```
replicate controlfile from '/tmp/foo.cf';
sql 'alter database mount';
set until time 'MMM DD YY HH24:MM:SS';
restore database;
recover database;
sql 'alter database open resetlogs';
release channel 'dev_0';
}
```

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

- For Oracle9i/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.

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:

- “Overview” on page 143
- “Prerequisites and Limitations” on page 145
- “Integration Concept” on page 146
- “Data Protector SAP R/3 Configuration File” on page 155
- “Configuring the Integration” on page 162
- “Configuring an SAP R/3 Backup” on page 176
- “Backing Up an SAP R/3 Database” on page 194
- “Restoring an SAP R/3 Database” on page 201
- “Monitoring an SAP R/3 Backup and Restore” on page 207
- “Troubleshooting” on page 209
- “Examples of SAP R/3 Database Restore” on page 234

Overview

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 Software Release Notes* or http://www.openview.hp.com/products/datapro/spec_0001.html 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 Software Release Notes* or http://www.openview.hp.com/products/datapro/spec_0001.html 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 Software Release Notes* or http://www.openview.hp.com/products/datapro/spec_0001.html 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).

Limitations

Refer to the *HP OpenView Storage Data Protector Software Release Notes* 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 [REDACTED]

as well as [REDACTED]

[REDACTED]. 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*.

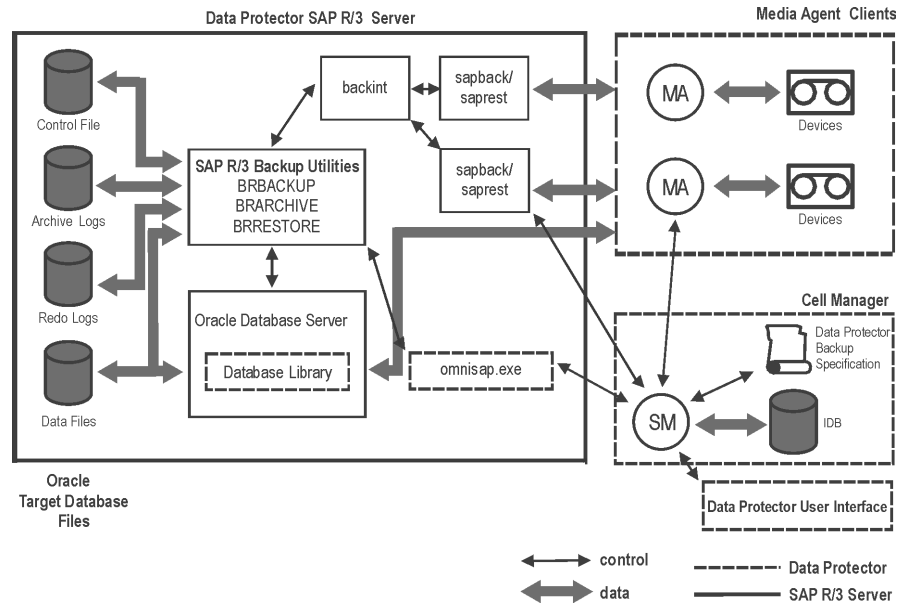
**Data Protector
Integration
Software**

The Data Protector integration software consists of the following components, as depicted in Figure 2-1 on page 148.

- 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

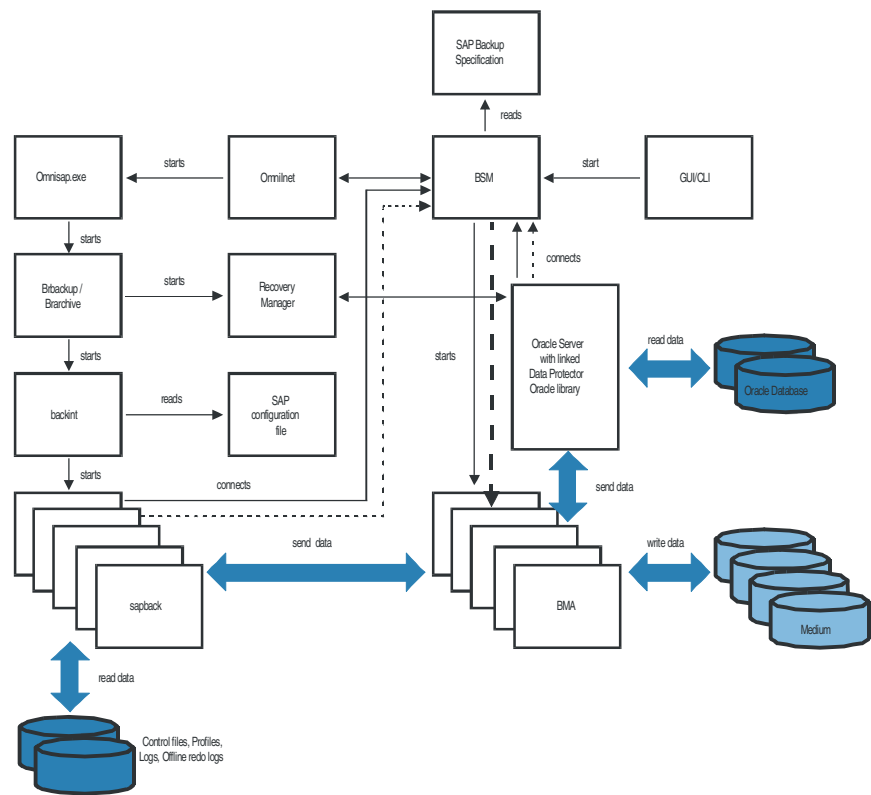
A backup session using RMAN mode differs from a backup session in backint mode in step 3. See Figure 2-3 on page 153 for details.

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_specification_name>={'<filename>'=<device_
number>;
}
}
```


Example

This is an example of the file:

```
ORACLE_HOME='/app/oracle805/product';
ConnStr='EIBBKIBBEIIBBFIBBGHBBBOHBBQDBBOFBBCFBBPFBBCFBBI FBBGFBBBDG BBB
FBB CFBBDFBB CFBB';

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



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

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.

Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for a detailed information on Data Protector user rights and how to add a user to a user group.

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.

Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for detailed instructions about filesystem backups.

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 155 for information on how to do that.

**Cluster-Aware
Clients on UNIX**

When configuring the Data Protector SAP R/3 integration, configure it only on 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. 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/sbin/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. Refer to *HP OpenView Storage Data Protector Administrator's Guide* for information on how to add a user to Data Protector.

For information on the Data Protector Cell Manager package configuration (if you want to install and configure an Data Protector Cell Manager in the MC/SG cluster), refer also to the *HP OpenView Storage Data Protector Administrator's Guide*.

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_DEA 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 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 183.

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 162.

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

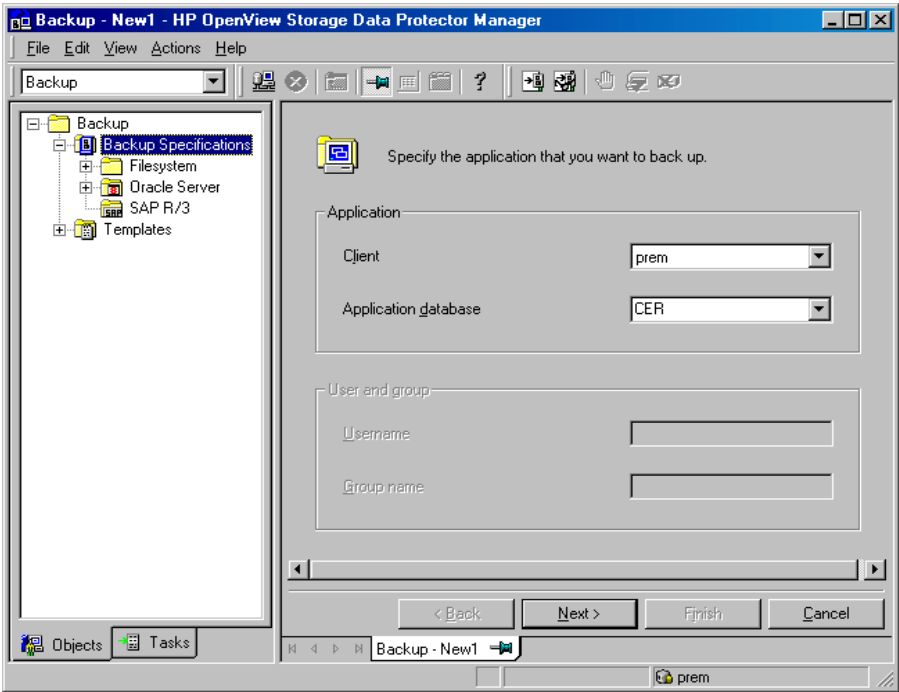
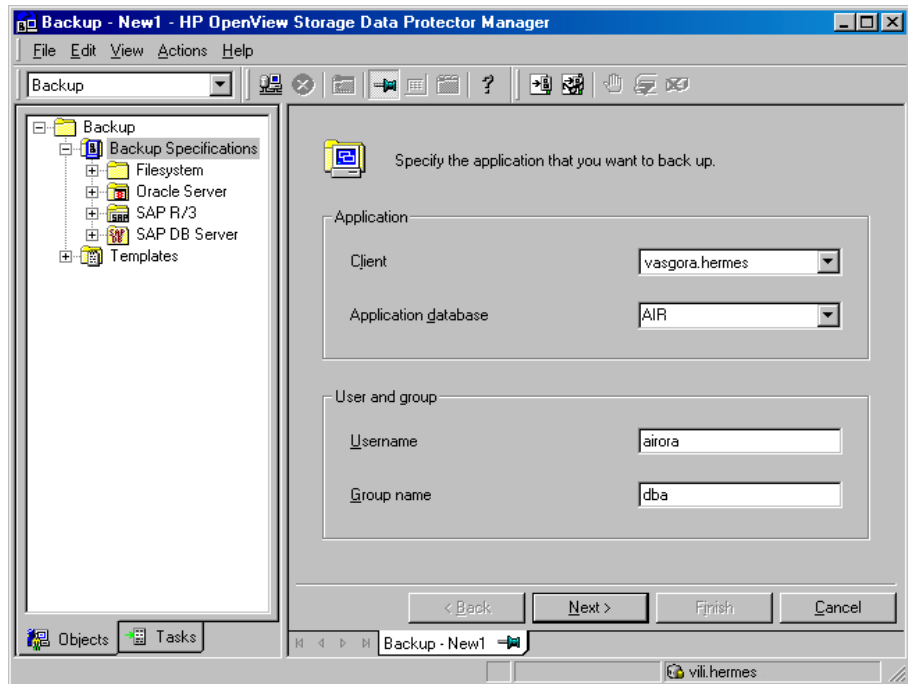


Figure 2-6 Specifying the SAP R/3 Database Server and the Oracle SID on UNIX



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**

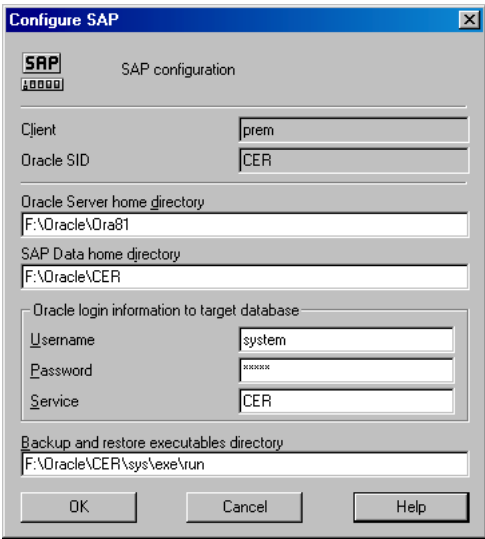
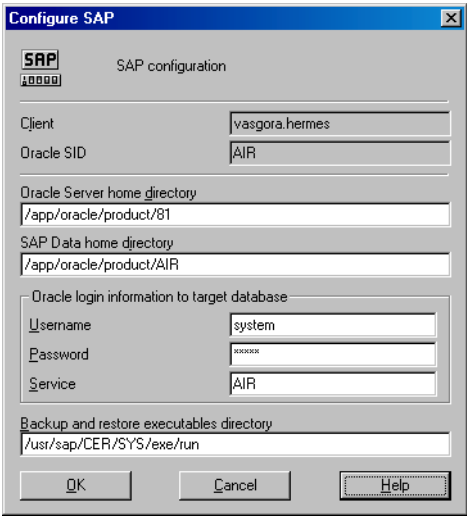


Figure 2-8 **Configuring an SAP R/3 Database Server on UNIX**



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/sbin` (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.

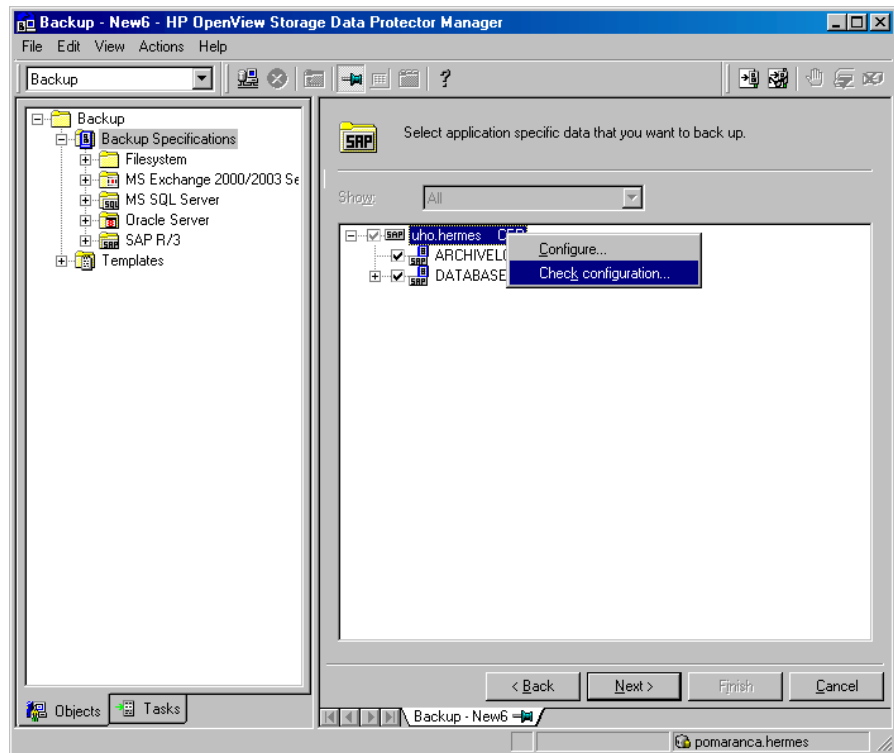
NOTE

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/sbin/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. Refer to 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 177.
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 187.

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	Archives the offline redo logs.
Brarchive_SaveDelete	Archives 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.

Brbackup_SMB_Offline

Used for an offline ZDB (split mirror or snapshot backup). The database is stopped during the creation of a replica.

Brbackup_SMB_Online

Used for an online ZDB (split mirror or snapshot backup). The database is active during the creation of a replica.

Click OK.

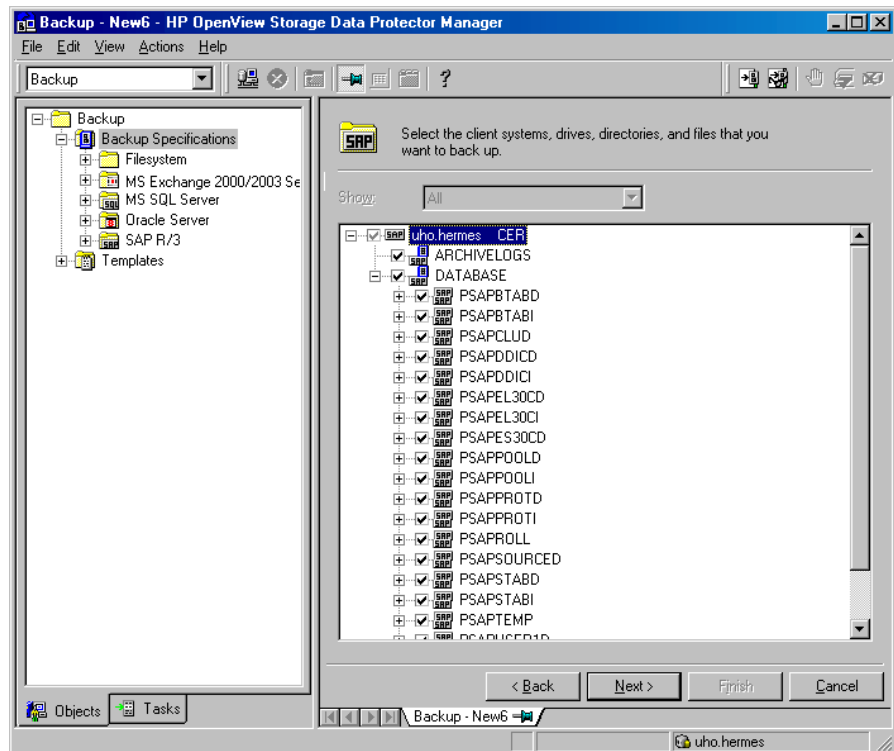
4. In the Results Area, provide the following information:

- 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 162.

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 164 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 183, 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 *HP OpenView Storage Data Protector Administrator's Guide*.

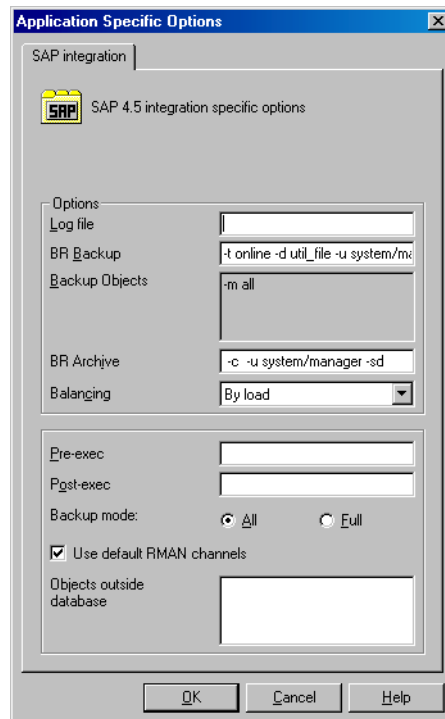
Click **Next**.

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 183 or online **Help**.

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 195.

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 194.

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 162 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 192 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. See the *HP OpenView Storage Data Protector Administrator's Guide* for more information on 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 155 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 190 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 and the *HP OpenView Storage Data Protector Administrator's Guide* 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 Backup**

In order to use SAP R/3 with the Oracle RMAN utility in order to back up the Oracle Target Database data files, you need to link the Oracle Server with the Data Protector Database Library. Refer to “Linking Oracle with the Data Protector Database Library on UNIX Systems” on page 22 for more information.

**Creating an SAP
R/3 Backup
Specification**

To create a backup specification that enables you to use RMAN for backing up, specify the following BRBACKUP option:

`-d rman_util` (using BRBACKUP) or `backup_dev_type` (in the SAP parameter file)

NOTE

If the BRBACKUP option `-d` is not specified, the default value is taken from the SAP parameter file. In this case you have to make sure that the value for `rman_params` is set correctly in the SAP parameter file.

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 155.

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 | <ul style="list-style-type: none">• Schedule a backup of an existing SAP R/3 backup specification using the Data Protector Scheduler. See “Scheduling a Backup” on page 195.• 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 197.• Start an interactive backup on SAP R/3 Database Servers using either the <code>brbackup</code> or the <code>sapdba</code> command. See “Using SAP R/3 Commands” on page 199. |
|-----------------------|--|

NOTE	<p>If you use <code>brbackup</code> or <code>sapdba</code> to start a backup session, you do not receive any Data Protector messages about the progress of the session.</p>
-------------	---

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	<p>Concurrent backups of systems with the same Oracle SID in the same cell are not supported.</p>
-----------------------	---

Backup Modes	<p>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.</p>
---------------------	--

Incremental Backups	<p>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):</p> <p><code>-m full</code> (using <code>BRBACKUP</code>) or <code>backup_mode=full</code> (in the SAP parameter file).</p>
----------------------------	---

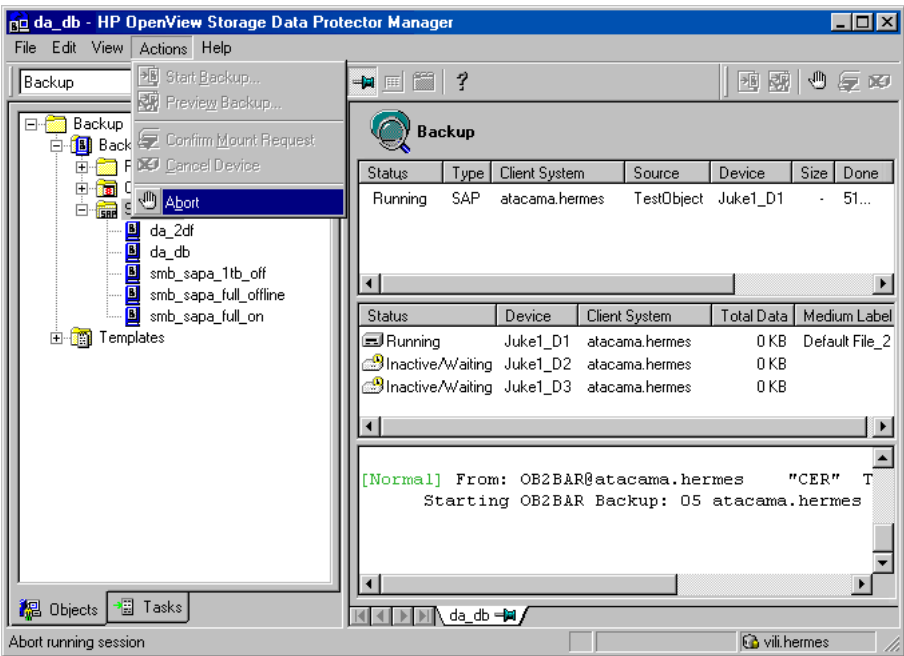
To start an incremental backup, specify Incremental mode in the Data Protector GUI or incr mode in the CLI, as follows:

```
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.

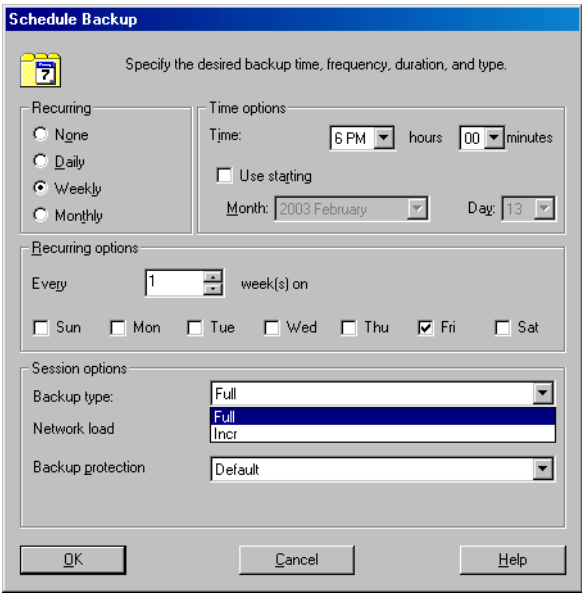
Some examples of scheduling backups of production databases:

- 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.

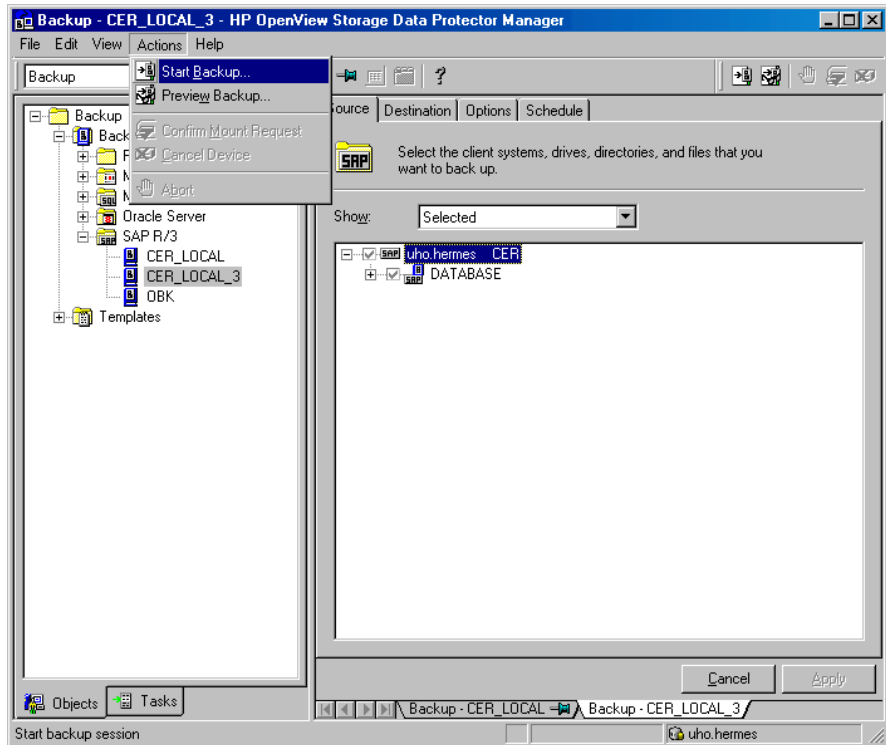
Starting a Backup Using the GUI To start an interactive backup of an SAP R/3 database using the Data Protector GUI, proceed as follows:

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>`

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

An SAP R/3 database can be restored:

- Using the Data Protector GUI or CLI. A standard filesystem restore is performed.
- Using SAP R/3 commands

Considerations



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.

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

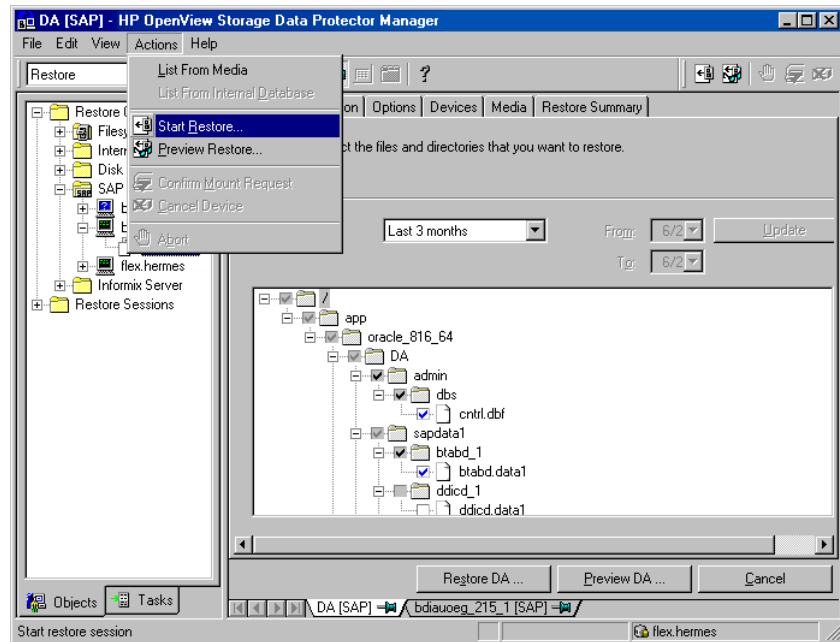
To restore the SAP R/3 objects using the Data Protector GUI, proceed as follows:

Integrating SAP R/3 and Data Protector

Restoring an SAP R/3 Database

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>`

The `omnir` Command

Using the CLI, execute the following command:

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

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.

Example (UNIX)

```
omnir -sap corsa.hermes.si:ABA.0 -session 2000/02/23-1 -tree  
/app/oracle805/ABA/sapdata1/btabd_1/btabd_1.dat
```

Example (Windows)

```
omnir -sap corsa.hermes.si:ABA.0 -session 2000/02/23-1 -tree  
C:\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, refer to the “Restoring Under Another Device” section of the *HP OpenView Storage Data Protector Administrator’s Guide* 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. Also see the Disaster Recovery chapter in the *HP OpenView Storage Data Protector Administrator's 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 either 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 channel 'dev0' type disk;  
  replicate controlfile from '<TMP_FILENAME>';  
  release channel 'dev0';  
}
```

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

Monitoring 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 the most common problems you may encounter while working with the integration.

The following sections provide some checking procedures you should perform before you call the Data Protector support. In this way, you may either resolve the problem yourself or identify the area where the difficulties are occurring.

Should you fail when performing a troubleshooting procedure, measures have been proposed to help you work around the problem.

Before You Begin

- Ensure that the latest official Data Protector patches are installed. Refer to the “Verifying Which Data Protector Patches Are Installed” section in the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

Refer to the *HP OpenView Storage Data Protector Software Release Notes* or

http://www.openview.hp.com/products/datapro/spec_0001.html for an up-to-date list of supported versions, platforms, and other information.

- Refer to the *HP OpenView Storage Data Protector Software Release Notes* for a description of general Data Protector limitations, problems and workarounds, as well as a list of available Data Protector patches.

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
```

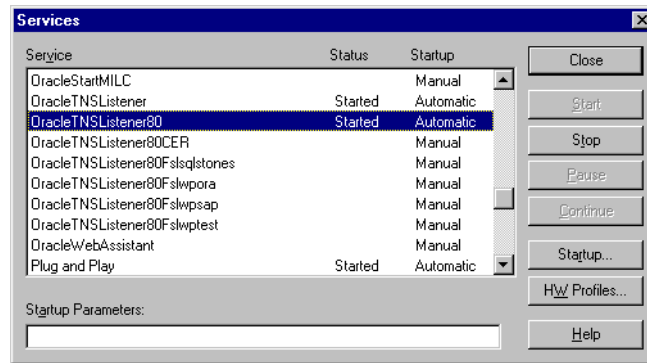
```
bin\lsnrctl status <service> (for Oracle8i/9i)
```

```
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 2-17 Checking the Status of the Oracle Listener



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

bin\svrmgr1 (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 220 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 164.

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.

Refer to the *HP OpenView Storage Data Protector Administrator's Guide* 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:**

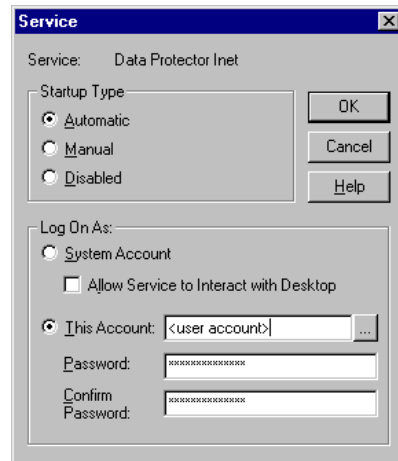
Check the Data Protector Inet service startup parameters on the SAP R/3 Database Server system. Proceed as follows:

- 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 2-18 **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 155.

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 Administrator's 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 158.
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 201.

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>
```

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 Systems” on page 108 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 220 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 220 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 220 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 220 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 220 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.

Refer to “Linking Oracle with the Data Protector Database Library on UNIX Systems” on page 22 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 162:

```
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 164.

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 220.

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.

Refer to the *HP OpenView Storage Data Protector Administrator's Guide* 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 155.

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 162.

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/utlins/testbar2 -type:SAP
-appname:<ORACLE_SID> -bar:<backup_specification_name>
-perform:backup (HP-UX and Solaris systems)

/usr/omni/bin/utlins/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 162 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 Administrator's 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 158.
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 162

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 201.

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/utlins/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/utlins/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>
```

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:

- “Example of Full Database Restore and Recovery” on page 236
- “Example of Partial Restore” on page 240
- “Example of Lost Files Restore” on page 241
- “Example of Archive Log Files Restore” on page 242

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: $n=1, \dots, 99$. 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 2-19

Starting the SAPDBA in Expert Mode

```
SAPDBA V4.6A - SAP Database Administration

ORACLE version: 8.0.5.0.0
ORACLE_SID      : ABA
ORACLE_HOME     : /app/oracle805/product
DATABASE        : open
SAPR3           : not connected

a - Startup/Shutdown instance    h - Backup database
b - Instance information         i - Backup offline redo logs
c - Tablespace administration   j - Restore/Recovery
d - Reorganization              k - DB check/verification
e - Export/import               l - Show/Cleanup
f - Archive mode                m - User and Security
g - Additional functions        n - SAP Online Help

q - Quit

Please select ==> m

User and Security

a - Expert mode
b - User information
c - Role information
d - Restricted mode
p - Change password

q - Return

Please select ==> a
```

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

Figure 2-20

Selecting the Restore/Recovery Option

```
SAPDBA V4.6A - SAP Database Administration

ORACLE version: 8.0.5.0.0
ORACLE_SID      : ABA
ORACLE_HOME     : /app/oracle805/product
DATABASE        : open
SAPR3           : not connected

a - Startup/Shutdown instance    h - Backup database
b - Instance information          i - Backup offline redo logs
c - Tablespace administration    j - Restore/Recovery
d - Reorganization              k - DB check/verification
e - Export/import               l - Show/Cleanup
f - Archive mode                m - User and Security
g - Additional functions         n - SAP Online Help

q - Quit

Please select ==> j
```

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 2-21

Selecting Full Restore and Recovery

```
Restore / Recovery (2001-10-09)

a - Partial restore and complete recovery (Check and repair,
redo logs and control files are prerequisites)
b - Full restore and recovery
(excl. redo logs, control files incl. if required)
c - Reset database
(incl. redo logs and control files)

d - Restore one tablespace
e - Restore individual file(s)

h - Help
q - Return

Please select ==> b
```

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 2-22

Selecting the backup type and version for restore

```
a0y0C _____
                                     Full Restore and Recovery (2001-10-09)
_____

DATABASE STATE   : open
RESTORE / RECOVER: disallowed (see status)

A - Select a backup of type           Current setting
    full online/offline (level 0) or   <not selected>
    whole online/offline (all)

c - Recover until                       now
d - Show status
e - Options
g - Restart restore/recover operation

S - Start restore and recover
q - Return

Please select ==> █
```

- 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 2-23

Starting the Restore Session

```
a0y0C _____
                                     Full Restore and Recovery (2001-10-09)
_____

DATABASE STATE   : open
RESTORE / RECOVER: allowed

A - Select a backup of type           Current setting
    full online/offline (level 0) or   bdgjp1a.anf
    whole online/offline (all)         2001-10-08 14.51.06

c - Recover until                       now
d - Show status
e - Options
f - Show/Delete datafiles younger than 2001-10-08 14.51.06
g - Restart restore/recover operation

S - Start restore and recover
q - Return

Please select ==> S█
```

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

Figure 2-24

Selecting Return to restore process and continue option

```
a0y0C _____
                                     Specify Restore Parameters for Backup Files
_____

      Selected  bdgjp1a.anf 2001-10-08 14.51.06

a - BRBACKUP profile           Current value
b - Use (choose) former restores  initABA.sap
c - Clear list of former restores rdgjwtty.rsb
g - Backup utility parameter file util_file
i - Language                    English

q - Return to restore process and continue
r - Cancel restore process

Please select ==> q
```

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 and Data Protector

In This Chapter

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

It is organized into the following sections:

“Prerequisites and Limitations” on page 247

“Overview” on page 248

“Integration Concept” on page 252

“Data Protector SAP DB Configuration File” on page 255

“Configuring the Integration” on page 259

“Backing Up an SAP DB Database” on page 276

“Restoring an SAP DB Database” on page 285

“Monitoring an SAP DB Backup and Restore” on page 305

“Troubleshooting” on page 308

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 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 and Data Protector systems. Refer to:
 - *HP OpenView Storage Data Protector Software Release Notes* 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 integration.
 - SAP DB documentation for information on the SAP DB Server.
- The SAP DB Automatic Log Backup must be activated for an SAP DB instance to enable transactional backup (log backup).

Limitations

Refer to the *HP OpenView Storage Data Protector Software Release Notes* 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 restore sessions.

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

Overview

Backup

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

- **SAP DB data,**
- **SAP DB configuration** and
- **SAP DB 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 database is backed up using the Data Protector SAP DB integration, it can be switched to either the Admin or to the Online mode, depending on the selected options.

Offline backup of SAP DB objects is not integrated. The standard Data Protector filesystem backup can be performed in such a case. Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for more information on how to perform a Data Protector filesystem backup. It is also not possible to perform an integrated Data Protector SAP DB restore from an offline filesystem backup.

SAP DB data and archive logs are backed up or restored in streams, whereas the SAP DB configuration is backed up or restored as ordinary files. After the backup has finished, the archive logs can either be deleted or kept on the SAP DB Server, depending on the selected options.

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

Table 3-1 on page 249 shows what is actually backed up with regards to the selected SAP DB backup type and Data Protector GUI object.

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

		SAP DB Backup Mode		
		Full	Diff	Trans
GUI Selections	Data	data	diff on data	archive logs
	Configuration	configuration	configuration	configuration
	Instance	data + configuration	diff on data + configuration	archive logs + configuration

Restore and Recovery

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

Only a *complete* SAP DB instance can be restored. Using the integration, SAP DB 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 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 Server) during the process of switching the database to the Online mode.

The integration supports SAP DB **migration**, meaning that an SAP DB instance can be restored to an SAP DB Server or instance other than the original. In such a case, if the SAP DB Server has not yet been configured for the Data Protector SAP DB 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 Server are never deleted.

Parallelism	The integration also takes advantage of the concept of SAP DB media and media groups, thus providing parallel backup and restore of SAP DB objects. Several SAP DB media are grouped in an SAP DB media group, which is then backed up or restored in streams. This is referred to as SAP DB parallelism. SAP DB Parallelism can be utilized only if the value of the Data Protector <code>Parallelism</code> 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 Specific Backup Options” on page 265 for more information on the Data Protector <code>Parallelism</code> option.
Permissions	Data Protector backup and restore operations on the SAP DB instance require the SAP DB Restoring backups (Recovery) and Saving backups (Backup) permissions , whereas Data Protector configuration operation requires also the SAP DB Installation management (<code>InstallMgm</code>) and Parameter access (<code>ParamCheckWrite</code>) permissions. The backup can be performed in either the SAP DB Online or Admin mode.
Supported Platforms	Refer to the <i>HP OpenView Storage Data Protector Software Release Notes</i> or http://www.openview.hp.com/products/datapro/spec_0001.html for up-to-date information about platforms supported by the integration.
Advantages	<p>Using Data Protector with the SAP DB Database Server offers several advantages over using SAP DB alone:</p> <ul style="list-style-type: none">• 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 backup using Data Protector, the location of devices is completely transparent to the user. They can be connected to the SAP DB 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 Software Release Notes* or http://www.openview.hp.com/products/datapro/spec_0001.html 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 configuration, providing you with a history of activities that can be queried at a later time.

Integration Concept

The Data Protector SAP DB integration links the SAP DB database management software with Data Protector. From the SAP DB point of view, Data Protector represents a media management utility. On the other hand, the SAP DB 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 database management server and of the `backint` interface for SAP DB 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 Server system, which controls activities between the SAP DB Server and Data Protector backup and restore processes.
- The `sapdb_backint` component, installed on the SAP DB Server system, is a binary interface between Data Protector and backup and restore functionality of the SAP DB.
- The DMA (Data Mover Agent) component, installed on the SAP DB Server system, is the actual data transferring module, called by the `sapdb_backint`.
- The `util_sapdb` utility, which is used by Data Protector to configure an SAP DB instance to use with Data Protector and check the instance configuration.

Supported Interfaces

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

- Data Protector GUI or CLI
- SAP DB 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 Server that is to be backed up. The `sapdbbar.exe` module then starts an SAP DB session using the SAP DB `dbmcli`. The `sapdbbar.exe` module issues `dbmcli` commands that configure SAP DB backup media (parallelism), configure `sapdb_backint` and then start the

backup using SAP DB `dbmcli`. SAP DB then starts the configured `sapdb_backint` component. For every SAP DB medium (pipe) `sapdb_backint` starts a DMA, which transfers the data from SAP DB media (pipes) to Data Protector media. This procedure is the same for full, differential, and transactional backup. Additionally, if the configuration (including media specification and the backup history) is selected for backup, it is backed up directly by the `sapdbbar.exe` module and DMA. The list of configuration files to be backed up is retrieved through `dbmcli`.

See Figure 3-1 on page 254.

NOTE

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

Restore Flow

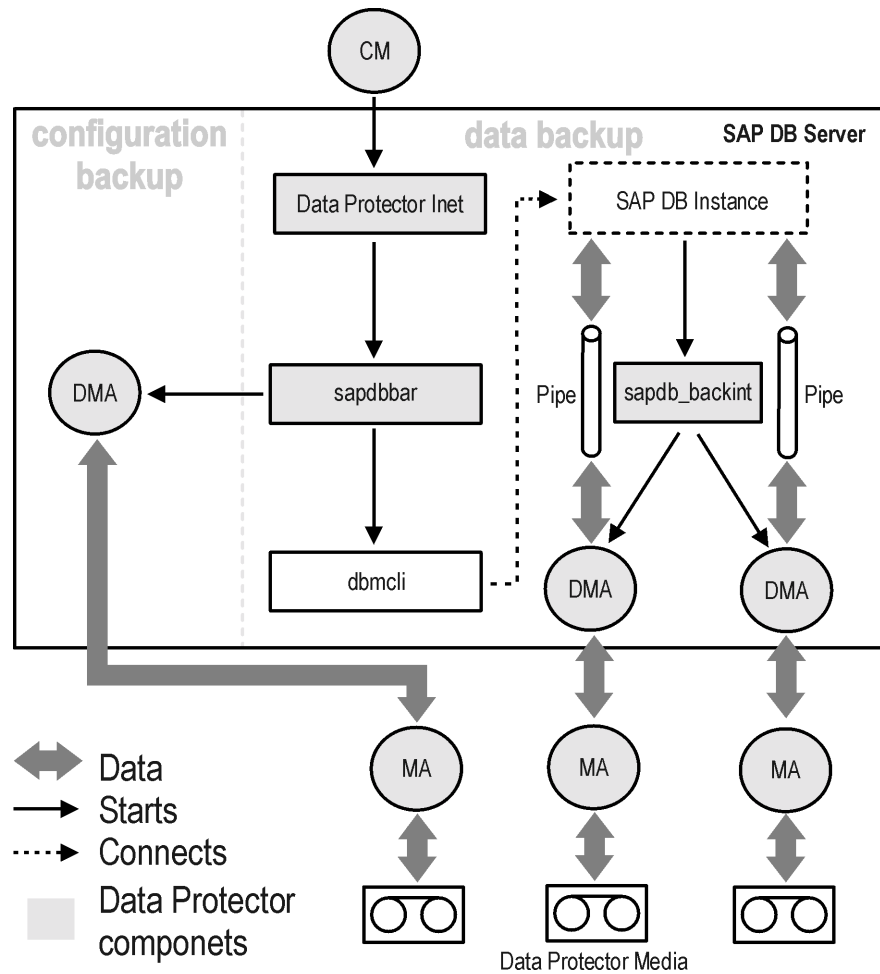
When a restore session is started, the Cell Manager starts the `sapdbbar.exe` module, which starts SAP DB `dbmcli`. The `sapdbbar.exe` module issues commands to SAP DB `dbmcli` to configure `sapdb_backint` and SAP DB backup media (parallelism). SAP DB then starts the configured `sapdb_backint`, which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

See Figure 3-1 on page 254.

NOTE

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

Figure 3-1 **SAP DB Backup and Restore Concept**



Data Protector SAP DB Configuration File

Data Protector stores the SAP DB integration parameters for every configured SAP DB *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 259. These parameters are:

- The username of the SAP DB user created or identified as described in “Configuring Users” on page 259.
- The password of the SAP DB user created or identified as described in “Configuring Users” on page 259.
- The SAP DB version.
- The SAP DB independent program path parameter. This parameter is the independent program path directory specified during the installation of the SAP DB application on the SAP DB Server.
- Data Protector SAP DB integration related environment variables.

NOTE

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

The configuration parameters are written to the Data Protector SAP DB 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 configuration file:

```
Username='dba';
Password='FHBBDHBBCHBB';
Version='7.4.3.27';
Home='/opt/sapdb/indep_prog';
```

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

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

The util_cmd Command You can set, retrieve, or list the Data Protector SAP DB 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 client. The command resides in the `/opt/omni/sbin` (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>] \<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 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 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 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 Server system was successful.

Configuration Overview

To run or schedule Data Protector SAP DB integration backups the following task must be performed:

- An SAP DB user with certain permissions must be added to the Data Protector `admin` group. See “Configuring Users” on page 259.
- Backup devices, media, and media pools must be configured. Refer to the and *HP OpenView Storage Data Protector Administrator's Guide* for information on how to do this.
- The SAP DB 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 Backup” on page 260.
- A Data Protector SAP DB integration backup specification must be created. See “Configuring an SAP DB Backup” on page 260.

SAP DB instance can be reconfigured once it has been configured to be used with Data Protector. See “Modifying the Configuration of an SAP DB Instance in Data Protector” on page 267.

Configuring Users

To perform a backup, restore, or other operation on an SAP DB Server, an SAP DB user granted specific SAP DB permissions must be configured on the SAP DB 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 Server, create or identify an SAP DB user with at least the following SAP DB permissions:
 - Backup
 - Recovery
 - InstallMgm
 - ParamCheckWrite

2. Add the OS user under whose account SAP DB is running to Data Protector admin group. Refer to online Help, index keyword “adding users” for more information on how to do this.
3. Add the operating system root user on the SAP DB 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 Backup

To configure an SAP DB backup, perform the following steps:

1. Configure backup devices, media, and media pools.
Refer to the *HP OpenView Storage Data Protector Administrator's Guide* or online Help for instructions.
2. Create an SAP DB backup specification.

The Data Protector backup specification is stored on the Cell Manager system and contains information needed for the Data Protector SAP DB 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 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 Server is running. In a cluster environment, select the virtual hostname for the systems on which the SAP DB Server is running.

In the Application database drop-down list, all the SAP DB instances located on the SAP DB Server and configured for Data Protector are listed. Enter or select the SAP DB instance to be configured.

On UNIX, enter the user name and the group name for the OS user, under whose account the SAP DB application is running on the SAP DB Server (for example, the sapdb user in the sapsys group).

Figure 3-2

Selecting an SAP DB Server and Instance on UNIX Systems

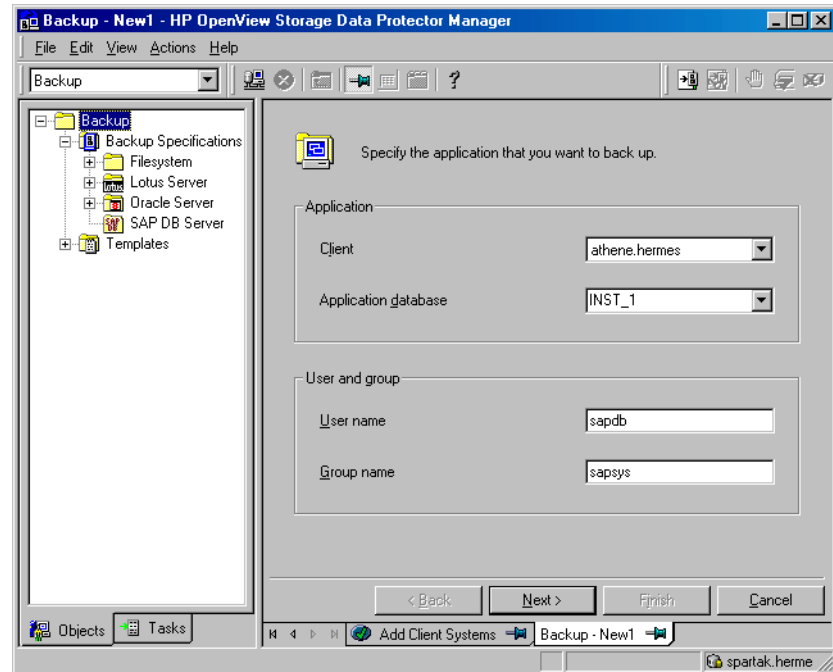
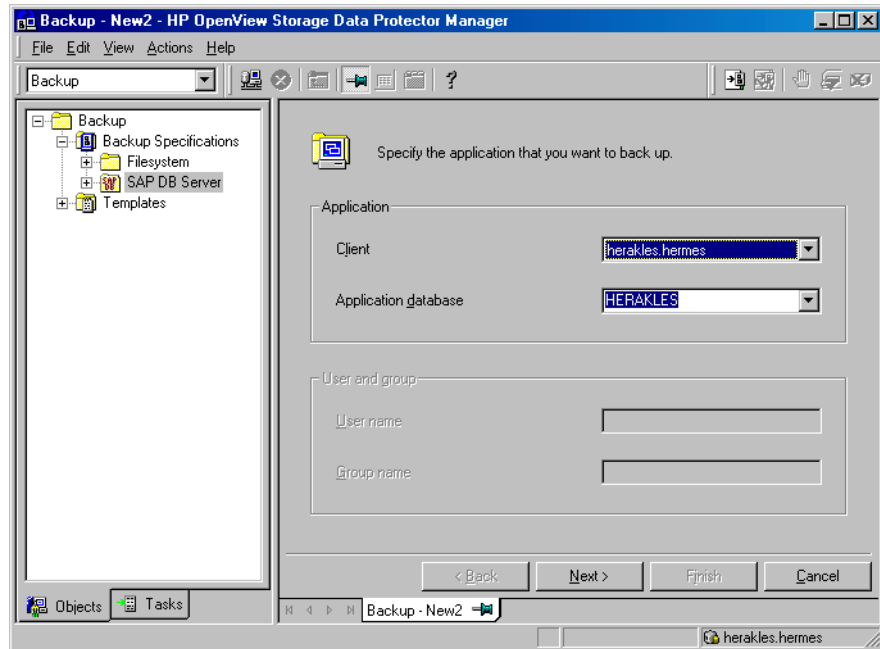


Figure 3-3 **Selecting an SAP DB Server and Instance on Windows Systems**



Click Next.

If the SAP DB 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 application on the SAP DB Server. You can leave the Auto-detect option selected to automatically detect the directory on the SAP DB Server.

Enter the username and the password of the user created or identified as described in “Configuring Users” on page 259.

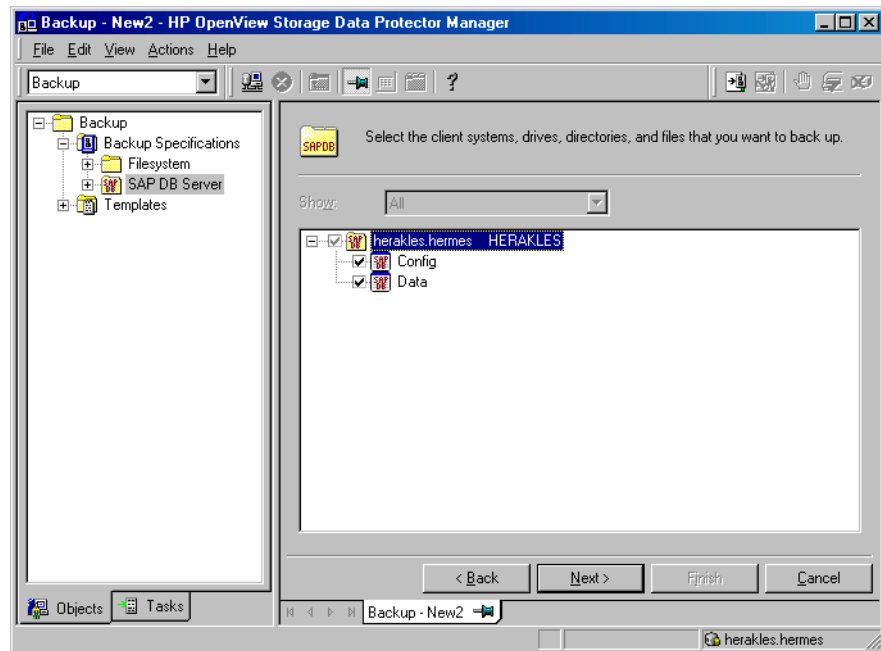
Click OK to confirm the configuration.

6. In the next step of the wizard, select the database objects you want to back up. See Table 3-1 on page 249 for more information on database objects selections.

IMPORTANT

To backup SAP DB 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 Database” on page 276 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 and the *HP OpenView Storage Data Protector Administrator's Guide* 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 integration objects equals the number of devices used for backup.

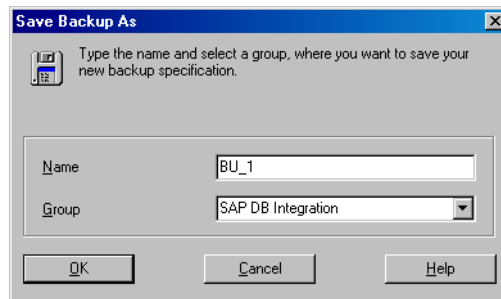
For detailed information on the object mirror functionality, refer to the *HP OpenView Storage Data Protector Administrator's Guide*.

See “SAP DB Specific Backup Options” on page 265 for details about the SAP DB backup options (application specific options).

See “Scheduling Example” on page 277 for information on how to schedule the backup specification.

8. Once you have defined all the backup options, name and save your SAP DB backup specification. It is recommended that you save all the SAP DB 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 Database” on page 276 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 273 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 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 Specific Backup Options

The SAP DB 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 specific backup options:

Change database state Selects the SAP DB 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 archive logs are kept on the SAP DB Server after the backup has finished.

If this option is not selected, the SAP DB archive logs on the SAP DB Server are deleted after the backup has finished.

Parallelism Sets the Parallelism option to set the number of SAP DB media created on the SAP DB Server and consequently the number of SAP DB 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 data volumes to be backed up.

The value of the `Parallelism` option must be equal as or lower than the SAP DB `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, refer to the *HP OpenView Storage Data Protector Administrator's Guide*.

Modifying the Configuration of an SAP DB Instance in Data Protector

The parameters that need to be specified during the configuration of an SAP DB instance in Data Protector are the username and the password of the SAP DB user created or identified as described in the section “Configuring Users” on page 259, and the SAP DB independent program path parameter (the latter is, by default, detected automatically). These parameters are also used for establishing the connection to the SAP DB 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 backup specification, see “Creating a Backup Specification” on page 260.

If properly configured, the SAP DB user entered is allowed to back up or restore SAP DB Server database objects. In order to start a backup of an SAP DB object using Data Protector, the user must also be the owner of the Data Protector backup specification.

Refer to the SAP DB documentation for further information on different types of connections, roles and authorities of SAP DB database administrators and security issues that must be considered.

The configuration of an SAP DB instance in Data Protector can be modified using the Data Protector GUI or CLI.

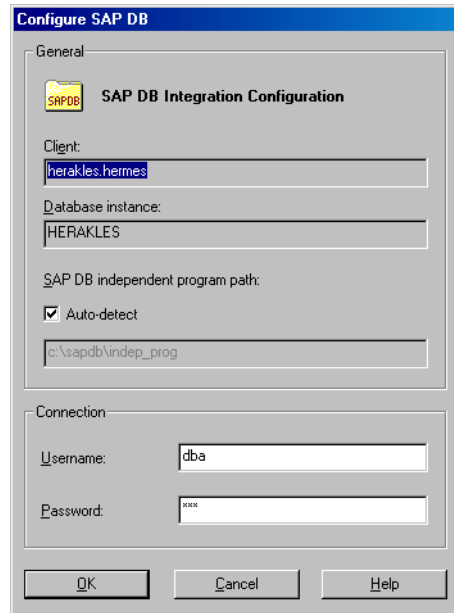
Modifying the Configuration of an SAP DB Instance Using the GUI

The procedure below describes the re-configuration of an SAP DB 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 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 application on the SAP DB Server. You can leave the Auto-detect option selected to automatically detect the directory on the SAP DB Server.

Enter the username and the password of the user created or identified as described in “Configuring Users” on page 259.

Figure 3-7 **SAP DB Configuration**



Click OK to confirm the configuration.

Modifying the Configuration of an SAP DB Instance Using the CLI

UNIX

Login to the SAP DB Server as the user under whose account the SAP DB application is running on the SAP DB 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 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 independent program path parameter. This parameter is the independent program path directory specified during the installation of the SAP DB application on the SAP DB Server.

This parameter is optional. If it is not specified, the directory is detected automatically.

<Instance_Name> The name of the SAP DB instance to be configured.

<username> The username of the SAP DB user created or identified as described in “Configuring Users” on page 259.

`<password>`

The password of the SAP DB user created or identified as described in “Configuring Users” on page 259

NOTE

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

Examples

In the example below, the SAP DB 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 Instance

To check the configuration of an SAP DB instance, use either the GUI or CLI.

Checking the Configuration of an SAP DB Instance Using the GUI

Follow the procedure below to check the configuration of an SAP DB 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 Instance Using the CLI

UNIX

Login to the SAP DB Server system as the user under whose account the SAP DB 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/sbin/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 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 Server system:

On HP-UX:

```
/opt/omni/sbin/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 Server system:

```
<Data_Protector_home>\bin\util_sapdb.exe -chkconf  
<Instance_Name>
```

If an error occurs, it is explained.

Testing the Integration

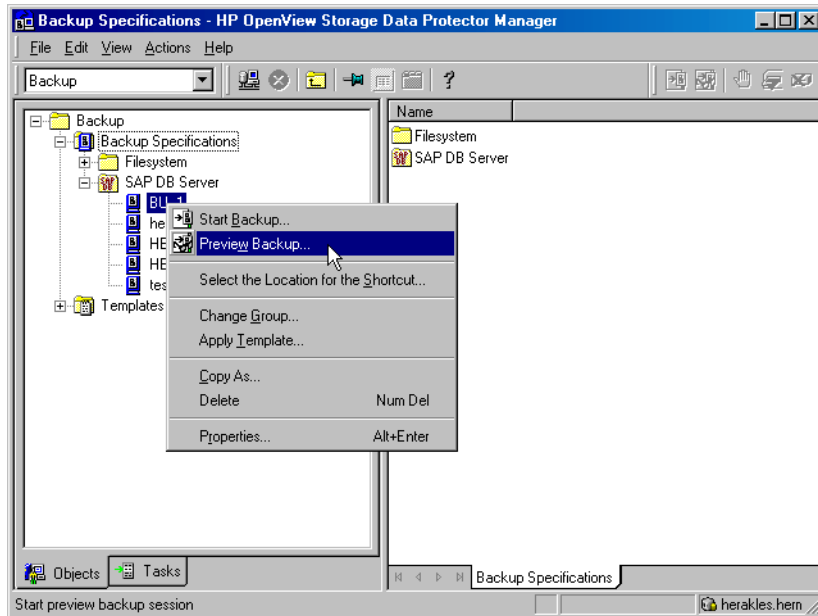
Once you have created and saved a backup specification, you should test it before running a backup.

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 Integration backup specification is correct
- if the devices are correctly configured
- if the required media reside in the devices

After that, the SAP DB 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 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 instance, a regular Data Protector filesystem backup should be configured. Refer to the *HP OpenView Storage Data Protector Administrator's Guide* 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 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 278.

IMPORTANT

To backup SAP DB 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 278.

Figure 3-9 **Scheduling the Backup Specification**

Schedule Backup

Specify the desired backup time, frequency, duration, and type.

Recurring

☐ None
☐ Daily
☒ Weekly
☐ Monthly

Time options

Time: 8 AM hours 00 minutes
☐ Use starting
Month: 2004 April Day: 28

Recurring options

Every 1 week(s) on
☐ Sun ☒ Mon ☒ Tue ☒ Wed ☒ Thu ☒ Fri ☐ Sat

Session options

Backup type: Full
Network load: ☒ High ☐ Medium ☐ Low
Backup protection: Default

OK Cancel Help

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 or the *HP OpenView Storage Data Protector Administrator's Guide* for scheduling details.

NOTE

When creating an SAP DB backup specification, you access the Data Protector Scheduler through the Backup Wizard. See “Creating a Backup Specification” on page 260 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 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 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 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 Utilities

Using this integration, you can run an integrated Data Protector backup of an SAP DB Server from SAP DB utilities.

Prerequisites

- The SAP DB Server must be configured for use with this integration. See “Configuring the Integration” on page 259 for information on how to configure an SAP DB Server for use with this integration.
- A Data Protector SAP DB integration backup specification must be created. See “Configuring an SAP DB Backup” on page 260 for information on how to create a Data Protector SAP DB integration backup specification.

In order to run an interactive backup using SAP DB utilities, the following must be done:

- SAP DB media must be created
- the SAP DB `bsi_env` file must be created

Follow the procedure on the next few pages to run a backup from SAP DB utilities using an existing Data Protector SAP DB 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 user created or identified as described in “Configuring Users” on page 259

`<location>` is the location of the `bsi_env` file

`<media_group_name>` is the name of the SAP DB media group

`<medium_name>` is the name of the SAP DB medium

<pipe_name> is the name of the SAP DB pipe

<medium_type> is the type of the SAP DB medium

1. On the SAP DB 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 259) 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 database manager as the SAP DB user created or identified as described in “Configuring Users” on page 259. On the SAP DB Server, execute the following command to login:

```
dbmcli -d <inst_name> -u <username>,<password>
```

3. In the SAP DB database manager, 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 media in an SAP DB media group. Execute the following command for every medium to be created:

```
medium_put <media_group_name>/<medium_name> <pipe_name>  
<medium_type> <backup_type>
```

Where <backup_type> can be one of the following:

- DATA for full backup
- PAGES for differential backup
- LOG for log backup

IMPORTANT

When creating SAP DB media for the purpose of a Data Protector backup and restore, the media group name must begin with the “BACK” string. 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 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 Backup and Restore” on page 305.

Restoring an SAP DB Database

An SAP DB database can be restored using the Data Protector GUI or CLI. An integrated restore is performed.

An SAP DB 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 289.
- The Data Protector CLI: see “Restoring Using the Data Protector CLI” on page 292.
- The SAP DB utilities: see “Restoring Using SAP DB Utilities” on page 294.

When performing an SAP DB migration using any of the above methods, some additional tasks must first be done in order to prepare the SAP DB Server or instance. These tasks are described in “SAP DB Migration Prerequisites” on page 289.

Restore and Recovery Overview

This section provides an overview of restore and recovery process with regard to Data Protector restore and recovery options selection. For a detailed description of these options, see “SAP DB Restore and Recovery Options” on page 298.

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

Depending on the type of restore and on the selected restore and recovery options, the SAP DB database can be switched to the following modes after the restore:

- If the Data Protector Recovery option is selected, the database is switched to the Online mode after the restore.
- If the Data Protector Recovery option is *not* selected and archive logs have not been restored (if restore from a full or diff backup session is performed), the database remains in the Admin mode after the restore.

- If the Data Protector Recovery option is *not* selected and archive logs have been restored, the database is, if the restored archive logs allow it, switched to the Online mode. If the database, however, cannot be switched to the Online mode (because the restored archive logs do not allow it), it remains in the Admin mode.

IMPORTANT

There are several scenarios, depending on the backup option `Keep archive logs` and the recovery option `Use existing archive logs`, in which a gap of transactions between the sequence of redo logs on the SAP DB Server and the restored volumes can occur. When performing recovery (when the database is switched to the Online mode), SAP DB always checks whether such a gap exists, regardless of the point in time selected for recovery. If such a gap exists, the recovery is not performed and the database remains in the Admin mode, unless the existing redo logs are manually deleted before starting the restore.

If a full or diff backup session is restored, only the data (no archive logs) from the selected backup session is restored. The data on the SAP DB Server is overwritten.

If a trans backup session is restored, only the archive logs (no data) from the selected backup session are restored.

During the restore, the redo logs that existed on the SAP DB Server before the restore are not deleted during the restore. Note that during the SAP DB migration, the redo logs that existed on the SAP DB Server before the restore are deleted during the restore.

When restoring, the existing redo logs on the SAP DB Server can be, depending on the Data Protector `Use existing archive logs` option selection (it can be selected only if the `Recovery` option is selected), handled as follows:

- If the `Use existing archive logs` option is selected, the existing archive logs on the SAP DB Server are applied to the redo logs.

When a transactional backup session is selected for restore, or when it is a part of the needed restore chain, and the `Use existing archive logs` option is selected at the same time, the archive logs from Data Protector media are applied to redo logs. Thereafter, the archive logs on the SAP DB Server are applied to redo logs.

- If the Use existing archive logs option is not selected, the backed up archive logs on backup media are applied to the redo logs (if trans backup session is restored), or the redo logs are left intact together with the existing archive logs on the SAP DB Server (if full or diff backup session is restored).

NOTE

The Use existing archive logs option is disabled in case of SAP DB migration, thus allowing only for the restore of redo logs from the backed up archive logs on backup media (if trans backup session is restored).

Figure 3-11 SAP DB Restore Process

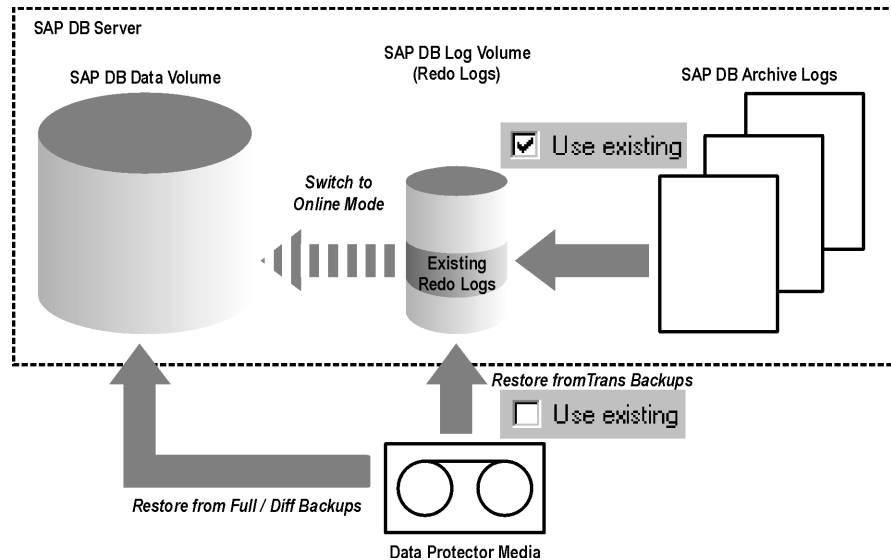
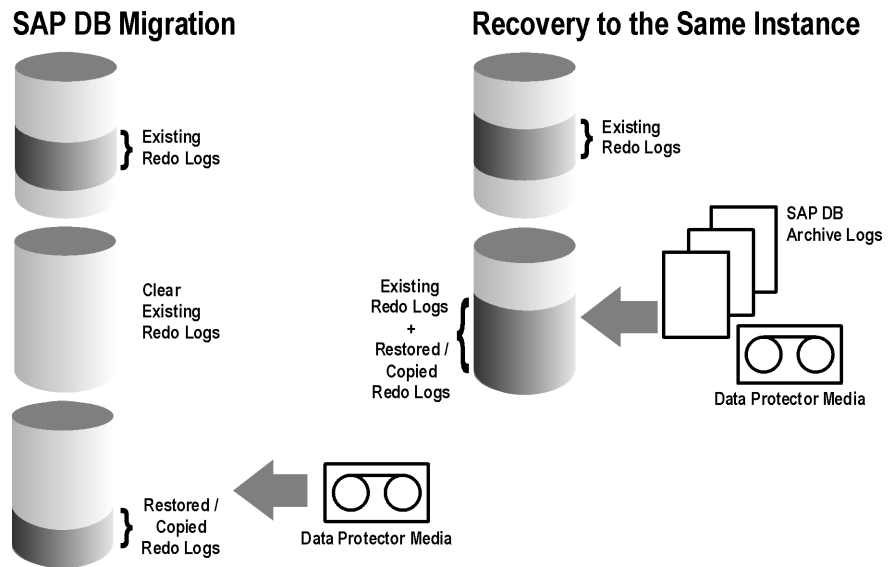


Figure 3-12

SAP DB Archive Logs Restore Process—Redo Logs Details



If you select a differential or a transactional backup session to be restored, you can set the integration to:

- Perform a full database restore. In this case, the integration automatically determines the chain of needed full, differential or transactional backup sessions when performing the restore. After the restore has finished, the database is, if the `Recovery` option is selected, switched to the `Online` mode.
- Restore only the selected differential or the selected transactional backup session. If the database is consistent after such a restore and if the `Recovery` option is selected, it is switched to the `Online` mode. Otherwise, the database is left in the `Admin` mode.

Restoring only the selected trans or diff backup session is useful if the database remains offline or in the Admin mode after a restore from full backup session, which is then followed by a restore from diff or trans backup session.

NOTE

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

SAP DB Migration Prerequisites

The integration supports SAP DB migration, meaning that an SAP DB instance can be restored to some other SAP DB Server or instance than the original.

If the SAP DB Server has not yet been configured for the Data Protector SAP DB 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 migration:

- Install the Data Protector SAP DB integration on the SAP DB Server to which you want to migrate the backed up SAP DB instance. When this is done, add the SAP DB Server in the Data Protector cell.
- Identify or create an OS user under whose account the SAP DB is running and add it to the Data Protector admin group. For information on how to do this, see “Configuring Users” on page 259.
- When performing an SAP DB 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 Instance in Data Protector” on page 267.

When performing an SAP DB 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 migration, some additional tasks must first be done in order to prepare the SAP DB Server or instance. These tasks are described in “SAP DB Migration Prerequisites” on page 289.

To restore your data, proceed as follows in the HP OpenView Storage Data Protector Manager:

1. In the Context List, select Restore.

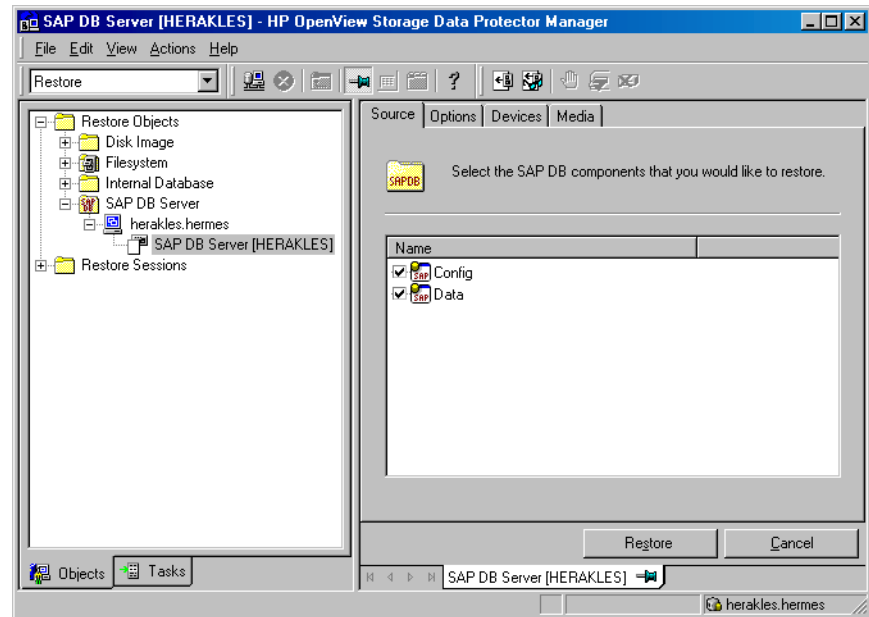
Expand Restore Objects, then SAP DB Server, and then the SAP DB Database Server which you want to restore.

Select the SAP DB instance. A list of backed up objects is displayed in the Results Area.

2. Select the SAP DB objects you want to restore.

Figure 3-13

Selecting SAP DB Objects for a Restore

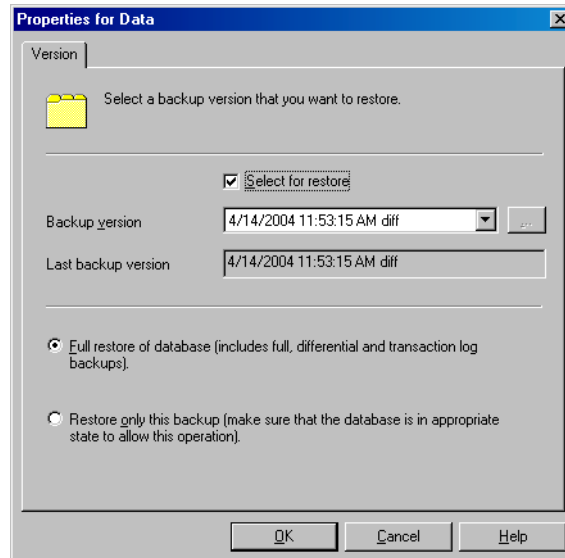


IMPORTANT

To restore SAP DB 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 Restore and Recovery Options” on page 298 for more information on SAP DB 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 migration, some additional tasks must first be done in order to prepare the SAP DB Server or instance. These tasks are described in the “SAP DB Migration Prerequisites” on page 289.

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 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 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 restore.

The `-barhost` option sets the name of the SAP DB Server that was backed up.

The `-instance` option sets the name of the SAP DB instance that was backed up.

The `-session` option selects the backup session to be restored. If this option is not specified, the last backup session is restored, regardless of the `-endlogs` or the `-time` option selection.

When restoring objects that have copies do not use the copy session ID, but the object's backup ID, which equals the object's backup session ID.

The `-nochain` option instructs the integration to restore only the selected or last backup session; the integration does not restore the whole restore chain of full, differential, and transactional backups.

For descriptions of all other options, see “SAP DB Restore and Recovery Options” on page 298. Refer also to the `omnir` man page.

Example

To restore an instance named “inst1” (together with configuration), backed up on an SAP DB Server named “srv1.company.com” from the last backup session and then perform a recovery until the end of logs, enter the following command:

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 Utilities

Using this integration, it is also possible to run an integrated Data Protector restore of an SAP DB Server from SAP DB utilities.

To perform a restore to an existing SAP DB Server instance, see “SAP DB Restore and Recovery” on page 294.

To migrate an SAP DB instance, see “SAP DB Migration” on page 298.

Finding Information Needed for Restore

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

Execute the following commands:

- `omnidb -sapdb`
to get a list of SAP DB objects.
- `omnidb -sapdb <object_name>`
to get details on a specific object, including the SessionID.

SAP DB Restore and Recovery

Follow the procedure on the next few pages to restore and recover a database using SAP DB utilities from existing Data Protector SAP DB backup session(s). In the procedure, the following conventions are used:

`<inst_name>` is the name of the instance to be restored

`<username>`, `<password>` is the connection string for the SAP DB user created or identified as described in “Configuring Users” on page 259

`<location>` is the location of the `bsi_env` file

`<media_group_name>` is the name of the SAP DB media group

`<medium_name>` is the name of the SAP DB medium

`<pipe_name>` is the name of the SAP DB pipe

`<medium_type>` is the type of the SAP DB medium

<SessionID> is the Data Protector session ID of the session to be restored

Restore

1. Skip this step if the bsi_env file is already present and configured on the SAP DB Server.

On the SAP DB Server create the bsi_env file in a directory of your choice. It must contain the following lines:

Windows

```
BACKINT <Data_Protector_home>\bin\sapdb_backint.exe
INPUT <Data_Protector_home>\tmp\<inst_name>.bsi_in
OUTPUT <Data_Protector_home>\tmp\<inst_name>.bsi_out
ERROROUTPUT
<Data_Protector_home>\tmp\<inst_name>.bsi_err
TIMEOUT_SUCCESS 60
TIMEOUT_FAILURE 30
```

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 database manager as the SAP DB user created or identified as described in “Configuring Users” on page 259. On the SAP DB Server, execute the following command to login:

```
dbmcli -d <inst_name> -u <username>,<password>
```

3. In the SAP DB database manager, switch the database to the Admin mode by executing the following command:

```
db_admin
```

4. Skip this step if the location of the bsi_env file is already registered on the SAP DB Server.

Register the location of the bsi_env file as follows:

Windows

```
dbm_configset -raw BSI_ENV <location>\<inst_name>.bsi_env
```

UNIX

```
dbm_configset -raw BSI_ENV <location>/<inst_name>.bsi_env
```

5. Skip this step if the SAP DB media and pipes to be used with Data Protector are already existing on the SAP DB Server.

Note that to restore a Data Protector SAP DB backup session, the number of SAP DB media and pipes required equals the parallelism value used during the backup session.

Create SAP DB media in an SAP DB media group. Execute the following command for every medium to be created:

```
medium_put <media_group_name>/<medium_name> <pipe_name>  
<media_type> <backup_type>
```

Where <backup_type> can be one of the following:

- DATA for full backup
- PAGES for differential (diff) backup
- LOG for transactional (trans) backup

IMPORTANT

When creating SAP DB media and pipes for the purpose of a Data Protector backup and restore, the media group name must begin with the “BACK” string. The commands below create two media and two pipes (parallelism = 2) in a media group:

Windows

```
medium_put BACKDP-Data[2]/1 \
\\.\Pipe\<inst_name>.BACKDP_Data[2].1 PIPE DATA

medium_put BACKDP-Data[2]/2 \
\\.\Pipe\<inst_name>.BACKDP_Data[2].2 PIPE DATA
```

UNIX

```
medium_put BACKDP-Data[2]/1 \
/var/opt/omni/tmp/<inst_name>.BACKDP_Data[2].1 PIPE
DATA

medium_put BACKDP-Data[2]/2 \
/var/opt/omni/tmp/<inst_name>.BACKDP_Data[2].2 PIPE
DATA
```

6. Start the SAP DB utility session by executing the following command:

```
util_connect
```

7. Start the restore from a Data Protector backup session by executing the following command:

```
recover_start <media_group_name> <backup_type> EBID
"<inst_name> <SessionID>:1 <pipe_name1>,<inst_name>
<SessionID>:2 <pipe_name2> [, ...]"
```

Windows

```
recover_start BACKDP-Data[2] DATA EBID "<inst_name>
<SessionID>:1 \\.\Pipe\<inst_name>.BACKDP-Data[2].1,TEST
<SessionID>:2 \\.\Pipe\<inst_name>.BACKDP-Data[2].2"
```

UNIX

```
recover_start BACKDP-Data[2] DATA EBID "<inst_name>
<SessionID>:1
/var/opt/omni/tmp/<inst_name>.BACKDP-Data[2].1,<inst_name>
> <SessionID>:2
/var/opt/omni/tmp/<inst_name>.BACKDP-Data[2].2"
```

Repeat this step for every session in the required chain of backup sessions.

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 database manager:

```
db_online
```

- b. To recover the database until the specified point in time, execute the following command in the SAP DB database manager:

```
db_warm -f -u <yyyymmdd> <hhmmss>
```

Where <yyyymmdd> and <hhmmss> parameters set the time for the last redo log to be applied.

SAP DB Migration

When performing an SAP DB migration, some additional tasks must first be done in order to prepare the SAP DB Server or instance. These tasks are described in the “SAP DB Migration Prerequisites” on page 289.

Follow the procedure in the section “SAP DB Restore and Recovery” on page 294 to migrate the SAP DB database using SAP DB utilities from existing Data Protector SAP DB backup session(s). When following the mentioned procedure, *before* executing the `recover_start` command, delete the existing redo logs on the SAP DB Server by executing the following command in the SAP DB database manager:

```
util_execute clear log
```

SAP DB Restore and Recovery Options

SAP DB 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 object has been selected.

Figure 3-15 SAP DB Restore and Recovery Options

The following are SAP DB specific backup options:

Migration Options To restore selected SAP DB object to the same SAP DB Server and instance, leave the migration options as they are. Use the migration options only in case of SAP DB migration (when restoring to some other SAP DB Server or to some other instance than those that were backed up).

The following are descriptions of the migration options. First the GUI option is given, followed by a slash (/), CLI equivalent, and then description.

Restore to client / -destination *<ClientName>* When using the GUI, in the drop-down list, select an SAP DB Server to which you want to restore the database.

When using the CLI, specify the -destination option and the name of the SAP DB Server as the *<ClientName>* argument.

The selected SAP DB Server must be a part of the Data Protector cell and must have the Data Protector SAP DB Integration software component installed.

Restore to instance / -newinstance

<DestinationInstanceName> When using the GUI, you can either:

- Select an instance in the Restore to instance drop-down list. The drop-down list shows only the instances that are already configured for use with this integration. See “Configuring the Integration” on page 259 for information on how to configure an SAP DB Server for use with this integration.
- Enter the name of an existing instance, not yet configured for use with this integration. In this case, click on the Settings button to configure the specified instance.

When using the CLI, the instance specified as the *<DestinationInstanceName>* argument to the -newinstance option must already be configured for use with this integration. See “Configuring the Integration” on page 259 for information on how to configure an SAP DB Server for use with this integration.

User name and **User group** / N/A On UNIX, you can change the user name and the group name for the OS system user, under whose account the SAP DB application is running on the SAP DB Server (for example, the sapdb user in the sapsys group). By default, the user that started the Data Protector GUI is set for this option.

When using the CLI, it is not possible to change the user name and the group name. The same user as used during the backup session is used.

Settings / N/A Click this button if the instance you are restoring to is not yet configured for use with this integration. See step 4. on page 268 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 Instance Using the CLI” on page 269.

Recovery Options Use the recovery options to recover the database by applying the redo logs until the latest version or until the specified date and time.

IMPORTANT

There are several scenarios, depending on the backup option `Keep archive logs` and the recovery option `Use existing archive logs`, in which a gap of transactions between the sequence of redo logs on the SAP DB Server and the restored volumes can occur. When performing recovery (when the database is switched to the `Online` mode), SAP DB always checks whether such a gap exists, regardless of the point in time selected for recovery. If such a gap exists, the recovery is not performed and the database remains in the `Admin` mode, unless the existing redo logs are manually deleted before starting the restore.

The following are descriptions of the recovery options. First the GUI option is given, followed by a slash (/), CLI equivalent, and then description.

Recovery / `-recover` When this option is selected, the database is recovered after the restore (it is switched to `Online` mode) by applying the redo logs until the latest version (if the `Latest version` option is selected) or until the specified date and time (if the `Until date` option is selected).

IMPORTANT

When using this option, make sure that the backup session selected in the `Properties for Data` dialog box (when using GUI) or by the `-session` option (when using CLI) will restore enough data for the integration to apply the redo logs until the latest version or until the specified date and time. For information on how to access the `Properties for Data` dialog box, refer to step 3 on page 290. For information on the `-session` option, refer to “Restoring Using the Data Protector CLI” on page 292.

When this option is not selected, all other recovery options are disabled and the following happens after the restore:

- If archive logs are not restored (if restore from a full backup session is performed), the database remains in the Admin mode after the restore.
- If archive logs are restored, the database is, if the restored archive logs allow it, switched to the Online mode. If the database, however, cannot be switched to the Online mode (because the restored archive logs do not allow it), it remains in the Admin mode.

Latest version / `-endlogs` Select this option to recover the database until the last log.

When using the CLI, this is the default option.

Until date / `-time: <YYYY-MM-DD.hh.mm.ss>` When using the GUI, select this option to recover the database until the point you select in the `Until date` drop-down menu.

When using the CLI, specify the `-time:` option if you want to recover the database until the point specified by the `<YYYY-MM-DD.hh.mm.ss>` argument.

NOTE

The selected time is the system time on the system running the Data Protector GUI or CLI. If the system to be recovered is not in the same time zone as the system running the Data Protector GUI or CLI, the point of recovery is adjusted to the local time setting on the system to be restored.

Use existing archive logs / `-from_disk` Select this option to copy the existing archive logs on the SAP DB Server to SAP DB Server redo logs.

If this option is not selected, the backed up archive logs on backup media are applied to the redo logs (if trans backup session is restored), or the redo logs are left intact together with the existing archive logs on the SAP DB Server (if full or diff backup session is restored).

When a transactional backup session is selected for restore or when it is a part of the needed restore chain, and the Use existing archive logs option is selected at the same time, the archive logs from Data Protector media are applied to the redo logs. Thereafter, the archive logs on the SAP DB Server are applied to redo logs.

NOTE

The Use existing archive logs option is disabled in case of SAP DB migration, thus allowing only for the restore of redo logs from the backed up archive logs on backup media (if trans backup session is restored).

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, refer to the “Restoring Under Another Device” section of the *HP OpenView Storage Data Protector Administrator’s Guide* 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 Disaster Recovery chapter in the *HP OpenView Storage Data Protector Administrator’s 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 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

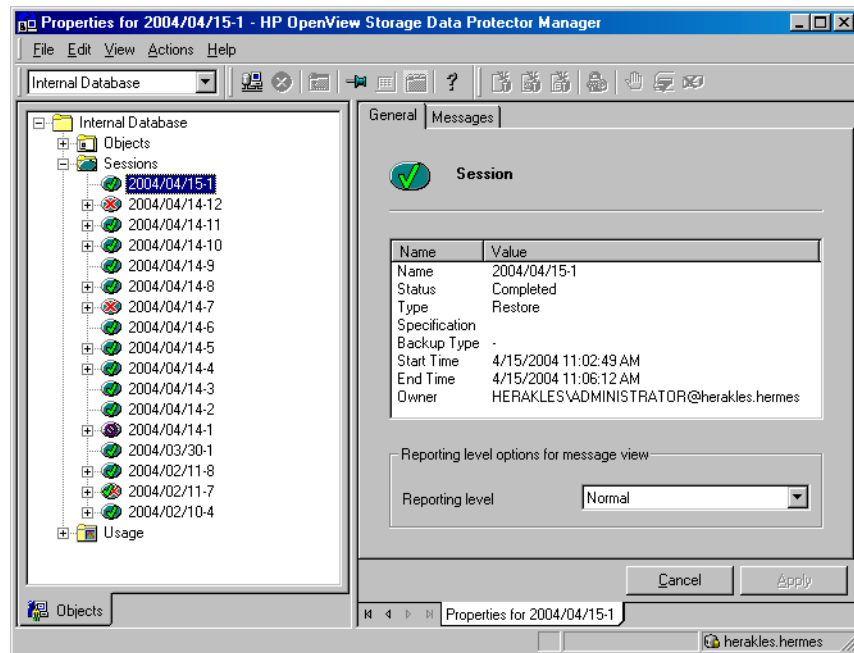
To view a previous session using the Data Protector GUI, proceed as follows:

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

Before you start troubleshooting the Data Protector SAP DB integration, check the following:

1. Ensure that the latest official Data Protector patches are installed. Refer to “Verifying Which Data Protector Patches Are Installed” in the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

Refer to the *HP OpenView Storage Data Protector Software Release Notes* or

http://www.openview.hp.com/products/datapro/spec_0001.html for an up-to-date list of supported versions, platforms, and other information.

2. Refer to the *HP OpenView Storage Data Protector Software Release Notes* for a description of general Data Protector limitations, supported versions, problems and workarounds, and a list of related Data Protector patches.

The following sections provide some checking procedures you should perform before you call Data Protector support. In this way you may either resolve the problem yourself or identify the area where the difficulties are occurring.

Should you fail when performing a troubleshooting procedure, measures have been proposed to help you work around the problem.

SAP DB Troubleshooting

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 MAXBACKUPDEVS parameter to a value that is greater than or equal to the value of the Data Protector Parallelism option, or reduce the value of the Data Protector Parallelism option.

Problem **An SAP DB instance cannot be started after restore.**

Action Using the SAP DB `db_restartinfo` command, check if the instance can be restarted.

- If the instance cannot be restarted, most probably the existing log volumes do not contain enough data to restart the instance from data volumes. The required differential or transactional backups might not have been restored.
- If the instance can be restarted, check the SAP DB instance kernel error file for errors.

If there was insufficient space for SAP DB logs at some point of time, logs might have been corrupted: delete the logs (using the `dbmcli util_execute clear log` command) or contact SAP DB or Data Protector support.

Problem **A restore from an object copy hangs.**

Action Before restarting the restore:

- Increase the number of Disk Agent buffers for the device used for the restore.
- If all objects of the backup are recorded in the IDB, perform the following steps:
 1. In the Internal Database context of the Data Protector GUI, search for all objects belonging to the same backup. The objects are identified by the same backup ID.
 2. Copy each object in a separate object copy session to a separate device, for example a file library. For each object, use a separate medium with the non-appendable media policy.
 3. Set the highest media location priority for the newly created copies.

Problem	Data Protector reports the following error: Error: SAPDB responded with: Error! Connection failed to node (local) for database CLUSTER: connection refused: x_server not running.
Action	Start the SAP DB x_server. Refer to the SAP DB documentation for information on how to do that.
Problem	Data Protector reports the following error: Error: SAPDB responded with: -24988,ERR_SQL: sql error 1,database not running
Action	Start the SAP DB instance. Refer to the SAP DB documentation for information on how to do that.
Problem	Data Protector reports the following error: Error: SAPDB responded with: -24988,ERR_SQL: sql error 1,utility session is already in use
Action	Some other user is connected to the SAP DB instance and is performing administrative tasks (utility session). Such SAP DB tasks are of the "Utility" type and can be displayed using the dbmcli show task command. Finish these tasks.
Problem	Data Protector reports the following error: Error: SAPDB responded with: -24950,ERR_USRFAIL: user authorization failed
Action	Reconfigure the SAP DB instance as described in the section “Modifying the Configuration of an SAP DB Instance in Data Protector” on page 267.

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 Configuration File” on page 255.

You can also set the TimeoutSuccess environment variable using the Data Protector GUI. Select the backup specification in the Scoping Pane, then right-click the SAP DB instance object in the Results Pane under the Source tab and select the Set Environment Variables from the pop-up menu.

SAP DB Cluster-Related Troubleshooting

In a cluster environment, the environment variable OB2BARHOSTNAME must be defined as the virtual hostname before performing some procedures run from the command line (on the client). The OB2BARHOSTNAME variable is set as follows:

UNIX

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

Windows

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

Glossary

access rights

See **user rights**.

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 one (or more) archived log destination(s). This copy is the Archived Redo Log. The presence or absence of an Archived Redo Log is determined by the mode that the database is using:

- **ARCHIVELOG** - The filled online redo log files are archived before they are reused. The database can be recovered from an instance and disk failure. The “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

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

This relates to a situation where full and incremental backups are performed. Based on the level of the incremental backups used (Incr, Incr 1, Incr 2, and so on), simple or rather complex dependencies of incrementals to previous incrementals can exist. The backup chain are all backups, starting from the full backup plus all the dependent incrementals up to the desired point in time.

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 backup device to perform the backup of the data in a replica.

See also **application system, target volume, and replica.**

backup types

See **incremental backup, differential backup, transaction backup, full backup and delta backup.**

backup 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 Continuity 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 Process (*EMC Symmetrix specific term*)

A protected storage environment solution that has defined specially configured EMC Symmetrix devices as mirrors or Business Continuity Volumes to protect data on EMC Symmetrix standard devices.

See also **BCV.**

BC VA (*HP StorageWorks Virtual Array specific term*)

Business Copy VA allows you to

maintain internal copies of HP StorageWorks Virtual Array LUNs for data backup or data duplication within the same virtual array. The copies (child or Business Copy LUNs) can be used for various purposes, such as backup, data analysis or development. When used for backup purposes, the original (parent) LUNs are connected to the application system and the Business Copy (child) LUNs are connected to the backup system.

See also **HP StorageWorks Virtual Array LUN, application system, and backup system.**

BCV (*EMC Symmetrix specific term*) Business Continuance Volumes, or BCV devices, are dedicated SLDs that are pre-configured in the ICDA on which the business continuation operation runs. BCV devices are assigned separate SCSI addresses, differing from the addresses used by the SLDs they mirror. The BCV devices are used as splittable mirrors of the primary EMC Symmetrix SLDs that need to be protected.

See also **BC** and **BC Process.**

Boolean operators

The Boolean operators for the full text search functionality of the online Help system are AND, OR, NOT, and NEAR. Used when searching, they enable you to define your query precisely by creating a relationship between search terms. If no operator is specified in a

multi-word search, AND is used by default. For example, the query manual disaster recovery is equivalent to manual AND disaster AND recovery.

boot volume/disk/partition

A volume/disk/partition with files required for the initial step of the boot process. Microsoft terminology defines the boot volume/disk/partition as a volume/disk/partition containing the operating system files.

BRARCHIVE (*SAP R/3 specific term*)

An SAP R/3 backup tool that allows you to archive redo log files. BRARCHIVE also saves all the logs and profiles of the archiving process.

See also **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**.

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 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 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 OnLine Server

(Informix specific term)

Windows CMD script that is created in INFORMIXDIR when Informix OnLine Server is configured. The CMD script is a set of system commands that export environment variables for OnLine 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, restore, and management tasks.

Command View (CV) EVA *(HP*

StorageWorks EVA specific term)

The user interface that allows 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, creating snapclones and snapshots of virtual disks. The Command View EVA software runs on the HP OpenView

Glossary

Storage Management Appliance, and is accessed by a Web browser.

See also **HP StorageWorks EVA Agent (legacy)** and **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.

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 Session Manager process controls the object copy session 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 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, the Oracle Server.

database parallelism

More than one database is backed up at a time if the number of available devices allows you to perform backups in parallel.

database server

A computer with a large database stored on it, such as the SAP R/3 or Microsoft SQL database. A server has a database that can be accessed by clients.

Dboject (*Informix specific term*)

An Informix physical database object. It can be a blob space, db space, 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 configuration of IP addresses and related information.

differential backup

An incremental backup (incr) based on any previous Data Protector backup (full or any incremental), which must still be protected.

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.

Directory Store (DS) (*Microsoft Exchange specific term*)

A part of the Microsoft Exchange Server directory. The Microsoft Exchange Server directory contains objects used by Microsoft Exchange applications in order to find and access services, mailboxes, recipients, public folders, and other addressable objects within the messaging system.

See also **Information Store (MDB)**.

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

Glossary

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 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 specific term)*

An Informix configuration file that resides in the <INFORMIXDIR>\etc directory (on HP-UX) or <INFORMIXDIR>/etc directory (on Windows) and is called ixbar.<server_id>, where <INFORMIXDIR> is the OnLine 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.

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 2000/2003 specific term)

A database technology used as a storage system for information exchange in Microsoft Exchange Server 2000/2003.

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.

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**.

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.

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

Glossary

<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 Agent (legacy)

A Data Protector software module that executes all tasks required for the HP StorageWorks Enterprise Virtual Array integration operating on HP StorageWorks EVA with Command View (CV) EVA software v3.1 or lower, and the EVA VCS firmware v3.01x or lower.

See also **Command View (CV) EVA** and **HP StorageWorks EVA SMI-S Agent**.

HP StorageWorks EVA SMI-S Agent

A Data Protector software module that executes all tasks required for the HP StorageWorks Enterprise Virtual Array integration operating on HP

StorageWorks EVA with Command View (CV) EVA software starting with v3.2. 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**, **HP StorageWorks SMI-S EVA provider**, and **HP StorageWorks EVA Agent (legacy)**.

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 which devices and libraries are configured.

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, allowing selective backup of only files that have changed since the last incremental backup.

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.

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 2000/2003 specific term*)

The Microsoft Exchange Server 2000/2003 service that is responsible for storage management. Information Store in Microsoft Exchange Server 2000/2003 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**.

Information Store (*Microsoft Exchange Server 5.5 specific term*)

This is the default message store provider for the Microsoft Exchange Server 5.5. Information Store consists of the following stores:

- public information store
- private information store
- personal folder store
- offline information store.

The public information store contains public folders and messages that can be shared among multiple users and applications. A single public store is shared by all users within an Exchange Server 5.5 organization, even if multiple Exchange Servers are used. The private information store consists of mail boxes that can belong to users or to applications. The mail boxes reside on the server running the Exchange Server 5.5.

See also **Directory Store (DS)**.

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 2000/2003 specific term)*

The Microsoft Exchange Server 2000/2003 service that provides encryption functionality for enhanced security. *See also Information Store and Site Replication Service.*

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 UNIX specific term*)

A script provided by ON-Bar that you can use to start backing up logical-log files when OnLine Server issues a log-full event alarm. The Informix ALARMPROGRAM configuration parameter defaults to the `<INFORMIXDIR>/etc/log_full.sh`, where `<INFORMIXDIR>` is the OnLine 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 or object copying. 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 been granted Oracle SYSDBA or SYSOPER rights.
- <password> is a string used for data security and known only to its owner. Passwords are entered to connect to an operating system or software application. The password has to 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 has to be the owner of the Oracle Recovery (Oracle) 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 2000/2003 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**.

MFS

The Migrating File System enables a standard JFS filesystem with migration capabilities (on HP-UX 11.00). The MFS is accessed via a standard filesystem interface (DMAPI), it is mounted to a directory the same way as any HP-UX filesystem. In an MFS, only the superblock, the inode and the 'extended attribute' information remain permanently on the hard disk and are never migrated. *See also* **VBFS**.

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 7.0/2000

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 MoM allows 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

An IDB file 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, on a separate physical disk from other IDB directories, and, additionally, to make a copy of the file and locate it where you want.

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 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**

OmniStorage

Software providing transparent migration of less frequently used data to the optical library while keeping more frequently used data on the hard disk. HP OmniStorage runs on HP-UX systems.

On-Bar (*Informix specific term*)

A backup and restore system for OnLine Server. ON-Bar enables you to create a copy of your OnLine Server data and later restore the data. The ON-Bar backup and restore system involves the following components:

- onbar utility
- Data Protector, as the backup solution
- XBSA interface
- ON-Bar catalog tables, which are used to back up dbobjects and track instances of dbobjects through multiple backups.

onbar utility (*Informix specific term*)

The Informix utility that communicates backup and restore requests to OnLine Server. The utility uses XBSA to exchange control data and back up and restore data with Data Protector.

ONCONFIG (*Informix specific term*)

An environment variable that specifies the name of the active ONCONFIG configuration file. If the ONCONFIG environment variable is not present, OnLine uses the configuration values from the file <INFORMIXDIR>/etc/onconfig (on HP-UX) or <INFORMIXDIR>\etc\onconfig (on Windows).

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**.

OnLine Server (*Informix specific term*)

Refers to INFORMIX-OnLine Dynamic Server.

OpC

See **OVO**.

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)**.

Private Information Store (Microsoft Exchange Server 5.5 specific term)

A part of the Information Store that maintains information in user mailboxes. A mailbox store consists of a binary rich-text .edb file.

protection

See **data protection** and also **catalog protection**.

public folder store (Microsoft

Exchange Server 2000/2003 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.

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 (RCU) (*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 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, 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 change 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 2000/2003 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, object copy, restore, 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 specific term*)

An Informix connectivity-information file 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 2000/2003 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.

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 7.0/2000 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.

VBFS (*OmniStorage specific term*)

A Very Big File System is an extension of the standard HP-UX file system on HP-UX 9.x. It is mounted to a directory the same way as any HP-UX file system. In a VBFS, only the superblock, the inode and the 'extended attribute'

information remain permanently on the hard disk and are never migrated.
See also **MFS**.

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 7.0/2000 specific term)

This is a SQL Server 7.0/2000 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 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 specific term)*

The onbar utility and Data Protector communicate with each other through the X/Open Backup Specification Services 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, 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.**

A

- aborting
 - SAP R/3 backup session, 195
- advantages
 - Oracle integration, 4
 - SAP DB integration, 250
 - SAP R/3 integration, 143
- architecture
 - SAP R/3 integration, 149

B

- backing up Oracle, 58–71
 - backup options, 48
 - backup templates, configuring, 43
 - examples, using RMAN, 68
 - offline, 58
 - online, 58
 - recovery catalog, 61
 - scheduled backup, 62
 - starting backup, 58–71
 - starting backup, using RMAN, 65
 - starting interactive backup, using Data Protector CLI, 65
 - starting interactive backup, using Data Protector GUI, 63
- backing up SAP DB, 276–284
 - backup concepts, scheme, 254
 - backup flow, 252
 - backup methods, 276
 - backup modes, 248, 249
 - backup options, 265
 - backup specification, creating, 260
 - differential backup, 248
 - full backup, 248
 - scheduled backup, 276
 - starting backup, 276–284
 - starting backup, using Data Protector CLI, 280
 - starting backup, using Data Protector GUI, 279
 - starting backup, using SAP DB utilities, 281
 - starting interactive backup, 279
 - transactional backup, 248
- backing up SAP R/3, 194–200
 - aborting backup session, 195
 - backup concepts, scheme, 148
 - backup flow, backint mode, 149
 - backup flow, RMAN mode, 152
 - backup methods, 194
 - backup modes, 194
 - backup options, 183
 - backup problems, on UNIX, 227
 - backup problems, on Windows, 216
 - backup specification, creating, 177
 - backup templates, configuring, 176
 - backup utilities, 146
 - incremental backup, 194
 - scheduled backup, 195
 - starting backup, 194–200
 - starting backup, using BRBACKUP, 199
 - starting backup, using sapdba, 199
 - starting interactive backup, using Data Protector CLI, 199
 - starting interactive backup, using Data Protector GUI, 197
 - using RMAN, 188
- backint mode
 - SAP R/3 integration, backup flow, 149
 - SAP R/3 integration, restore flow, 153
- backup concepts
 - SAP DB integration, 254
 - SAP R/3 integration, 148
- backup flow
 - Oracle, 10–11
 - SAP DB integration, 252
 - SAP R/3 integration, backint mode, 149
 - SAP R/3 integration, RMAN mode, 152
- backup methods
 - SAP DB integration, 276
 - SAP R/3 integration, 194
- backup modes
 - SAP DB integration, 248, 249
 - SAP R/3 integration, 194
- backup options
 - Oracle integration, 48
 - SAP DB integration, 265
 - SAP R/3 integration, 183
- backup problems
 - Oracle integration, on UNIX systems, 128
 - Oracle integration, on Windows systems, 132
 - SAP R/3 integration, on UNIX, 227
 - SAP R/3 integration, on Windows, 216
- backup specification
 - Oracle integration, creating, 44
 - Oracle integration, ownership, 29
 - Oracle integration, scheduling, 62

- SAP DB integration, creating, 260
- SAP DB integration, scheduling, 276
- SAP R/3 integration, creating, 177
- SAP R/3 integration, creating for manual balancing, 191
- SAP R/3 integration, ownership, 162
- SAP R/3 integration, scheduling, 195
- backup templates
 - Oracle integration, 43
 - SAP R/3 integration, 176
- backup types
 - Oracle integration, 3
- backup utilities
 - SAP R/3 integration, 146
- BRARCHIVE, 146, 184
- BRBACKUP, 146, 183, 199
- BRRESTORE, 147, 154, 204

C

- checking
 - Oracle client configuration, 38, 42
 - Oracle configuration, 37
 - SAP DB configuration, 271
 - SAP R/3 configuration, 172
- concepts
 - SAP DB integration, 252–254
 - SAP R/3 integration, 146–154
- configuration file
 - SAP DB integration, 255–258
 - SAP DB integration, parameters, 256
 - SAP R/3 integration, 155–161
 - SAP R/3 integration, parameters, 158
- configuration files
 - Oracle integration, 14–16
 - Oracle integration, parameters, 16
- configuration parameters
 - Oracle integration, 40
- configuration problems
 - Oracle integration, 123, 126
 - SAP R/3 integration, on UNIX, 225
 - SAP R/3 integration, on Windows, 213
- configuring Oracle, 21–42
 - backup specification, 44
 - backup templates, 43
 - checking configuration, 37
 - client, 38, 42
 - server, 30
 - server, on HP-UX, 40
 - server, on SUN Solaris, 40
 - server, using CLI, 39

- server, using Data Protector GUI, 31
- users, 29
- configuring SAP DB, 259–275
 - backup, 260
 - backup specification, 260
 - checking configuration, 271
 - overview, 259
 - users, 259
- configuring SAP R/3, 162–175
 - backup specification, 177
 - backup templates, 176
 - checking configuration, 172
 - configuration problems, on UNIX, 225
 - configuration problems, on Windows, 213
 - Database Server, 164
 - Inet user account, 166
 - users, 162
- control file
 - Oracle integration, restore, 78
- conventions, ix
- creating
 - Oracle backup specification, 44
 - SAP DB backup specification, 260
 - SAP R/3 backup specification, 177
 - SAP R/3 backup specification for manual balancing, 191
 - SAP R/3 parameter file, 187

D

- data objects
 - SAP R/3 integration, 146
- Data Protector Database Library, 10
 - linking with Oracle, 22
- database items restore, 72
- database recovery
 - Oracle integration, options, 86
- differential backup
 - SAP DB integration, 248
- disaster recovery
 - Oracle integration, 105
 - restoring databases after, 74
 - SAP DB integration, 303
 - SAP R/3 integration, 205

E

- examples
 - Oracle integration, backing up using RMAN, 68
 - Oracle integration, restore using RMAN, 91

SAP R/3 integration, restore, 234–243

F

finding

- Oracle integration, users, 29
- SAP R/3 integration, Oracle users, 163
- SAP R/3 integration, SAP R/3 users, 163

full backup

- SAP DB integration, 248

I

incremental backup

- Oracle integration, 62
- SAP R/3 integration, 194

Inet user account

- SAP R/3 integration, configuring, 166
- SAP R/3 integration, setting, 199

instance objects

- SAP DB integration, 248

integration concepts

- Oracle, 8

interactive backup

- Oracle integration, 63
- SAP DB integration, starting, 279
- SAP R/3 integration, starting, 197

L

limitations

- Oracle integration, 7
- SAP DB integration, 247
- SAP R/3 integration, 145

linking

- Oracle with the Database Library, 22

M

manual balancing

- SAP R/3 backup specification, creating for
 - manual balancing, 191
- SAP R/3 integration, 190

migration

- SAP DB restore, 249, 289

modifying

- SAP R/3 parameter file, 187

monitoring

- Oracle sessions, 106
- SAP DB sessions, 305
- SAP R/3 sessions, 207

O

Oracle backup, 58–71

- backup concepts, scheme, 12
- backup templates, configuring, 43
- scheduling backup, 62
- starting backup, 58–71
- starting backup, using RMAN, 65
- starting interactive backup, using Data Protector CLI, 65
- starting interactive backup, using Data Protector GUI, 63

Oracle client

- checking configuration, 42

Oracle configuration

- backup specification, 44
- backup templates, 43
- checking, 37
- users, 29

Oracle integration

- backing up, 58–71
- backup types, 3
- concepts, 8
- configuration files, 14–16
- configuration overview, 22
- configuration parameters, 40
- configuring, 21–42
- configuring the Oracle server using Data Protector GUI, 31
- configuring, Oracle server, 30
- disaster recovery, 105
- limitations, 7
- linking with the Database Library, 22
- monitoring sessions, 106
- overview, 3
- prerequisites, 6
- removing, 108
- removing, on HP-UX, 108
- removing, on Solaris and other UNIX systems, 109
- restore flow, 11
- restore methods, 72
- restore types, 3
- restoring, 72–104
- testing, 55
- troubleshooting, 113–139
- viewing sessions, 107

Oracle restore, 72–104

- control file, 78
- database items, 72

- database objects, 80
- disaster recovery, 105
- examples, using RMAN, 91
- preparing database for restore, 91
- recovery catalog database, 75
- recovery catalog, using CLI, 103
- restorable items, 72
- restore options, 86
- tablespaces and datafiles, 85
- using another device, 104
- using Data Protector GUI, 74
- using RMAN, 91
- Oracle RMAN metadata, 111
- Oracle RMAN script, 50
- Oracle troubleshooting, 113–139
- overview
 - Oracle integration, 3
 - SAP DB integration, 248
 - SAP DB restore, 285
 - SAP R/3 integration, 143
- ownership
 - Oracle integration, backup specification, 29
 - SAP R/3 integration, backup specification, 162
- P**
- parallelism
 - SAP DB integration, 250, 252, 253, 266
- parameter file
 - SAP R/3 integration, creating, 187
 - SAP R/3 integration, modifying, 187
- parameters
 - SAP DB integration, configuration file, 256
 - SAP R/3 integration, configuration file, 158
- post-exec command
 - specific Oracle backup option, 50
- pre-exec command
 - specific Oracle backup option, 49
- prerequisites
 - Oracle integration, 6
 - SAP DB integration, 247
 - SAP R/3 integration, 145
- R**
- RAC
 - configuring the Oracle server, on HP-UX, 31
 - configuring the Oracle server, other UNIX systems, 31
 - configuring the Oracle server, using CLI, 40
- recovery
 - Oracle integration, options, 86
- recovery catalog
 - Oracle integration, backup, 61
- recovery catalog database
 - Oracle integration, restore, 75
- Recovery Manager *See* RMAN
- removing
 - Oracle integration, 108
 - Oracle integration, on HP-UX, 108
 - Oracle integration, on Solaris and other UNIX systems, 109
- restore flow
 - SAP DB integration, 253
 - SAP R/3 integration, backint mode, 153
 - SAP R/3 integration, RMAN mode, 154
- restore options
 - SAP DB integration, 298
- restore problems
 - Oracle integration, 134
 - SAP R/3 integration, on UNIX, 230
 - SAP R/3 integration, on Windows, 217
- restore types
 - Oracle integration, 3
- restoring Oracle, 72–104
 - control file, 78
 - database objects, 80
 - disaster recovery, 105
 - methods, 72
 - recovery catalog database, 75
 - recovery catalog, using CLI, 103
 - restore flow, 11
 - tablespaces and datafiles, 85
 - using another device, 104
 - using Data Protector GUI, 74
 - using RMAN, 91
- restoring SAP DB, 285–304
 - disaster recovery, 303
 - migration, 249, 289
 - overview, 285
 - restore flow, 253
 - restore options, 298
 - using another device, 303
 - using CLI, 292
 - using GUI, 289
 - using SAP DB utilities, 294
- restoring SAP R/3, 201–205
 - archive log files, example, 242
 - disaster recovery, 205
 - examples, 234–243

- full database, example, 236
- lost files, example, 241
- partial, example, 240
- preparing database for restore, 234
- restore flow, backint mode, 153
- restore flow, RMAN mode, 154
- restore problems, on UNIX, 230
- restore problems, on Windows, 217
- using another device, 204
- using BRRESTORE, 204
- using CLI, 203
- using GUI, 201
- using sapdba, 204
- RMAN, 65
 - backing up SAP R/3, 188
 - Oracle integration, restoring using, 91
 - SAP R/3 integration, backup flow, 152
 - SAP R/3 integration, restore flow, 154
 - scripts, examples, 68
- running backup *See* starting backup

S

- SAP DB backup, 276–284
 - backup concepts, scheme, 254
 - backup flow, 252
 - backup methods, 276
 - backup modes, 248, 249
 - backup options, 265
 - backup specification, creating, 260
 - differential backup, 248
 - full backup, 248
 - scheduling backup, 276
 - starting backup, 276–284
 - starting backup, using Data Protector CLI, 280
 - starting backup, using Data Protector GUI, 279
 - starting backup, using SAP DB utilities, 281
 - starting interactive backup, 279
 - transactional backup, 248
- SAP DB configuration, 259–275
 - backup, 260
 - backup specification, 260
 - checking, 271
 - overview, 259
 - users, configuring, 259
- SAP DB configuration file, 255–258
- SAP DB integration

- advantages, 250
- backing up SAP DB, 276–284
- backup flow, 252
- concepts, 252–254
- configuration file, 255–258
- configuring, 259–275
- instance objects, 248
- limitations, 247
- monitoring sessions, 305
- overview, 248
- parallelism, 250, 252, 253, 266
- prerequisites, 247
- restore, 285–304
- restore flow, 253
- restore, migration, 249, 289
- restore, overview, 285
- testing, 273
- troubleshooting, 308–311
- util_cmd, 256
- viewing sessions, 306
- SAP DB restore, 285–304
 - disaster recovery, 303
 - migration, 249, 289
 - overview, 285
 - restore flow, 253
 - restore options, 298
 - using another device, 303
 - using CLI, 292
 - using GUI, 289
 - using SAP DB utilities, 294
- SAP DB troubleshooting, 308–311
- SAP DB utilities, 281
 - restore, 294
- SAP R/3 backup, 194–200
 - aborting backup session, 195
 - backup concepts, scheme, 148
 - backup flow, backint mode, 149
 - backup flow, RMAN mode, 152
 - backup methods, 194
 - backup modes, 194
 - backup options, 183
 - backup problems, on UNIX, 227
 - backup problems, on Windows, 216
 - backup specification, creating, 177
 - backup templates, configuring, 176
 - backup utilities, 146
 - incremental backup, 194
 - scheduling backup, 195
 - starting backup, 194–200

- starting backup, using BRBACKUP, 199
- starting backup, using sapdba, 199
- starting interactive backup, using Data Protector CLI, 199
- starting interactive backup, using Data Protector GUI, 197
- using RMAN, 188
- SAP R/3 configuration, 162–175
 - backup specification, 177
 - backup templates, 176
 - checking, 172
 - configuration problems, on UNIX, 225
 - configuration problems, on Windows, 213
 - Database Server, 164
 - Inet user account, 166
 - users, 162
- SAP R/3 configuration file, 155–161
- SAP R/3 integration
 - advantages, 143
 - architecture, 149
 - backing up, 194–200
 - concepts, 146–154
 - configuration file, 155–161
 - configuring, 162–175
 - data objects, 146
 - disaster recovery, 205
 - Inet user account, setting, 199
 - limitations, 145
 - manual balancing, 190
 - monitoring sessions, 207
 - overview, 143
 - parameter file, creating, 187
 - parameter file, modifying, 187
 - prerequisites, 145
 - restoring, 201–205
 - testing, 192
 - troubleshooting, 209–233
 - util_cmd, 158
 - viewing sessions, 208
- SAP R/3 restore, 201–205
 - archive log files, example, 242
 - disaster recovery, 205
 - examples, 234–243
 - full database, example, 236
 - lost files, example, 241
 - partial, example, 240
 - preparing database for restore, 234
 - restore flow, backint mode, 153
 - restore flow, RMAN mode, 154
 - restore problems, on UNIX, 230
 - restore problems, on Windows, 217
 - using another device, 204
 - using BRRESTORE, 204
 - using CLI, 203
 - using GUI, 201
 - using sapdba, 204
- SAP R/3 troubleshooting, 209–233
 - backup problems, on UNIX, 227
 - backup problems, on Windows, 216
 - configuration problems, on UNIX, 225
 - configuration problems, on Windows, 213
 - on UNIX, 220–233
 - on Windows, 210–219
 - restore problems, on UNIX, 230
 - restore problems, on Windows, 217
- sapdba, 147, 199, 204
- scheduling backup
 - Oracle integration, 62
 - SAP DB integration, 276
 - SAP R/3 integration, 195
- setting
 - SAP R/3 integration, Inet user account, 199
- starting backup
 - Oracle integration, 58–71
 - Oracle integration, using Data Protector CLI, 65
 - Oracle integration, using Data Protector GUI, 63
 - Oracle integration, using RMAN, 65
 - SAP DB integration, 276–284
 - SAP DB integration, interactive, 279
 - SAP DB integration, using Data Protector CLI, 280
 - SAP DB integration, using Data Protector GUI, 279
 - SAP DB integration, using SAP DB utilities, 281
 - SAP R/3 integration, 194–200
 - SAP R/3 integration, interactive, 197
 - SAP R/3 integration, using BRBACKUP, 199
 - SAP R/3 integration, using Data Protector CLI, 199
 - SAP R/3 integration, using Data Protector GUI, 197
 - SAP R/3 integration, using sapdba, 199

T

testing

- Oracle integration, 55
- SAP DB integration, 273
- SAP R/3 integration, 192

transactional backup

- SAP DB integration, 248

troubleshooting Oracle, 113–139

troubleshooting SAP DB, 308–311

troubleshooting SAP R/3, 209–233

- backup problems, on UNIX, 227
- backup problems, on Windows, 216
- configuration problems, on UNIX, 225
- configuration problems, on Windows, 213
- on UNIX, 220–233
- on Windows, 210–219
- restore problems, on UNIX, 230
- restore problems, on Windows, 217

typographical conventions, ix

U

users

- Oracle integration, configuring, 29
- Oracle integration, finding, 29
- SAP DB integration, configuring, 259
- SAP R/3 integration, configuring, 162
- SAP R/3 integration, finding Oracle users, 163
- SAP R/3 integration, finding SAP R/3 users, 163

util_cmd

- Oracle integration, 17
- SAP DB integration, 256
- SAP R/3 integration, 158

V

viewing

- Oracle sessions, 107
- SAP DB sessions, 306
- SAP R/3 sessions, 208

