

HP Data Protector

Software Version: 8.10

Integration guide for Oracle and SAP

Document Release Date: November 2016

Software Release Date: November 2016



Legal Notices

Warranty

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.

The information contained herein is subject to change without notice.

Restricted Rights Legend

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.

Copyright Notice

© Copyright 2016 Hewlett-Packard Development Company, L.P.

Trademark Notices

Adobe® is a trademark of Adobe Systems Incorporated.

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

UNIX® is a registered trademark of The Open Group.

Documentation Updates

The title page of this document contains the following identifying information:

- Software Version number, which indicates the software version.
- Document Release Date, which changes each time the document is updated.
- Software Release Date, which indicates the release date of this version of the software.

To check for recent updates or to verify that you are using the most recent edition of a document, go to: <http://h20230.www2.hp.com/selfsolve/manuals>

This site requires that you register for an HP Passport and sign in. To register for an HP Passport ID, go to: <http://h20229.www2.hp.com/passport-registration.html>

Or click the **New users - please register** link on the HP Passport login page.

You will also receive updated or new editions if you subscribe to the appropriate product support service. Contact your HP sales representative for details.

Support

Visit the HP Software Support Online web site at: <http://www.hp.com/go/hpsoftwaresupport>

This web site provides contact information and details about the products, services, and support that HP Software offers.

HP Software online support provides customer self-solve capabilities. It provides a fast and efficient way to access interactive technical support tools needed to manage your business. As a valued support customer, you can benefit by using the support web site to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Download software patches
- Manage support contracts
- Look up HP support contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training

Most of the support areas require that you register as an HP Passport user and sign in. Many also require a support contract. To register for an HP Passport ID, go to:

<http://h20229.www2.hp.com/passport-registration.html>

To find more information about access levels, go to:

http://h20230.www2.hp.com/new_access_levels.jsp

HP Software Solutions Now accesses the HPSW Solution and Integration Portal Web site. This site enables you to explore HP Product Solutions to meet your business needs, includes a full list of Integrations between HP Products, as well as a listing of ITIL Processes. The URL for this Web site is <http://h20230.www2.hp.com/sc/solutions/index.jsp>

Contents

Contents	3
Chapter 1: Data Protector Oracle Server integration	11
Introduction	11
Integration concepts	12
Configuring the integration	16
Prerequisites	16
Limitations	17
Before you begin	18
Cluster-aware systems	19
Linking Oracle Server with the Data Protector MML	19
Linking on HP OpenVMS systems	19
Configuring Oracle user accounts	20
Configuring Oracle operating system user accounts	20
Clusters	21
Configuring Oracle database user accounts	21
Configuring user accounts on HP OpenVMS systems	22
Configuring Oracle databases	23
Using the Data Protector GUI	24
Using the Data Protector CLI	26
Configuring multiple Oracle databases simultaneously	29
XLS files	31
CSV files	31
Checking the configuration	33
Using the Data Protector GUI	34
Using the Data Protector CLI	34
Handling errors	34
Setting environment variables	35
Using the Data Protector GUI	36
Using the Data Protector CLI	36
Backup	37

Creating new templates	37
Creating backup specifications	38
Examples of pre-exec and post-exec scripts on UNIX systems	44
Editing the Oracle RMAN script	45
Creating copies of backed up objects	48
Testing the integration	49
Testing using the Data Protector GUI	49
Testing using the CLI	50
Starting backup sessions	51
Scheduling backup sessions	54
Running an interactive backup	55
Starting a backup using the GUI	55
Starting a backup using the CLI	56
Starting Oracle backup using RMAN	57
Examples of the RMAN scripts	60
Restore	63
Prerequisites	65
Restoring Oracle using the Data Protector GUI	65
Restoring database items in a disaster recovery	65
Changing the database state	66
Restoring the recovery catalog database	66
Restoring the control file	68
Restoring Oracle database objects	69
Restoring tablespaces and datafiles	74
Restoring and recovering an Oracle database in Oracle Data Guard environment	75
Restoring and recovering a primary database	75
Restoring and recovering a standby database	75
Duplicating an Oracle database	75
Restore, recovery, and duplicate options	78
Restore action options	78
General options	78
Duplicate options	79

Restore and recovery options	80
Restoring Oracle using RMAN	82
Preparing the Oracle database for restore	83
Connection strings used in the examples	84
SBT_LIBRARY parameter	85
Example of full database restore and recovery	85
Example of point-in-time restore	86
Example of tablespace restore and recovery	87
Example of datafile restore and recovery	89
Example of archive log restore	91
Example of database restore using a different device (with the automatic device selection functionality disabled)	92
Restoring using another device	93
Disaster recovery	93
Monitoring sessions	94
Monitoring current sessions	95
Viewing previous sessions	95
Resuming sessions	96
Using the Data Protector GUI	97
Using the Data Protector CLI	98
Aborting sessions	99
Oracle RMAN metadata and Data Protector Media Management Database synchronization	99
Troubleshooting	100
Before you begin	100
Checks and verifications	100
Problems	107
Chapter 2: Data Protector SAP R/3 integration	112
Introduction	112
Integration concepts	113
Backup flow	116
Restore flow	117
Data Protector SAP R/3 configuration file	118

Setting, retrieving, listing, and deleting Data Protector SAP R/3 configuration file parameters using the CLI	120
Configuring the integration	123
Prerequisites	123
Before you begin	124
Cluster-aware clients	124
Configuring user accounts	125
Checking the connection	125
Authentication password file	126
Enabling archived logging	126
Linking Oracle Server with the Data Protector MML	128
Choosing authentication mode	128
Configuring SAP R/3 databases	128
Before you begin	129
Using the Data Protector GUI	129
Using the Data Protector CLI	132
Handling errors	133
Checking the configuration	134
Using the Data Protector GUI	134
Using the Data Protector CLI	134
Backup	135
Considerations	137
Creating backup specifications	137
Modifying backup specifications	142
Scheduling backup sessions	142
Scheduling example	143
Previewing backup sessions	143
Using the Data Protector GUI	144
Using the Data Protector CLI	144
What happens during the preview?	144
Starting backup sessions	144
Backup methods	144

Using the Data Protector GUI	145
Using the Data Protector CLI	145
Using the SAP BRTOOLS	145
Backing up using Oracle Recovery Manager	146
Manual balancing	147
Restore	148
Considerations	148
Restoring using the Data Protector GUI	148
Restoring using the Data Protector CLI	150
Restoring using the SAP commands	151
Restoring using another device	152
Using the Data Protector GUI	152
Using the Data Protector CLI or SAP commands	152
Localized SAP R/3 objects	152
Sparse files	153
Disaster recovery	153
Restoring the control file	153
Monitoring sessions	154
Troubleshooting	154
Before you begin	154
General troubleshooting	154
Troubleshooting on Windows systems	155
Prerequisites concerning the Oracle side of the integration	155
Prerequisites on the SAP side of the integration	157
Configuration problems	158
Backup problems	160
Restore problems	161
Troubleshooting on UNIX systems	164
Prerequisites concerning the Oracle side of the integration	164
Prerequisites on the SAP side of the integration	166
Configuration problems	167
Backup problems	169

Restore problems	171
Chapter 3: Data Protector SAP MaxDB integration	174
Introduction	174
Integration concepts	175
Backup flow	176
Restore flow	176
Configuring the integration	176
Prerequisites	177
Limitations	177
Before you begin	177
Cluster-aware clients	177
Configuring SAP MaxDB users	178
Configuring SAP MaxDB instances	178
Before you begin	178
Using the Data Protector GUI	179
Using the Data Protector CLI	180
Handling errors	181
Checking the configuration	181
Using the Data Protector GUI	181
Using the Data Protector CLI	182
Backup	182
Creating backup specifications	182
Modifying backup specifications	185
Scheduling backup sessions	185
Scheduling example	185
Previewing backup sessions	186
Using the Data Protector GUI	186
Using the Data Protector CLI	186
What happens during the preview?	187
Starting backup sessions	187
Backup methods	187
Using the Data Protector GUI	187

Using the Data Protector CLI	188
Using SAP MaxDB utilities	188
Restore	191
Restore and recovery overview	191
Before you begin	194
Restoring using the Data Protector GUI	194
Restoring using the Data Protector CLI	197
Restoring using SAP MaxDB utilities	198
SAP MaxDB restore and recovery	198
SAP MaxDB migration	202
Finding information for restore	202
SAP MaxDB restore options	202
Restoring using another device	206
Monitoring sessions	206
Troubleshooting	206
Before you begin	206
Problems	207
SAP MaxDB cluster-related troubleshooting	209
Chapter 4: Data Protector SAP HANA Appliance integration	210
Introduction	210
Integration concepts	211
Integration limitations	212
Configuring the integration	213
Prerequisites	213
Configuration steps	213
Linking Data Protector backint agent with SAP HANA Appliance	213
Adjusting SAP HANA backup settings for redo logs	214
Backup	215
Creating backup specifications	215
Creating Data Protector parameter files for SAP HANA	217
Modifying backup specifications	218
Starting backup sessions	218

Restore	218
Finding information for restore	218
Starting restore and recovery sessions	219
Monitoring sessions	219
Troubleshooting	219
Before you begin	219
Problems	220
Glossary	222
Index	263
We appreciate your feedback!	272

Chapter 1: Data Protector Oracle Server integration

Introduction

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

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

Backup types

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

- Online backup of a whole database or parts of it
- Online incremental backup (*Oracle* differential incremental backup 1 to 4)
- Offline backup of a whole database
- Backup of archived redo logs only
- Backup of the Oracle database recovery catalog
- Backup of the Oracle control files
- Backup of **recovery files** residing in the **flash recovery area**.

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

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

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

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

Restore types

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

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

Duplicating a database

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

Integration concepts

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

Components

The software components involved in backup and restore processes are:

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

Integration functionality overview

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

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

How does the integration work?

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

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

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

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

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

Using the Data Protector integration with RMAN, you can back up and restore the Oracle control files, datafiles, and archived redo logs.

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

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

Oracle backup types handled by the integration

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

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

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

The full backup type is not related to the number of datafiles included in the backup, and can therefore be performed per single datafile. The data being backed up, regardless of the backup type (full or incremental), is selected and controlled by Oracle.

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

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

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

These features can be used for online or offline backup of the Oracle target database. However, you must ensure that the backup objects (such as tablespaces) are switched into the appropriate state before and after a backup session. For online backup, the database instance must operate in the ARCHIVELOG mode;

whereas for offline backup, objects need to be prepared for backup using the Pre-exec and Post-exec options in the backup specification.

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

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

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

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

Backup flow

A Data Protector scheduled or interactive backup is triggered by the Data Protector Backup Session Manager, which reads the backup specification and starts the `ob2rman.pl` command on the Oracle Server under the operating system user account specified in the backup specification. Further on, `ob2rman.pl` prepares the environment to start the backup, and issues the RMAN backup command. RMAN instructs the Oracle Server processes to perform the specified command.

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

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

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

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

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

Restore flow

A restore session can be started using:

- Data Protector GUI
- RMAN CLI
- Oracle Enterprise Manager GUI

You must specify which objects are to be restored.

A restore from the Data Protector user interface is triggered by the Data Protector Restore Session Manager, which starts the `ob2rman.pl` command. `ob2rman.pl` prepares the environment to start the restore, and issues the RMAN restore command. RMAN checks the recovery catalog (if one is used) or the control file to gather the information about the Oracle backup objects. It also contacts the Oracle

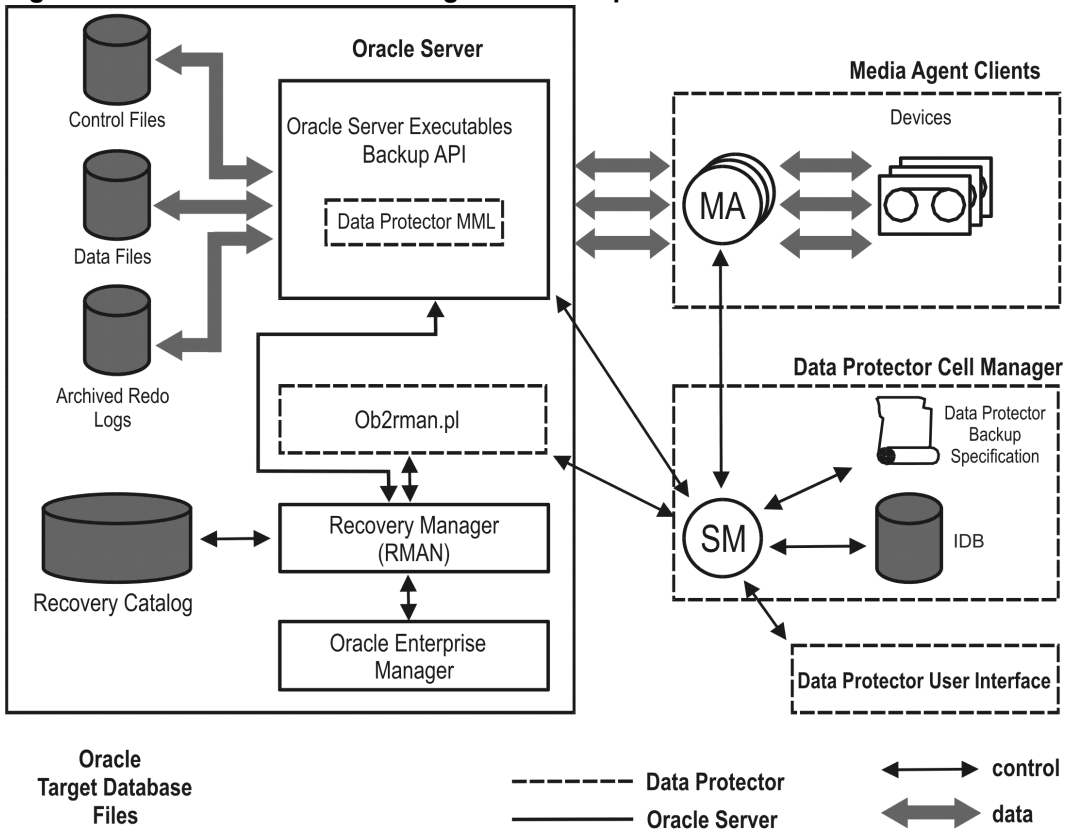
Server processes, which initialize the restore through MML. MML establishes a connection with the Restore Session Manager and passes along the information about which objects and object versions are needed.

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

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

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

Figure 1: Data Protector Oracle integration concept



Database files can also be managed by **Automatic Storage Management (ASM)**. They can reside in the flash recovery area.

Legend

<i>SM</i>	The Data Protector Session Manager, which can be the Data Protector Backup Session Manager during a backup session and the Data Protector Restore Session Manager during a restore session.
<i>RMAN</i>	The Oracle Recovery Manager.

<i>Data Protector MML</i>	The Data Protector Oracle integration Media Management Library, which is a set of routines that enables data transfer between the Oracle Server and Data Protector.
<i>Backup API</i>	The Oracle-defined application programming interface.
<i>IDB</i>	The Data Protector Internal Database where all the information about Data Protector sessions, including session messages, objects, data, and used devices and media, is written.
<i>MA</i>	The Data Protector General Media Agent, which reads and writes data from and to media devices.

Configuring the integration

Prerequisites

- It is assumed that you are familiar with the Oracle database administration and the basic Data Protector functionality.
- You need a license to use the Data Protector Oracle integration. For information on licensing, see the *HP Data Protector Installation and Licensing Guide*.
- Before you begin, ensure that you have correctly installed and configured the Oracle Server and Data Protector client systems. See the:
 - Latest support matrices at <http://support.openview.hp.com/selfsolve/manuals> for an up-to-date list of supported versions, platforms, devices, and other information.
 - *HP Data Protector Installation and Licensing Guide* for instructions on how to install Data Protector on various architectures and how to install the Data Protector Oracle integration.
 - *Oracle Recovery Manager User's Guide and References* for Oracle concepts and backup/recovery strategies.
 - *Oracle Backup and Recovery Guide* for the configuration and use of Recovery Manager, as well as for Oracle backup terminology and concepts.
 - *Oracle Enterprise Manager User's Guide* for information on backup and recovery with the Oracle Enterprise Manager, as well as information about SQL*Plus.
- The Oracle Server software must be installed and the Oracle target database must be open or mounted.
- Starting with Oracle 12c, Oracle Database on Microsoft Windows supports the use of an Oracle Home User, specified at the time of installation. This Oracle Home User is used to run the Windows services for a Oracle home, and is similar to the oracle user on Oracle Database on Linux.

For Oracle 12c database backups and restores, if the Oracle integration agent and Media agent are run on the same Windows host, then to avoid shared memory allocation issues Oracle Home user should be added to the Windows Backup Operator group.

For more information, see the Oracle documentation.

- If the Oracle recovery catalog database is used, ensure that it is properly configured and open.
- Oracle net services must be properly configured and running for the Oracle target database and the recovery catalog, if you use it.

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

For details on checking the prerequisites listed above, see [Troubleshooting on page 100](#).

- To successfully back up the recovery files residing in the flash recovery area, ensure that you have correctly configured the flash recovery area.
- **Oracle Real Application Clusters (RAC):** Each node must have a dedicated disk for storing archive logs. Such disks must be NFS mounted on all other RAC nodes.

However, if the archive logs are not on a NFS mounted disk, you must modify the archive log backup specification. See [Backup of archive logs on RAC cannot be performed on page 108](#).

- **RAC:** With Oracle version 11.2.0.2 and later, the control file must be created on a shared disk and be accessible from all RAC nodes, and the `OB2_DPMCTL_SHRLOC` environment variable must point to this location, from where the control file is backed up.
- On Windows systems, when using the Oracle backup set ZDB method, set the `omnirc` option `ZDB_SMISA_AUTOMOUNTING` on the backup system to 2, in order to enable automatic volume mounting on the local system.

Limitations

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

- The Data Protector Oracle integration (Oracle 12 c) does not support Container Database (CDB) and Pluggable Database (PDB) modes.
- Oracle database identifiers (DBIDs) of all databases must be unique within a Data Protector cell.
- **Oracle Data Guard:**
 - You cannot configure only a standby database (without configuring primary database).
 - Only physical standby database backup is supported.
 - Recovery catalog database is required for standby configurations.
 - For other limitations regarding RMAN backup, restore, recovery, and duplication in Oracle Data Guard environment, see the Oracle documentation.
- The Data Protector Oracle integration does not support non-ASCII characters in backup specification names.

Before you begin

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

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

- On Windows systems, if the Oracle target database and the Oracle recovery catalog are installed on two different systems, configure a *domain* user account that is a member of the Administrators group on both systems.

On Windows Server 2003 systems with the Oracle target database installed, you need to restart the Data Protector Inet service under a Windows domain user account that has the appropriate Oracle database permissions for running backups and restores.

For information on how to change the Data Protector Inet service account, see the *HP Data Protector Help* index: "Inet, changing account".

However, for other supported Windows operating systems, you can use user impersonation instead. For details on setting accounts for the Inet service user impersonation, see the *HP Data Protector Help* index: "Inet user impersonation".

Cluster-aware systems

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

Windows systems: `set OB2BARHOSTNAME=virtual_server_name`

UNIX systems: `export OB2BARHOSTNAME=virtual_server_name`

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

HP-UX with RAC: If you want to use virtual hostname, create an HP Serviceguard package containing only the virtual IP and the virtual hostname parameters and distribute it among the RAC nodes.

Linking Oracle Server with the Data Protector MML

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

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

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

Linking on HP OpenVMS systems

On Oracle Server running on HP OpenVMS systems, link the MML `SYS$SHARE:LIBOBK2SHR64.EXE` with the Oracle Server:

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

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

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

2. Check the presence of the MML `LIBOBK2SHR64.EXE` in the `SYS$SHARE:` directory.

Note: The logical definition for `SYS$SHARE:LIBOBK2SHR64.EXE` is `$DEFINE/SYSTEM DP_SBT SYS$SHARE:LIBOBK2SHR64.EXE`.

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

After relinking

To test the MML (SBT) interface, configure Oracle using the GUI (see [Configuring Oracle databases on page 23](#)).

Configuring Oracle user accounts

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

- Oracle operating system user account

For details, see [Configuring Oracle operating system user accounts below](#).

- Oracle database user accounts

For details, see [Configuring Oracle database user accounts on the next page](#).

Configuring Oracle operating system user accounts

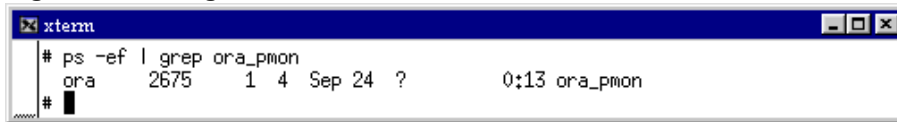
For each Oracle database, Data Protector requires an operating system user account that has Oracle rights to back up the database. This user account usually belongs to the DBA user group (**OSDBA user**). The user account under which the Oracle database is running has these rights. For example, to find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

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

or

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

Figure 2: Finding the Oracle user



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

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

Client system	Description
UNIX system	<p>Ensure that the Oracle user <code>oracle</code> from the Oracle Inventory group (<code>oinstall</code>) has been added to the Data Protector <code>admin</code> user group. For details on adding users, see the <i>HP Data Protector Help</i> index: "adding users".</p> <p>Add the OSDBA user account to the Data Protector <code>admin</code> or <code>operator</code> user group.</p> <div style="background-color: #e0e0e0; padding: 5px;"> <p>Note: If you plan to configure Oracle databases using the <code>omniintconfig.pl</code> command, note that specified OSDBA user accounts are automatically added to the Data Protector <code>admin</code> user group. For details, see Configuring multiple Oracle databases simultaneously on page 29.</p> </div>
Windows system	<p>On Windows systems, Data Protector connects to the Oracle database using the Data Protector <code>Inet</code> service on the related system. By default, the service runs under the <code>Local System</code> account, which is automatically added to the Data Protector <code>admin</code> user group. However, if you have restarted the Data Protector <code>Inet</code> service under an OSDBA user account, you need to add the new user to the Data Protector <code>admin</code> or <code>operator</code> user group.</p>
HP OpenVMS system	<p>Configure a Data Protector <code>admin</code> user with the username <code><Any></code> and the group name <code><Any></code>.</p>

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

Clusters

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

- OSDBA user for all physical nodes
- OSDBA user for the virtual server (applicable for HP Serviceguard clusters)

Configuring Oracle database user accounts

Identify or create the following Oracle database user accounts. You need to provide these user accounts when you configure the Oracle database as described in [Configuring Oracle databases on page 23](#).

Table 1: Oracle database user accounts

User	Description
Primary database user	Required to log in to the primary database.

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

Configuring user accounts on HP OpenVMS systems

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

1. Modify the location of `ORAUSER.COM` and `ORATAB` files as instructed in `OMNI$ROOT:[LOG]LOGIN.COM` based on the Oracle version used.

For example:

- `$PIPE@DKA0:[ORACLE]ORAUSER.COM > NLA0:`

Suppose `ORAUSER.COM` is located in `DKC0:[ORACLE10g]`, then change and uncomment the above statement to `$PIPE@DKC0:[ORACLE10g]ORAUSER.COM > NLA0:`.

- `$DEFINE/NOLOG/JOB ORATAB_LOC DKA0:[ORACLE]ORATAB`

Suppose `ORATAB` is located in `DKC0:[ORACLE10g]`, then change and uncomment the above statement to `$DEFINE/NOLOG/JOB ORATAB_LOC DKCF0:[ORACLE10g]ORATAB`.

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

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

```
$@OMNI$ROOT:[BIN.PERL1]PERL_SETUP.COM
```

```
$DEFINE /NOLOG /PROCESS PERL_ENV_TABLES "LNM$PROCESS", "LNM$JOB", "LNM$SERVER",  
"LNM$GROUP", "LNM$SYSTEM"
```

3. Uncomment the following line:

```
$_@OMNI$ROOT:[BIN]OMNI$ORA_OCI_SETUP.COM
```

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

- a. Log in as a privileged user.
- b. Execute:

```
$set def sys$system
```

```
$mcr authorize
```

```
UAF> show omniadmin
```

```
UAF> show oracle_user
```

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

```
UAF> modify omniadmin /UIC=[Group_ID_of_Oracle_user, User_ID]
```

- d. Verify the changes of the group ID.

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

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

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

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

Configuring Oracle databases

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

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

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

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

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

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

Tip: In large environments with multiple Oracle databases, consider using the configuration procedure described in [Configuring multiple Oracle databases simultaneously on page 29](#). However, note that this procedure cannot be used to configure standby databases.

Using the Data Protector GUI

Configure an Oracle database when you create the first backup specification for the database. Start with the procedure described in [Creating backup specifications on page 38](#) and at [If the Oracle database is not configured yet for use with Data Protector, the Configure Oracle dialog box is displayed](#). Configure the Oracle database for use with Data Protector as described in [Configuring Oracle databases on page 23](#). on [page 40](#) proceed as follows:

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

Figure 3: Configuring Oracle - General (Windows)

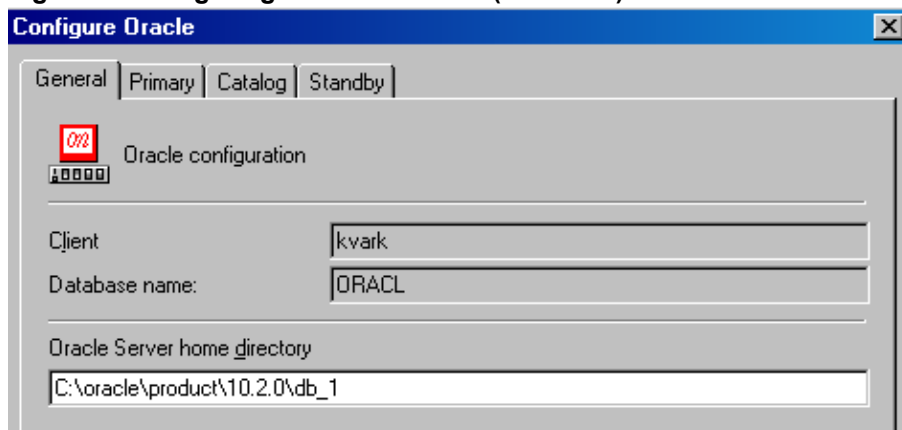
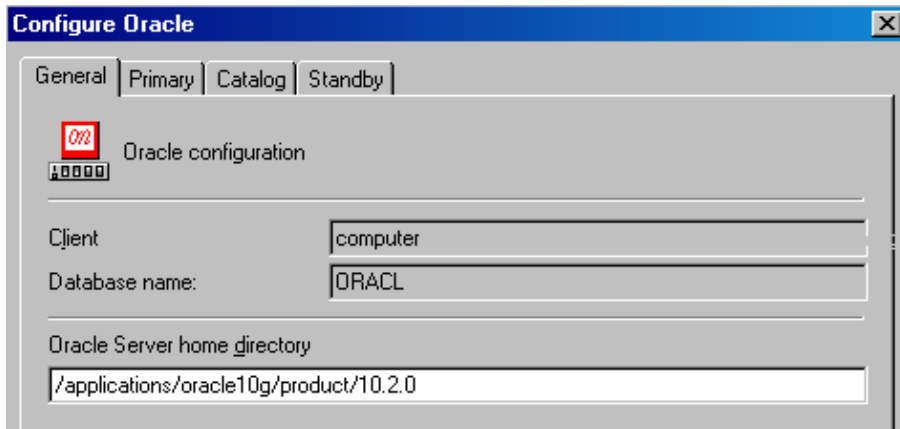


Figure 4: Configuring Oracle - General (UNIX)



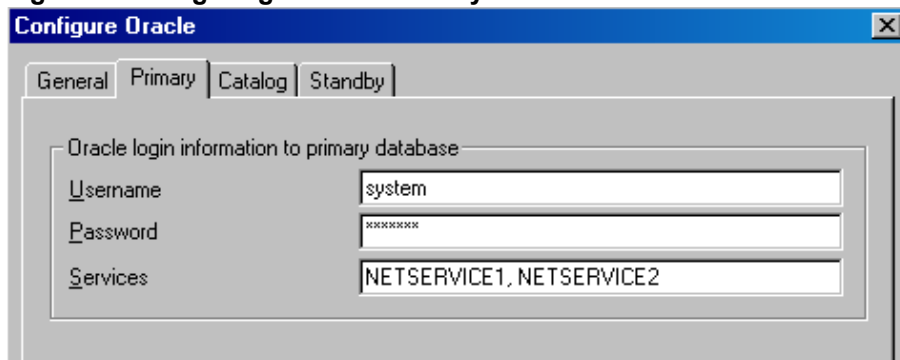
2. In the **Primary** page, specify the login information to the primary database.

Note that the user must have the SYSDBA privilege granted.

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

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

Figure 5: Configuring Oracle - Primary



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

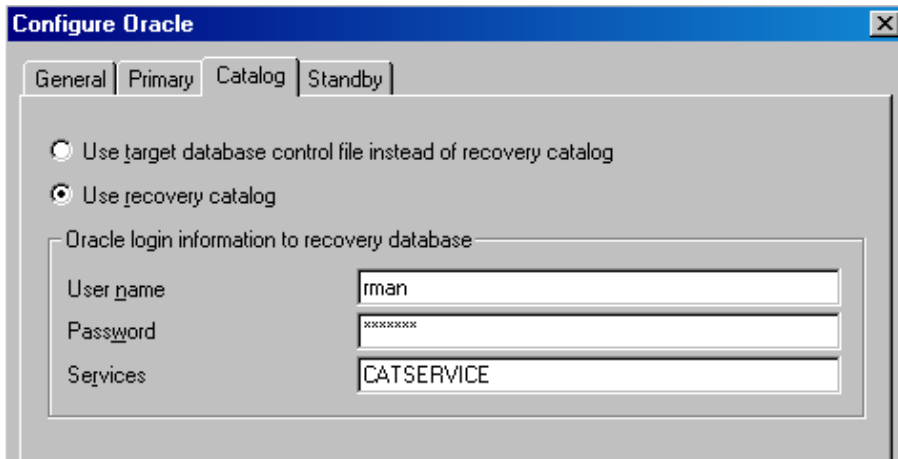
To use the recovery database catalog as an RMAN repository for backup history, select **Use recovery catalog** and specify the login information to the recovery catalog.

Oracle Data Guard: If you intend to back up a standby database, you must use the recovery catalog.

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

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

Figure 6: Configuring Oracle - Catalog



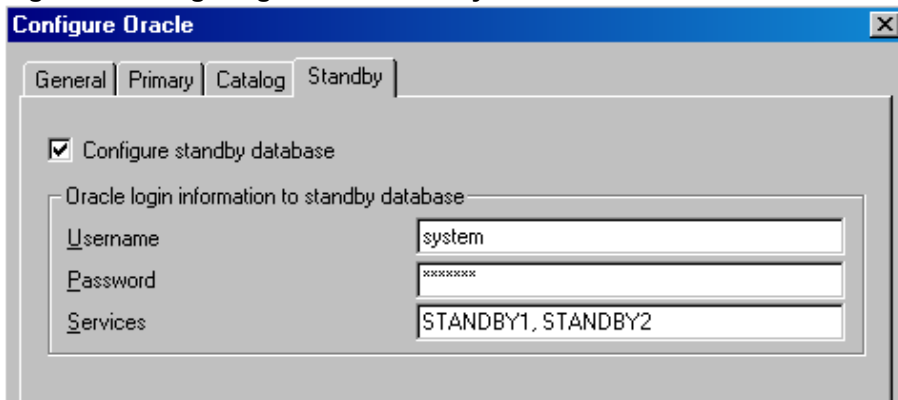
4. In Oracle Data Guard environments, if you intend to back up a standby database, configure also the standby database:

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

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

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

Figure 7: Configuring Oracle - Standby



5. Click **OK**.

The Oracle database is configured. Exit the GUI or proceed with creating the backup specification at [Select the Oracle database objects to be backed up. on page 40.](#)

Using the Data Protector CLI

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

1. On UNIX systems, log on to the Oracle Server system with an OSDBA user account.
2. On the Oracle Server system, execute:

Windows systems:

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

UNIX systems:

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

HP OpenVMS systems:

```
util_oracle8 -config -dbname DB_NAME -orahome ORACLE_HOMEPRIMARY_DB_LOGIN
[CATALOG_DB_LOGIN] [STANDBY_DB_LOGIN] [-client CLIENT_NAME]
```

where:

PRIMARY_DB_LOGIN is:

```
-prmuser PRIMARY_USERNAME
-prmpasswd PRIMARY_PASSWORD
-prmservice PRIMARY_NET_SERVICE_NAME_1[,PRIMARY_NET_SERVICE_NAME_2 ...]
```

CATALOG_DB_LOGIN is:

```
-rcuser CATALOG_USERNAME
-rcpasswd CATALOG_PASSWORD
-rcservice CATALOG_NET_SERVICE_NAME
```

STANDBY_DB_LOGIN is:

```
-stbuser STANDBY_USERNAME
-stbpasswd STANDBY_PASSWORD
-stbservice STANDBY_NET_SERVICE_NAME_1[,STANDBY_NET_SERVICE_NAME_2 ...]
```

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

Parameter description

<i>CLIENT_NAME</i>	Name of the Oracle Server system with the database to be configured. It must be specified in a cluster environment. RAC: The virtual server of the Oracle resource group. Oracle Data Guard: Name of either a primary system or secondary (standby) system.
<i>DB_NAME</i>	Name of the database to be configured.
<i>ORACLE_HOME</i>	Pathname of the Oracle Server home directory.
<i>PRIMARY_USERNAME</i> <i>PRIMARY_PASSWORD</i>	Username and password for login to the target or primary database. Note that the user must have the SYSDBA privilege granted.
<i>PRIMARY_NET_SERVICE_NAME_1</i> [, <i>PRIMARY_NET_SERVICE_NAME_2</i> , ...]	Net services names for the primary database. RAC: Each net service name must resolve into a specific database instance.
<i>CATALOG_USERNAME</i> <i>CATALOG_PASSWORD</i>	Username and password for login to the recovery catalog. This is optional and is used only if you use the recovery catalog database as an RMAN repository for backup history.
<i>CATALOG_NET_SERVICE_NAME</i>	Net service name for the recovery catalog.
<i>STANDBY_USERNAME</i> <i>STANDBY_PASSWORD</i>	This is used in Oracle Data Guard environment for backing up a standby database. Username and password for login to the standby database.
<i>STANDBY_NET_SERVICE_NAME_1</i> [, <i>STANDBY_NET_SERVICE_NAME_2</i> , ...]	Net services names for the standby database.

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

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

Example

The following example represents configuration on a UNIX system of an Oracle database and its recovery catalog in Oracle Data Guard environment.

The following names are used in the example:

- database name: oracle
- Oracle Server home directory: /app10g/oracle10g/product/10.1.0
- primary user name: system

- primary password: manager
- primary net service name 1: netservice1
- primary net service name 2: netservice2
- recovery catalog user name: rman
- recovery catalog password: manager
- recovery catalog net service name: catservice
- standby user name: system
- standby password: manager
- standby net service name 1: netservicesb1
- standby net service name 2: netservicesb2

Syntax

```
/opt/omni/lbin/util_oracle8.pl -config -dbname oracle -orahome  
/app10g/oracle10g/product/10.1.0 -prouser system -prmpasswd manager -prmservice  
netservice1,netservice2 -rcuser rman -rcpasswd manager -rcservice catservice -stbuser  
system -stbpasswd manager -stbservice netservicesb1,netservicesb2
```

Configuring multiple Oracle databases simultaneously

In large environments with multiple Oracle databases, it can be time-consuming to configure each database separately, especially if the configuration parameters need to be updated frequently.

For these reasons, Data Protector enables you to keep configuration parameters of multiple databases in a single file. In this way, you can do all necessary updates in one place. Once the file is ready, you execute the Data Protector `omniintconfig.pl` command, which reads the file and configures all the Oracle databases specified. It means that, for each Oracle database, a separate Data Protector configuration file is created or updated (if it already exists), similarly as if the standard configuration method were used. If specified, Data Protector also performs a configuration check.

In your configuration file, you specify the following parameters for each Oracle database:

Table 2: Oracle database configuration parameters

Parameter	Description
MoM (optional)	Manager of managers

Parameter	Description
CellManager	Data Protector Cell Manager Default: Cell Manager of the local client
Client	Client with the Oracle Server installed. In cluster environments, specify the virtual server or, in RAC, one of the cluster nodes. Default: local client
Instance	Oracle database instance (mandatory)
OSUSER (UNIX and Windows Server 2008 systems only)	An operating system user account (user name and group or domain) under which you want the configuration and browsing of Oracle databases to start. This user will be automatically added to the Data Protector admin user group for the client specified in Client. On Windows Server 2008, it is not mandatory to specify the user account.
OSGROUP (UNIX and Windows Server 2008 systems only)	
ORACLE_HOME	Oracle Server home directory
TGTUser	Login information for the target database (username and password)
TGTPasswd	
TGTService	Target database service(s). If there is more than one service, separate them with a semicolon (service1;service2...).
RCUser (optional)	Login information for the recovery catalog database (username and password)
RCPasswd (optional)	
RCSERVICE (optional)	Recovery catalog database service

Parameter	Description
ClusterNodes (optional)	<p>Cluster nodes (applicable in cluster environments). The user OSUSER, OSGROUP will be automatically added to the Data Protector admin user group for each cluster node listed here. Separate cluster nodes with a semicolon (node1;node2...).</p> <p>If you do not specify this parameter, you need to add these users manually as described in Configuring Oracle user accounts on page 20.</p>

File formats

Your file must be created in one of the following formats:

- XLS (Microsoft Office Excel file)
- CSV (comma separated values file)

When creating the file, consider the following:

- In the first line, list parameters that you want to specify. In subsequent lines, list parameter values for Oracle databases that you want to configure.
- Parameter names in the first line are not case-sensitive.
- Empty columns are not allowed.
- Empty rows are allowed.
- Empty cells are allowed only for optional parameters.

XLS files

In XLS files, you can format cells as you like. However, you are not allowed to add any information in extra cells. See [Keeping parameters in an XLS file below](#).


Figure 8: Keeping parameters in an XLS file

	A	B	C	D	E	F	G	H	I
1									
2	CellManager	Client	Instance	OSUSER	OSGROUP	ORACLE_HOME	TGTUser	TGTPasswd	TGTService
3	galaxy	helios	IN1	ora101	dba	appls/ora10	system	ZUIOZUIOW	IN1
4	galaxy	star	IN2			C:\Oracle\product\10.	system	GHUJKGHJK	IN2_1;IN2_2

CSV files

A CSV file is created by saving a text file in CSV format (for example, C:\My_documents\Oracle_databases.csv). Parameters in the file must be separated with commas. You can omit the specification of parameters that are not applicable by leaving the place between two commas empty. See [Keeping parameters in a CSV file below](#).

Figure 9: Keeping parameters in a CSV file




Encoding passwords

Data Protector requires that passwords in Data Protector Oracle database configuration files are encoded. You can achieve this in two different ways:

- Encode the passwords before you save them in your XLS or CSV file, using the Data Protector `util_cmd` command. For example, to encode the password `BlueMoon`, execute:

```
util_cmd -encode BlueMoon
```

Figure 10: Encoding a password



Once you receive the encoded password, copy it to your file. [Keeping parameters in an XLS file on the previous page](#) shows an example of a file in which all the passwords are encoded.

If you keep your passwords encoded, you do not need to specify the `-encode` option when you execute the `omniintconfig.pl` command.

- If your passwords are not encoded, specify the `-encode` option when you execute the `omniintconfig.pl` command.

Important: Ensure that the passwords in your XLS or CSV file are either all encoded or all plain-text.

omniintconfig.pl command syntax

Note: The `omniintconfig.pl` command can be run on any Data Protector client that has the User Interface component installed.

1. Log on to the client system under an operating system user account that is added to the Data Protector admin user group (actually, it suffices if the user has the Data Protector User configuration and See private objects user rights).
2. Go to the following default Data Protector administrative commands directory.
3. Execute:

Windows systems: `perl omniintconfig.pl Options`

UNIX systems: `omniintconfig.pl` *Options*

where *Options* are:

```
[-encode]
    [-chkconf]
    [-force]
    {-passwordfile FileName|Param=Value [Param=Value...]}
```

For the options description, see the `omniintconfig.pl` man page or the *HP Data Protector Command Line Interface Reference*.

Examples

1. Suppose you are logged in to the Windows system on which you have created the file `C:\My_documents\Oracle_instances.xls`. To configure the Oracle databases IN1 and IN2 using the information from the file, execute:

```
perl omniintconfig.pl -passwordfile
C:\My_documents\Oracle_instances.xls
```

2. Suppose you are logged in to a UNIX system. To configure the Oracle database IN2 by specifying parameters at run time, execute:

```
omniintconfig.pl -encode CellManager=galaxy Client=star
Instance=IN2 ORACLE_HOME=C:\oracle\product\10.2.0\db_1 TGTUser=system
TGTSERVICE=IN2_1;IN2_2 TGTPasswd=BlueMoon
```

Note that the password `BlueMoon` is not encoded. Therefore, you must specify the option `-encode`.

Parameters can be specified only for one Oracle database at a time.

3. Suppose you are logged in to a Windows system. To configure and check the configuration of all Oracle databases specified in `C:\My_documents\Oracle_instances.xls`, execute:

```
perl omniintconfig.pl -chkconf -force -passwordfile
C:\My_documents\Oracle_instances.xls
```

The `-force` option instructs Data Protector to continue configuring Oracle databases if the configuration check for an Oracle database fails.

4. Suppose you are logged in to a UNIX system. To check the configuration of the Oracle database IN2, execute:

```
omniintconfig.pl -chkconf CellManager=galaxy Client=star Instance=IN2
```

Checking the configuration

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

Using the Data Protector GUI

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

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

Using the Data Protector CLI

1. On UNIX systems, log on to the Oracle Server system with an OSDBA user account.
2. Execute:

Windows systems:

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

UNIX systems:

```
util_oracle8.pl -chkconf -dbname DB_NAME
```

HP OpenVMS systems:

```
util_oracle8 -chkconf -dbname DB_NAME
```

Handling errors

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

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

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

HP-UX and Linux systems: `/opt/omni/lbin/omnigetmsg 12 error_number`

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

HP OpenVMS systems:

Set up the Data Protector CLI environment by executing:

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

Execute:

```
$OMNIGETMSG 12 error_number
```

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

Setting environment variables

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

For details of how environment variables affect your environment, see [Environment variables below](#).

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

Table 3: Environment variables

Environment variable	Default value	Description
OB2_RMAN_COMMAND_TIMEOUT	300 s	This variable is applicable when Data Protector tries to connect to a target or catalog database. It specifies how long (in seconds) Data Protector waits for RMAN to respond that the connection succeeded. If RMAN does not respond within the specified time, Data Protector aborts the current session.
OB2_SQLP_SCRIPT_TIMEOUT	300 s	This variable is applicable when Data Protector issues an SQL*Plus query. It specifies how long Data Protector waits for SQL*Plus to respond that the query completed successfully. If SQL*Plus does not respond within the specified time, Data Protector aborts the current session.
OB2_DPMCTL_SHRLOC	N/A	Defines the location at which the control file is created and from where it is backed up in Data Protector managed control file backup. Data Protector copies the control file to its temporary files directory. This variable overrides the default directory with a customer-specified directory. In an Oracle Real Application Clusters (RAC) environments with Oracle version 11.2.0.2 or later, to enable Data Protector managed control file backups and the corresponding restore sessions, ensure this directory resides on a shared disk that all RAC nodes can access.

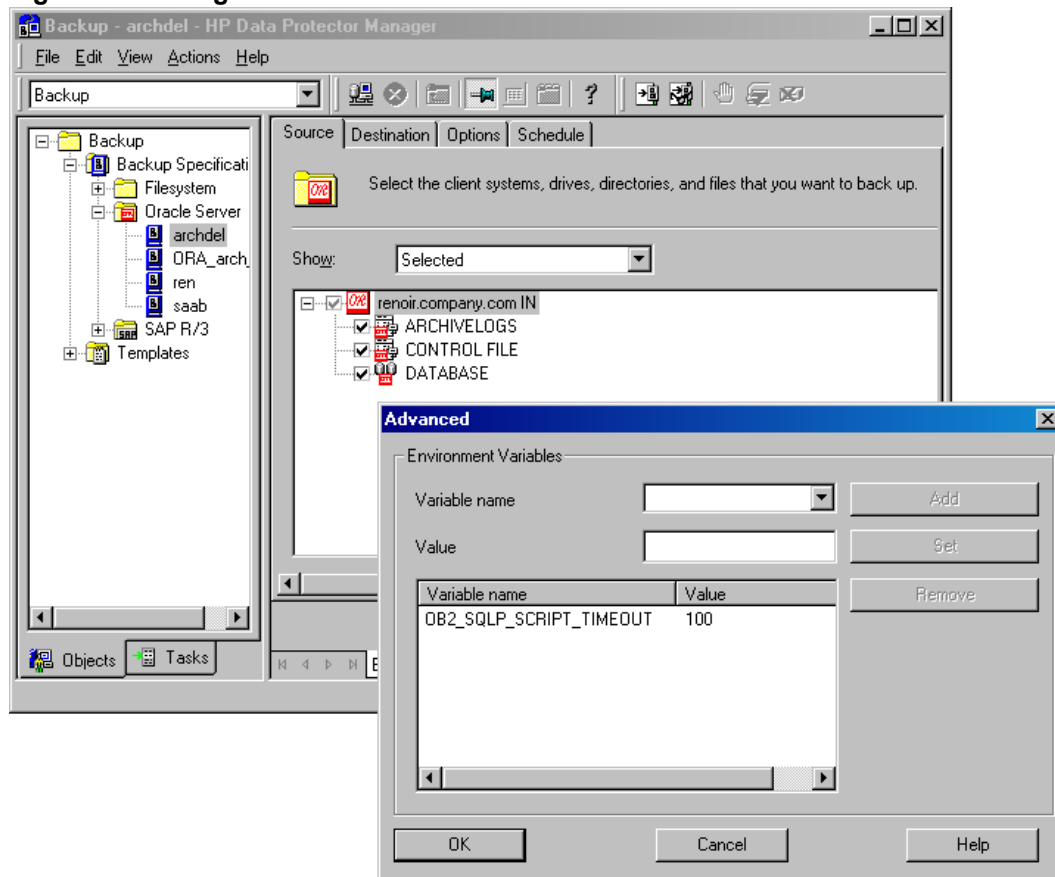
To set environment variables, use the Data Protector GUI or CLI.

Using the Data Protector GUI

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

1. In the Source page of the backup specification, right-click the Oracle database at the top and click **Set Environment Variables**.
2. In the Advanced dialog box, specify the variable name, its value, and click **Add**.
See [Setting environment variables below](#).

Figure 11: Setting environment variables



Click **OK**.

Using the Data Protector CLI

Execute:

```
util_cmd -putopt Oracle8 DatabaseNameVariableValue -sublist Environment
```

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

Example

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

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

Backup

To configure an Oracle backup, perform the following steps:

1. Configure the devices you plan to use for a backup. For instructions, see the *HP Data Protector Help* index: “configuring devices”.
2. Configure media pools and media for a backup. For instructions, see the *HP Data Protector Help* index: “creating media pools”.
3. Ensure you are able to connect to the database.
4. Create a Data Protector Oracle backup specification.
See [Creating backup specifications on the next page](#).

HP OpenVMS systems

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

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

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

Creating new templates

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

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

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

To create a new backup template, proceed as follows:

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

Creating backup specifications

Cluster-aware systems

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.

Procedure

To create an Oracle backup specification:

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **Oracle Server**, and click **Add Backup**.
3. In the Create New Backup dialog box, double-click **Blank Oracle Backup** to create a backup specification without predefined options, or use one of the pre-defined templates given below:

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

Click **OK**.

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

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

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

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

The database name can be obtained using SQL*Plus:

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

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

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

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

- **UNIX systems:** In **Username** and **Group/Domain name**, specify the OSDBA user account under which you want the backup to start (for example, the user name `ora`, group `DBA`). This user must be configured as described in [Configuring Oracle user accounts on page 20](#).
- **Windows Server 2008 systems:** It is not mandatory to specify these options and if they are not specified, the backup runs under the Local System Account.

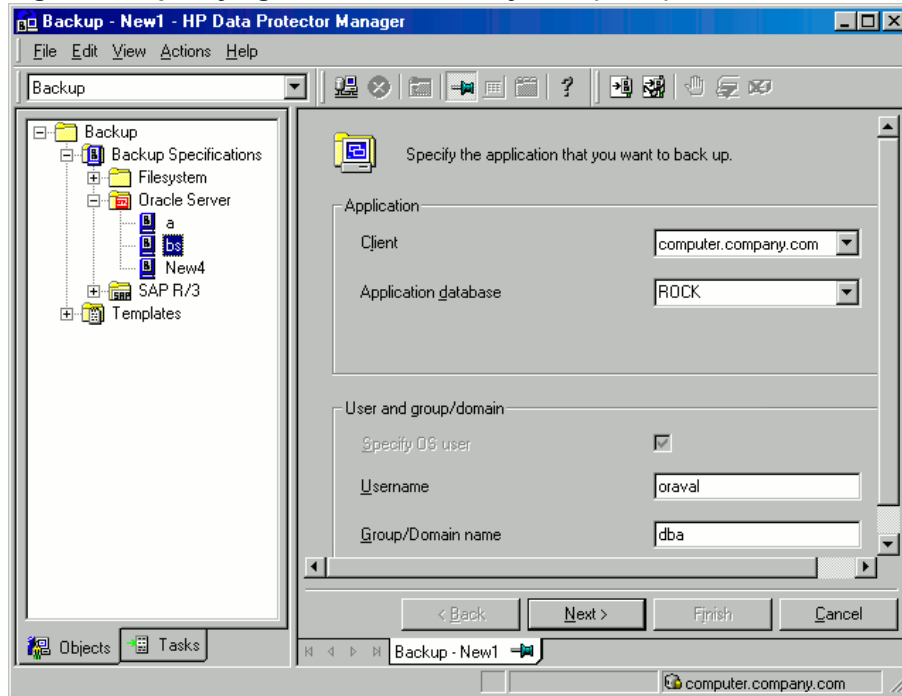
In **Username** and **Group/Domain name**, specify the operating system user account under which you want the backup session to run (for example, the user name `Administrator`, domain `DP`). This user must be set up for the Data Protector Inet service user impersonation.

For details on setting accounts for the Inet service user impersonation, see the *HP Data Protector Help* index: "Inet user impersonation".

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

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

Figure 12: Specifying an Oracle Server system (UNIX)



Click **Next**.

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

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

5. If the Oracle database is not configured yet for use with Data Protector, the Configure Oracle dialog box is displayed. Configure the Oracle database for use with Data Protector as described in [Configuring Oracle databases on page 23](#).

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

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

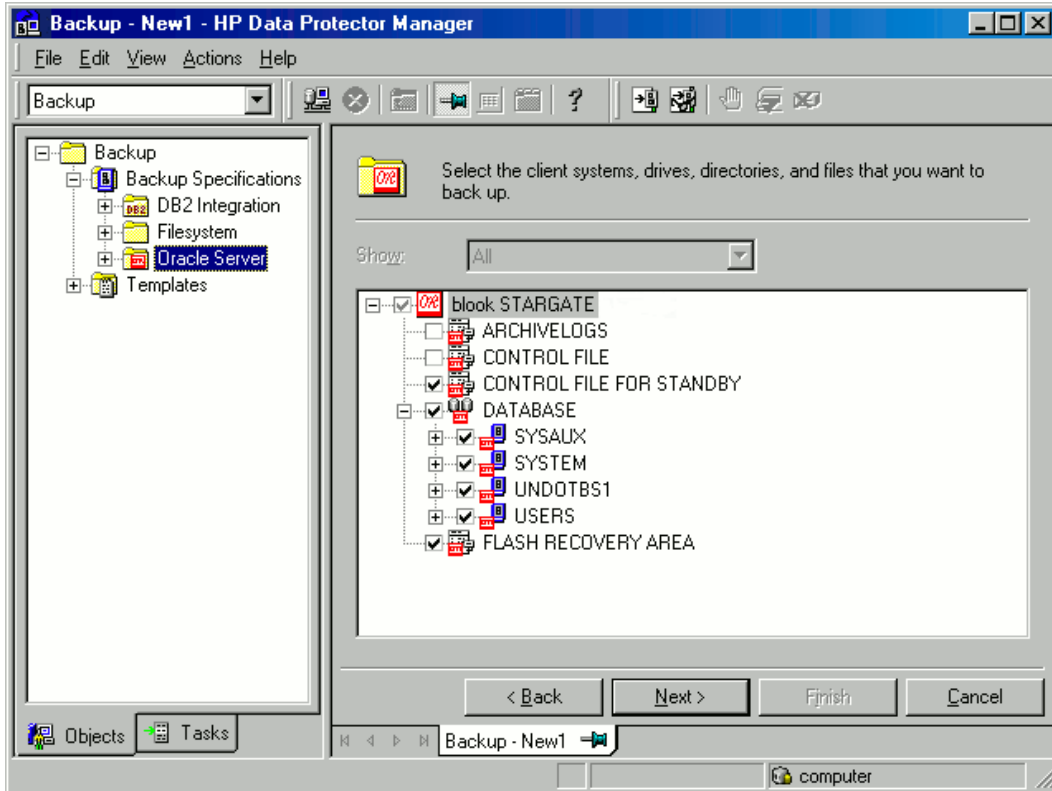
The archived logs can reside in the flash recovery area. In this case, if you select the **FLASH RECOVERY AREA** to be backed up, you do not need to select also **ARCHIVELOGS**.

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

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

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

Figure 13: Selecting backup objects



Click **Next**.

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

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

For detailed information on the object mirror functionality, see the *HP Data Protector Help* index: "object mirroring".

Click **Next** to proceed.

Set the backup options.

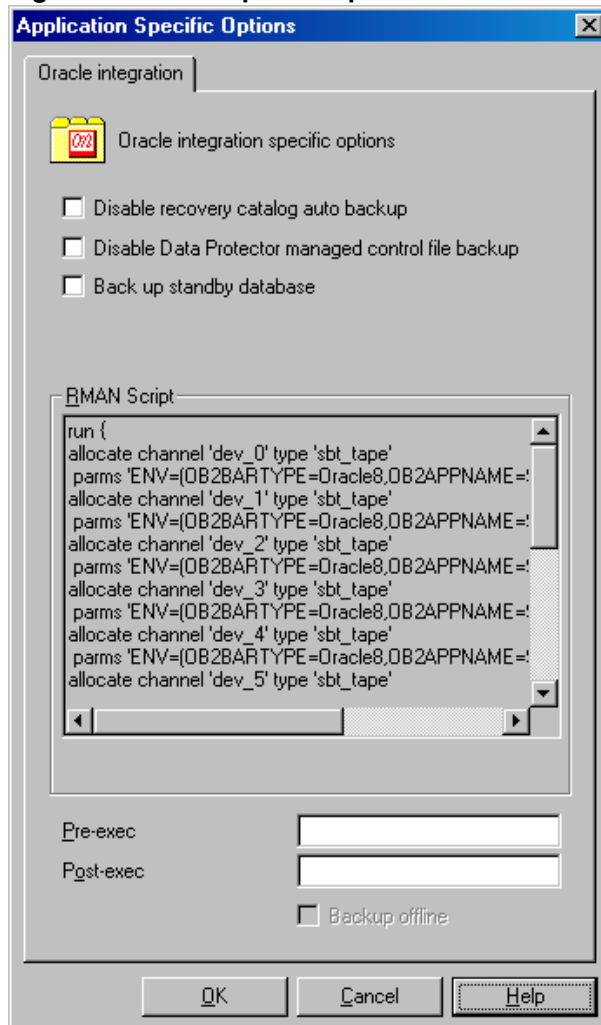
For information on other the Backup Specification Options and Common Application Options, press **F1**.

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

For information on the Application Specific Options, see [Oracle backup options on page 44](#) or press **F1**.

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

Figure 14: Oracle-specific options



8.

Click **Next**.

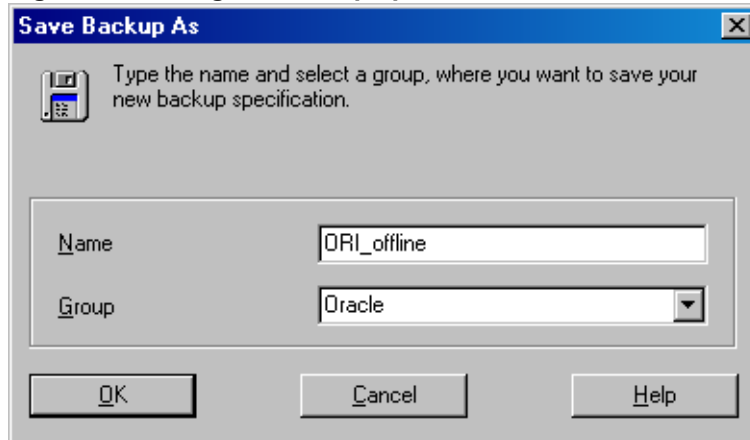
9. Optionally, schedule the backup. For more details, see [Scheduling backup sessions on page 54](#).

Click **Next**.

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

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

Figure 15: Saving the backup specification



Click **OK**.

To start the backup, see [Starting backup sessions on page 51](#).

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

Windows systems: `Data_Protector_program_data\Config\server\Barlists\Oracle8\Backup_Specification_Name`

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

12. It is recommended to test the backup specification. See [Testing the integration on page 49](#) for details.

Table 4: Oracle backup options

<p>Disable recovery catalog auto backup</p>	<p>By default, Data Protector backs up the recovery catalog in every backup session. Select this option to disable backup of the recovery catalog.</p>
<p>Disable Data Protector managed control file backup</p>	<p>By default, Data Protector backs up the Data Protector managed control file in every backup session. Select this option to disable backup of the Data Protector managed control file.</p>
<p>Back up standby database</p>	<p>Oracle Data Guard: This option is applicable if the database is configured with the standby connection. By default, RMAN backs up the database files and archived redo logs on the primary system. Select this option to enable backup of the database files and archive logs on standby system. However, only the archive logs created after the standby database was configured can be backed up at standby site. Archive logs created before the standby database was configured must be backed up on the primary database.</p> <p>Note that the current control file or the control file for standby will still be backed up from the primary system.</p>
<p>RMAN Script</p>	<p>You can edit the Oracle RMAN script section of the Data Protector Oracle backup specification. The script is created by Data Protector during the creation of a backup specification and reflects the backup specification's selections and settings. You can edit the script only after the backup specification has been saved. For information on how to edit the RMAN script section, see Editing the Oracle RMAN script on the next page.</p>
<p>Pre-exec, Post-exec</p>	<p>Specify a command or RMAN script that will be started by <code>ob2rman.pl</code> on the Oracle Server system before the backup (pre-exec) or after it (post-exec). RMAN scripts must have the <code>.rman</code> extension. Do not use double quotes.</p> <p>For example, you can provide scripts to shut down and start up an Oracle instance. For examples of shut-downing and starting an Oracle instance on a UNIX system, see Examples of pre-exec and post-exec scripts on UNIX systems below.</p> <p>Provide the pathname of the command or RMAN script.</p> <p>HP OpenVMS systems : Provide the pathname of the command (OMNI\$ROOT:[BIN]).</p>

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

Pre-exec example

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

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

Post-exec example

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

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

Editing the Oracle RMAN script

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

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

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

Limitations

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

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

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

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

To edit an Oracle RMAN script, click **Edit** in the **Application Specific Options** window (see [Recovery catalog settings dialog on page 67](#)), edit the script, and then click **Save** to save the changes to the script.

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

Data Protector RMAN script structure

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

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

The number of allocated channels is the same as the sum of concurrency numbers for all devices selected for backup.

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

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

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

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

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

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

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

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

- The RMAN datafile *tablespace_name*datafile_name* command.
- If the archived redo logs were selected for a backup, **an RMAN backup statement for the backup of Oracle archive logs.**

If an appropriate template was selected, or if the statement was manually added, the RMAN sql statement to switch the online redo logs before backing up the archived redo logs:

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

The backup statement consists of the following:

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

```
format 'Backup_Specification_NameDB_NAME_%s:%t:%p>.dbf'
```

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

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

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

- The RMAN `current controlfile` command.

Example of the RMAN script

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

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

Creating copies of backed up objects

Oracle duplex mode

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

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

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

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

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

```
%c
```


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

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

Example

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

Important: If the MIN and MAX load balancing parameters are set to lower values, the backup session will get blocked.

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.

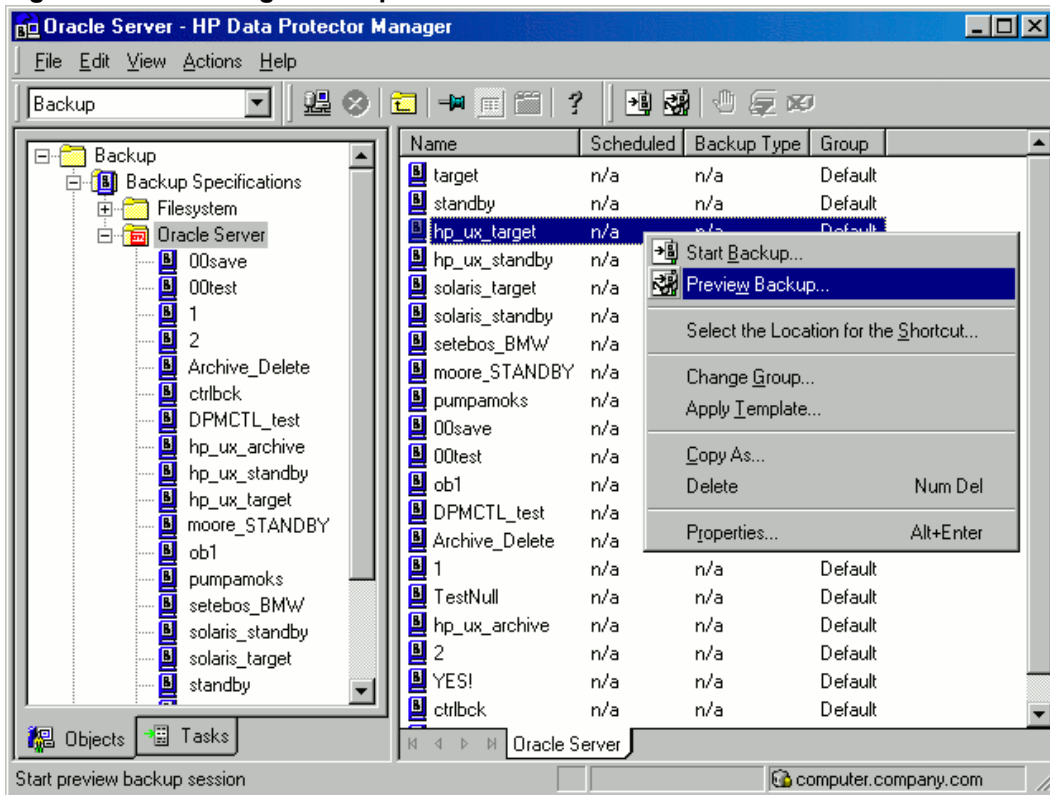
Details of the test backup, such as media protection, backup user and backup status are registered in the Data Protector database and in the Oracle control files. Set the **Protection** option of your test backup specification to **None**.

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 16: Previewing a backup



Testing using the CLI

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

Note: On HP OpenVMS systems, to invoke the Data Protector CLI, execute: `$_OMNIB$ROOT:[BIN] OMNIB$CLI_SETUP.COM`

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

Windows systems: `Data_Protector_home\bin\omnib -oracle8_list backup_specification_name -test_bar`

HP-UX, Solaris, and Linux systems: `/opt/omni/bin/omnib -oracle8_list bimbackup_specification_name -test_bar`

Other UNIX systems: `/usr/omni/bin/omnib -oracle8_list backup_specification_name -test_bar`

HP OpenVMS systems: `$omnib -oracle8_list backup_specification_name -test_bar`

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

Starting backup sessions

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

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

Oracle offline

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

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

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

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

The whole offline database backup is performed as follows:

1. Shut down the database cleanly.

A clean shutdown means that the database is not shut down using the ABORT option.
2. Mount the database if you are backing it up using RMAN.
3. Back up all datafiles, control files and, optionally, parameter files.
4. Start up the database again in the normal online mode.

Oracle online

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

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

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

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

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

The database must operate in the ARCHIVELOG mode so that the current online redo logs are archived to the archived redo logs.

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

```
archive log list;
```

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

When SPFILE is used:

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

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

When PFILE is used:

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

```
log_archive_start = true
```

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

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

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

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

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

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

Backup methods

- Schedule a backup of an existing Oracle backup specification using the Data Protector Scheduler. See [Scheduling backup sessions on the next page](#).
- 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 55](#).
- Start a backup on the Oracle server using either Oracle Recovery Manager or Oracle Enterprise Manager. See [Starting Oracle backup using RMAN on page 57](#).

Backup procedure

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

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

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

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

Deleting data from the recovery catalog

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

Scheduling backup sessions

For more information on scheduling, see the *HP Data Protector Help* index: “scheduled backups”.

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

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

An example of scheduling backups of production databases:

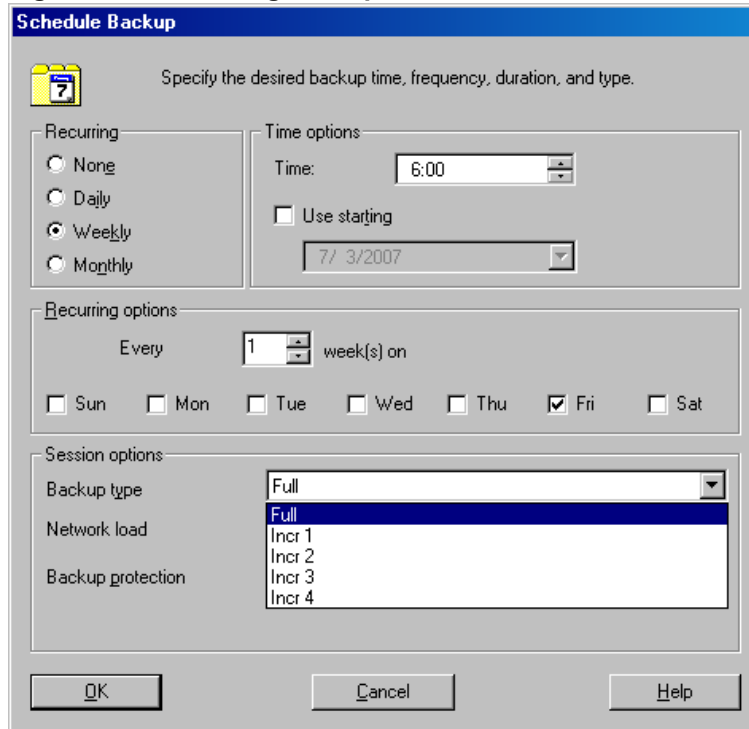
- Weekly full backup
- Daily incremental backup
- Archived Log backups as needed

To schedule an Oracle backup specification, proceed as follows:

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

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

Figure 17: Scheduling backup sessions



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

Running an interactive backup

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

Starting a backup using the GUI

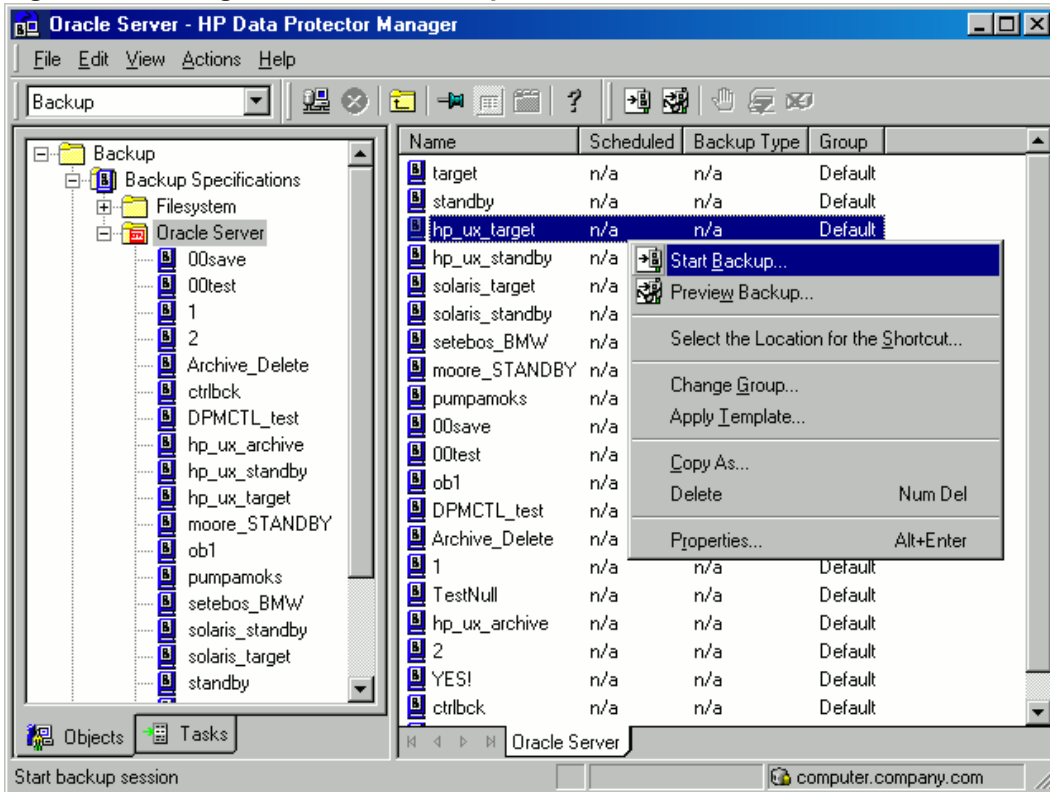
To start an interactive backup of an Oracle database using the Data Protector GUI, proceed as follows:

1. In the Context List, click **Backup** context.
2. In the Scoping Pane, expand **Backup Specifications** and then **Oracle Server**. Right-click the backup specification you want to use and click **Start Backup**.
3. In the **Start Backup** dialog box, select the **Backup type** and 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 [Scheduling backup sessions above](#). See the RMAN documentation for details on incremental backup levels.

Click **OK**.

Figure 18: Starting an interactive backup



Starting a backup using the CLI

1. On an Oracle Server, switch to the default Data Protector user commands directory.

HP OpenVMS systems: To set up the CLI, execute:

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

2. Execute:

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

You can select among the following *list_options*:

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

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

```
-crc
```

```
-no_monitor
```

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

See the omnib man page for details.

Example

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

```
omnib -oracle8_list RONA
```

Starting Oracle backup using RMAN

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

For information on how to create an Oracle backup specification, see [Backup on page 37](#).

To start an Oracle backup using RMAN:

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

If you use the recovery catalog, execute:

Windows systems: `ORACLE_HOME\bin\rman target Target_Database_Login catalog Recovery_CatLog_Login`

UNIX systems: `ORACLE_HOME/bin/rman target Target_Database_Login catalog Recovery_CatLog_Login`

HP OpenVMS systems:

- a. Execute `ORAUSER.COM` using `$_OMNI$ROOT:[LOG]LOGIN.COM`.
- b. Execute `$rman target target_connect_string catalog catalog_connect_string`.

Target database login

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

where:

user_name is the name by which a user is known to Oracle Server and to other users. Every user name is associated with a password and both have to be entered to connect to an Oracle target database. This user must have been granted Oracle SYSDBA or SYSOPER rights.

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

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

Recovery catalog login

The format of the Recovery Catalog Database login is `user_name/password@service`,

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

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

2. Allocate the Oracle channels.

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

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

or

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

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

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

Windows and UNIX systems: Specify the path to the Data Protector MML at run time by setting the `SBT_LIBRARY` RMAN script parameter. For details, see [Specify the parms operand: below](#).

HP OpenVMS system: Ensure that a symbolic link to the Data Protector MML exists.

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 32 or 64 (if device is local) channels can be allocated.

3. Specify the parms operand:

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

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

On Windows and UNIX systems, set the `SBT_LIBRARY` parameter to point to the correct platform-specific Data Protector MML. The location and the filename of the Data Protector MML depend on the platform:

Windows systems: `Data_Protector_home\bin`

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

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

Table 5: MML filenames on different platforms

Platform	32-bit	64-bit
HP-UX	libob2oracle8.sl	libob2oracle8_64bit.sl
HP-UX on Itanium	libob2oracle8.so	libob2oracle8_64bit.so
Solaris	libob2oracle8.so	libob2oracle8_64bit.so
AIX	libob2oracle8.a	libob2oracle8_64bit.a
Other UNIX systems	libob2oracle8.so	libob2oracle8_64bit.so
Windows	orasbt.dll	orasbt.dll
HP OpenVMS	N/A	LIBOBK2SHR_64.EXE

For example, on 32-bit Solaris system, set
 SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so.

4. Specify format:

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

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

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

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

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

5. Optionally, specify backup incremental level.

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

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

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

Windows systems: bin\rman target *Target_Database_Login* catalog *Recovery_Catalog_Login*

UNIX systems: bin/rman target *Target_Database_Login* catalog *Recovery_Catalog_Login*

HP OpenVMS systems:

1. Execute ORAUUSER.COM using `$_@OMNI$ROOT:[LOG]LOGIN.COM`.
2. Execute `$_rman target target_connect_string catalog catalog_connect_string`.

Examples of the RMAN scripts

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

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

Backing up a single channel

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

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

Backing up three channels in parallel

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

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

Backing up all archived logs and tablespaces

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

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

Backing up particular archived logs

To back up all archived redo logs from sequence #5 to sequence #105 and delete the archived redo logs after backup of the instance named ora1 is complete, execute the following script:

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

If the backup fails, the logs are not deleted.

Backing up the flash recovery area

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

```
run {
allocate channel 'dev_0' type 'sbt_tape' parms
'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';
allocate channel 'dev_1' type 'sbt_tape' parms
'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';
```

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

Including control file in a backup specification

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

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

Backing up while allowing for some corrupted blocks

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

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

On UNIX systems

```
run {
set maxcorrupt for datafile
'/oracle/data1.dbs' to 10;
allocate channel 'dev_0' type 'sbt_tape' parms
  'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
  ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';
allocate channel 'dev_1' type 'sbt_tape' parms
  'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
  ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';
allocate channel 'dev_2' type 'sbt_tape' parms
  'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
  ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORACL,OB2BARLIST=ora1)';
```

```
backup
incremental level 0
format 'ora1<ORACL_%s:%t>.dbf'
database;
}
```

On Windows systems

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

Restore

You can restore the database objects using:

- Data Protector GUI. See [Restoring Oracle using the Data Protector GUI on page 65](#).
- RMAN. See [Restoring Oracle using RMAN on page 82](#).

Restorable items

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

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

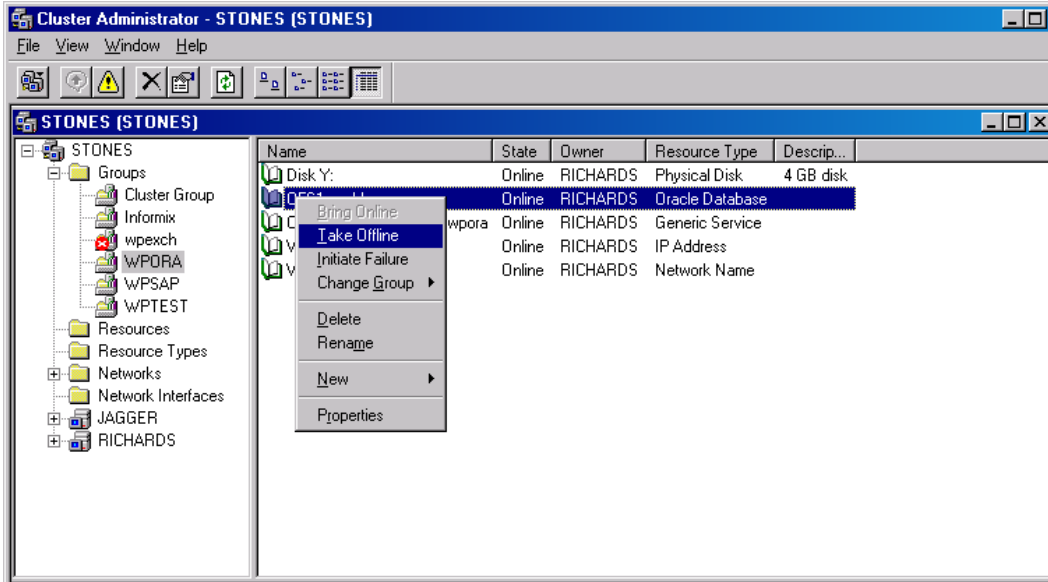
Duplicating databases

Using the Data Protector GUI, you can also **duplicate** a production database. See [Duplicating an Oracle database on page 75](#).

Microsoft Cluster Server systems

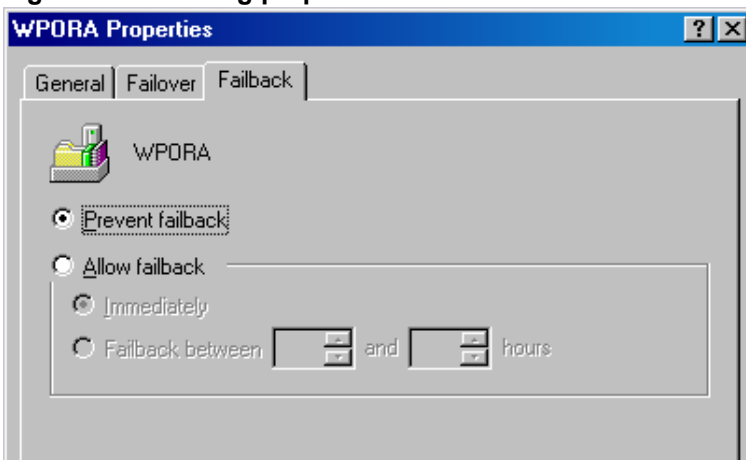
Before you start restoring a cluster-aware Oracle server, take the Oracle Database resource offline using, for example, the Cluster Administrator utility.
See [Taking the Oracle resource group offline below](#).

Figure 19: Taking the Oracle resource group offline



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

Figure 20: Checking properties



HP Serviceguard systems

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

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

Prerequisites

- An instance of Oracle must be created on the system to which you want to restore or duplicate the database.
- The database must be in the `Mount` state if the whole database is being restored, or in the `NoMount` state if the control file is being restored or a database duplication is performed.
- You must be able to connect to the database.

Restoring Oracle using the Data Protector GUI

For restore, RMAN scripts are generated with necessary commands, depending on selections made in the GUI. To use additional commands, use them manually from RMAN itself. You can also use the workaround described in [How to modify the RMAN restore script on page 109](#).

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:

1. Restore the control file from automatic backup.

If no automatic backup of the control file is available, see [The Recovery Catalog was lost and the control file cannot be restored from Data Protector managed backup on page 109](#).

2. Restore the database or data items.

Changing the database state

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

Table 6: Required database states

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

To put the database into the correct state, execute:

```
sqlplus /nolog
```

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

```
SQL>shutdown immediate;
```

To put the database into NoMount state, execute:

```
SQL>startup nomount;
```

To put the database into Mount state, execute:

```
SQL>startup mount;
```

Restoring the recovery catalog database

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

To restore the recovery catalog database:

1. Ensure that the recovery catalog database is in the **Open** state.
2. Remove the recovery catalog from the database (if it exists), using the RMAN command `DROP CATALOG`.

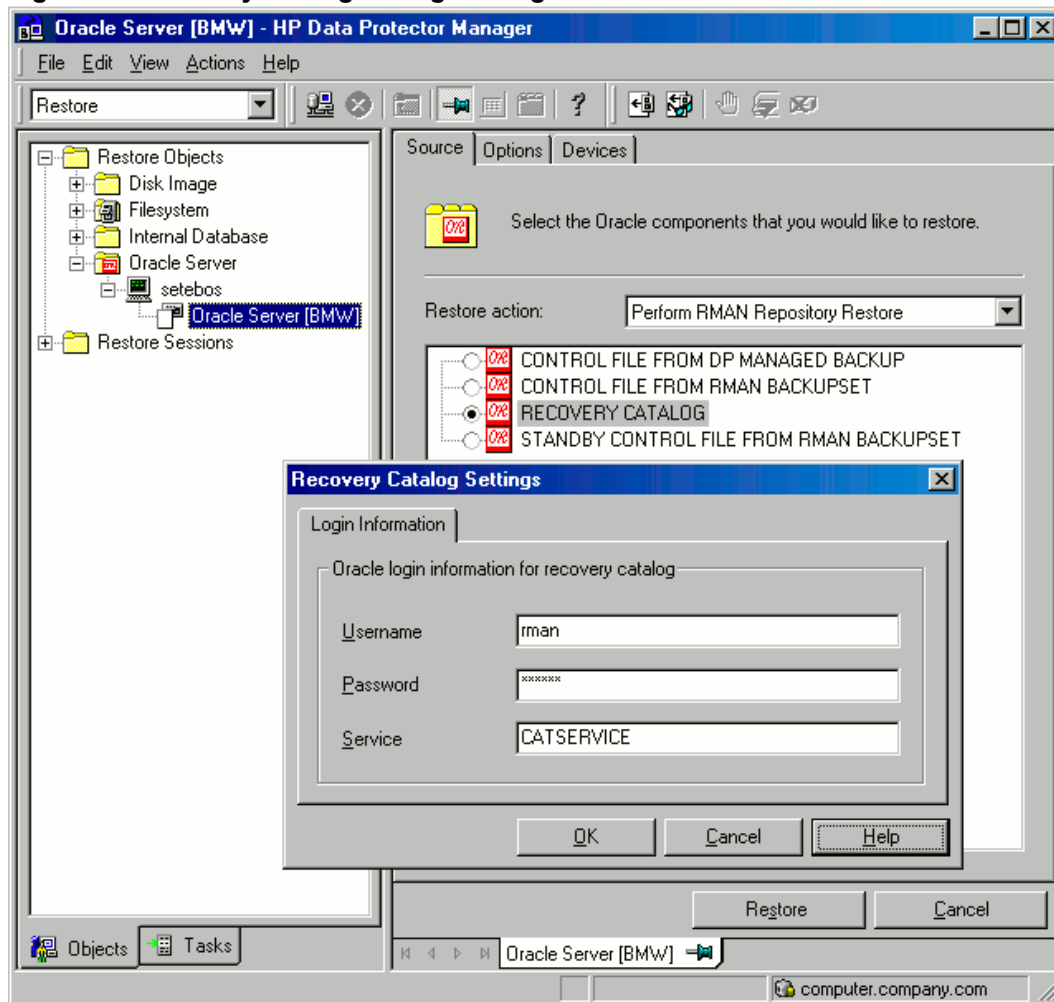
¹When restoring only a few tablespaces or datafiles, then the database can be open with the tablespaces or datafiles to be restored offline.

3. In the Data Protector GUI, switch to the **Restore** context.
4. Under **Restore Objects**, expand **Oracle Server**, expand the system on which the database, for which you want to restore the recovery catalog, resides, and then click the database.
5. In the **Restore action** drop-down list, select **Perform RMAN Repository Restore**.

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

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

Figure 21: Recovery catalog settings dialog



6. In the **Options** page:

In **User name** and **User group**, specify the user name and password to the recovery catalog database.

From the **Session ID** drop-down list, select the Session ID.

For further information, see [Restore, recovery, and duplicate options on page 78](#).

7. Click **Restore**.

Proceed to restore the control file.

Restoring the control file

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

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

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

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

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

The control files (`ctrlDB_NAME.dbf`) are restored to the default Data Protector temporary files directory.

Note: In Oracle Real Application Clusters (RAC) environments with Oracle versions 11.2.0.2 and later, the control files are created at, backed up from, and restored to the location defined by the `OB2_DPMCTL_SHRLOC` variable. This directory must reside on a shared disk and be accessible from all RAC nodes in order for restore sessions to succeed.

After the restore, execute the following script:

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

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

- Restoring from RMAN autobackup (CONTROLFILE FROM RMAN AUTOBACKUP)

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

Important: Ensure that you have properly configured the RMAN autobackup and that the correct backup version is available. If the RMAN autobackup session is not found during the restore, the

procedure is aborted. See the Oracle documentation on how to set up RMAN AUTOBACKUP.

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

The recovery catalog *is* required.

- **Oracle Data Guard:** Restoring standby control file from RMAN backup set (STANDBY CONTROL FILE FROM RMAN BACKUPSET)

If you restore a *standby* database (not using duplication), you must restore this type of control file.

This type of restore is available only in standby configurations if you selected the **CONTROL FILE FOR STANDBY** database object in the backup specification.

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

To restore the control file:

1. Open the `sqlplus` window and put the database in the nomount state. See [Changing the database state on page 66](#).
2. In the Data Protector GUI, switch to the **Restore** context.
3. Under **Restore Objects**, expand **Oracle Server**, expand the system on which the database, for which you want to restore the control file, resides, and then click the database.
4. In the **Restore Action** drop-down list, select **Perform RMAN Repository Restore**.

In the Results area, select the control file for restore.

5. In the **Options** page, from the **Client** drop-down list, select the system on which the Data Protector Oracle integration agent (`ob2rman.pl`) will be started. To restore the control file to a different database than it is selected, click **Settings** and specify the login information for the target database.

Set the other restore options. For information, see [Restore, recovery, and duplicate options on page 78](#).

6. Click **Restore**.

Proceed with restoring the Oracle database objects.

Restoring Oracle database objects

Before you restore Oracle database objects, ensure that you have an up-to-date version of the recovery catalog database and the control file. They contain the database structure information. If you do not have up-to-date versions of these files, restore them as described in [Restoring the recovery catalog database on page 66](#) and [Restoring the control file on the previous page](#).

To restore Oracle database objects:

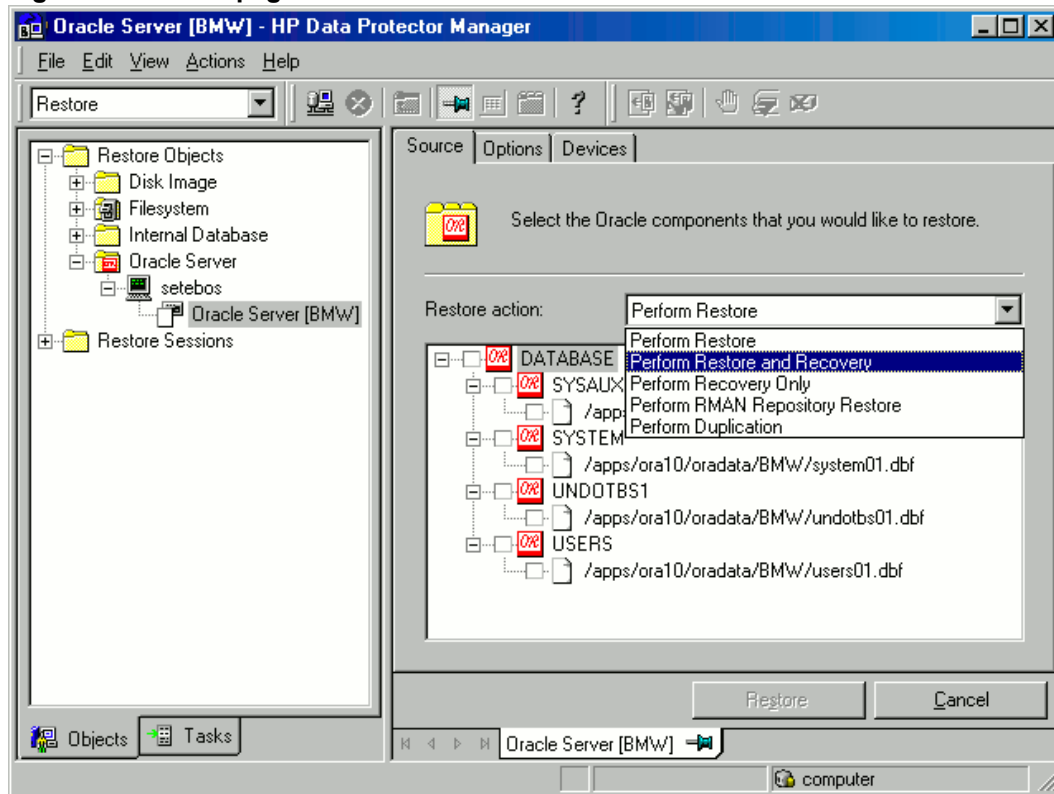
1. In Oracle Data Guard environments, if you restore a *standby* database, stop the managed recovery process (log apply services):

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

2. Put the database in the mount state. See [Changing the database state on page 66](#).
3. In the Data Protector GUI, switch to the **Restore** context.
4. Under **Restore Objects**, expand **Oracle Server**, expand the system on which the database, for which you restore the database objects, resides, and then click the database.
5. In the **Restore action** drop-down list, select the type of restore you wish to perform. For information on the options, see [Restore, recovery, and duplicate options on page 78](#).

Important: If you do not select **Perform Restore and Recovery** or **Perform Recovery Only**, you will have to recover the database objects manually using RMAN. For information, see [Restoring Oracle using RMAN on page 82](#).

Figure 22: Source page



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

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

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

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

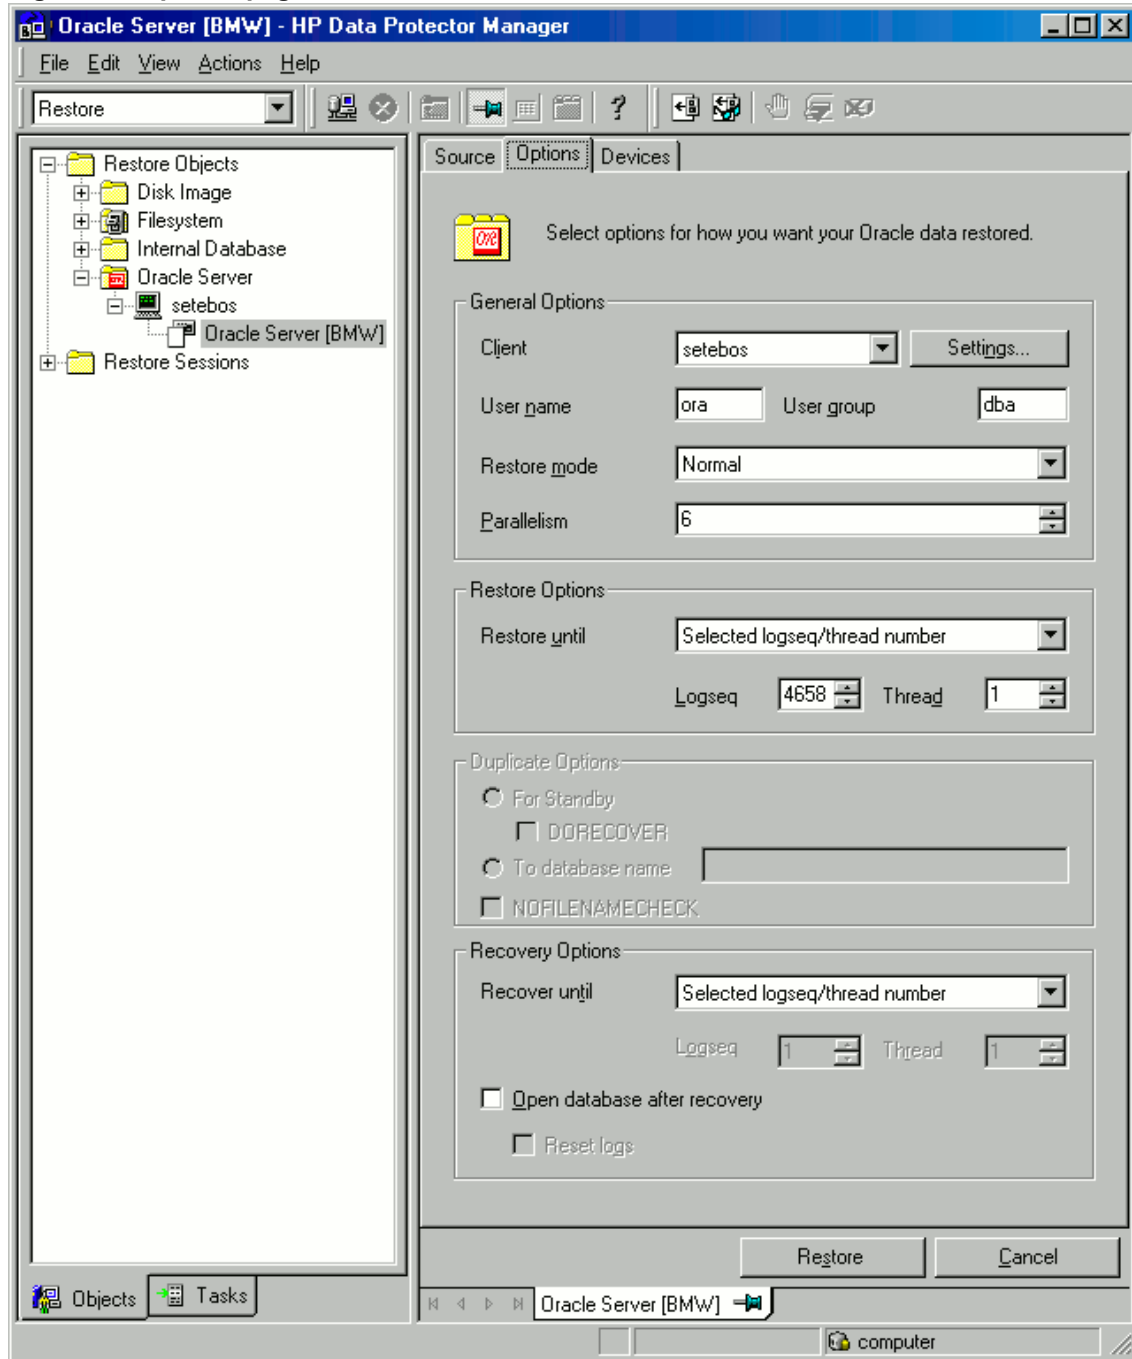
Tip: The same can be done if you set the `DB_FILE_NAME_CONVERT` initialization parameter. This parameter captures all the target datafiles and converts them appropriately.

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

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

Set the other restore options. For information, see [Restore, recovery, and duplicate options on page 78](#).

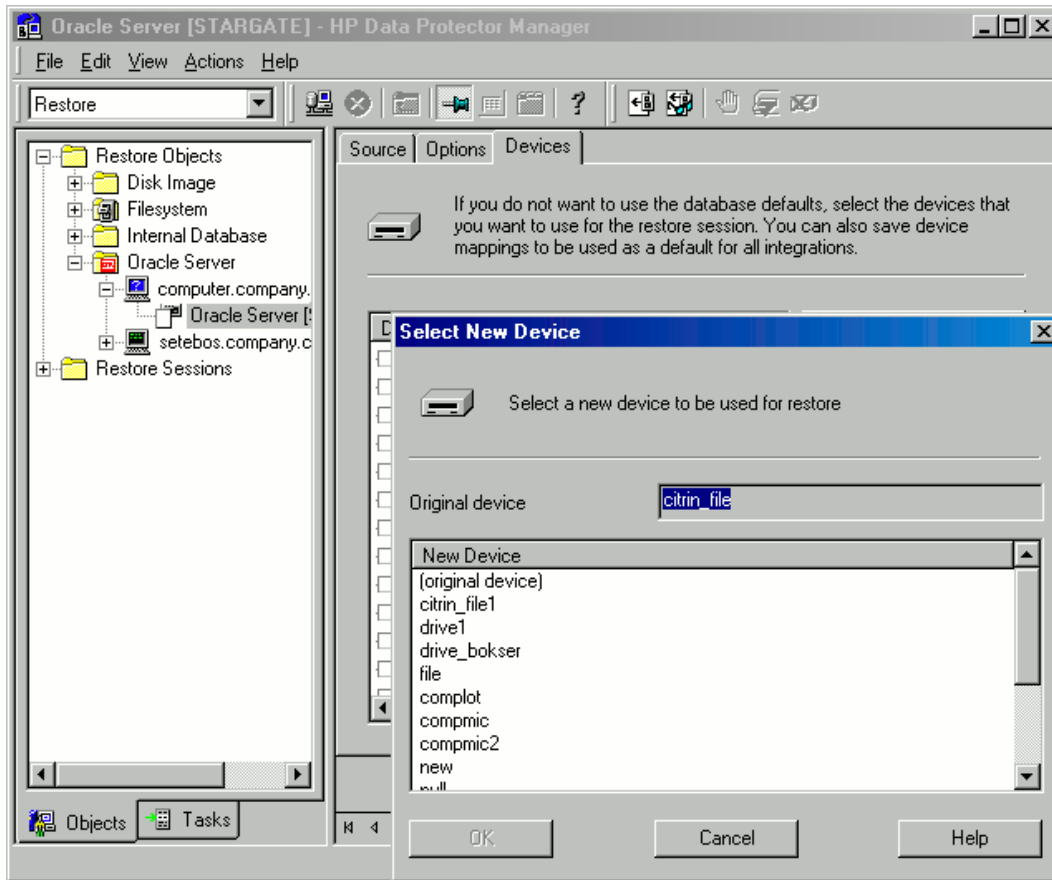
Figure 23: Options page



8. In the **Devices** page, select the devices to be used for the restore.

For more information on how to specify devices for a restore, see the *HP Data Protector Help* index: "restore, selecting devices for".

Figure 24: Devices page



9. Click **Restore**.

After the restore:

1. Put the database in the correct state.

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

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

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

```
rman target Target_Database_Login catalog Recovery_CatLog_Login
```

```
RMAN> RESET DATABASE;
```

```
RMAN> exit
```

3. If you did not choose to use Data Protector to recover the database objects and if you have all

archived redo logs on disk, perform the following after the database is restored:

Open a command line window and enter the following commands:

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

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

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

Restoring tablespaces and datafiles

To restore tablespaces and datafiles:

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

```
sqlplus /nolog  
  
SQL>connect user/password@service as sysdba;  
  
SQL>alter database datafile 'datafile name' offline;
```

If you are restoring a tablespace enter:

```
SQL>alter tablespace tablespace_name offline;
```

2. When the restore has been completed put the datafiles and tablespaces back online with the following procedures:

Open a command line window and enter the following commands:

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

If you are restoring a tablespace enter:

```
SQL>alter tablespace tablespace_name online;
```

Restoring and recovering an Oracle database in Oracle Data Guard environment

Restoring and recovering a primary database

You can restore and recover a primary database from backups done on either a primary or standby database. The restore and recover is almost the same as restore and recover of a database in a standalone configuration. For information, see [Restoring Oracle using the Data Protector GUI on page 65](#).

Restoring and recovering a standby database

You can restore and recover a standby database from backups of either a primary or standby database. The restore and recover is almost the same as restore and recover of a database in a standalone configuration. For information, see [Restoring Oracle using the Data Protector GUI on page 65](#).

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

Obtain UNTIL_SCN:

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

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

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

Perform duplication of the database also when:

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

Duplicating an Oracle database

Perform a production database duplication to create:

- A standby database which has the same DBID as the production (primary) database. With this, you can:
 - Create a new standby database.
 - Re-create a standby database after:

- Loss of entire standby database
 - Primary database control file was restored or recreated
 - Database point-in-time recovery was performed on the primary database
 - Switchover or failover of database roles occurred
- An independent copy, with a unique DBID, which can be used for data mining or testing purposes.

Prerequisites

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

Limitations

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

To duplicate a production database:

1. On the system where the selected database will be duplicated, put the Oracle auxiliary database instance in the nomount state. See [Changing the database state on page 66](#).
2. In the Context List of the Data Protector GUI, click **Restore**.
3. Under **Restore Objects**, expand **Oracle Server**, expand the system on which the production database resides, and then click the production database which you want to duplicate. If there are several such systems, select the system on which you want the Data Protector Oracle integration agent (ob2rman.pl) to be started.
4. In the **Restore Action** drop-down list, select **Perform Duplication**.
5. In the **Options** page, from the **Client** drop-down list, select the system on which the Data Protector

Oracle integration agent (ob2rman.pl) will be started.

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

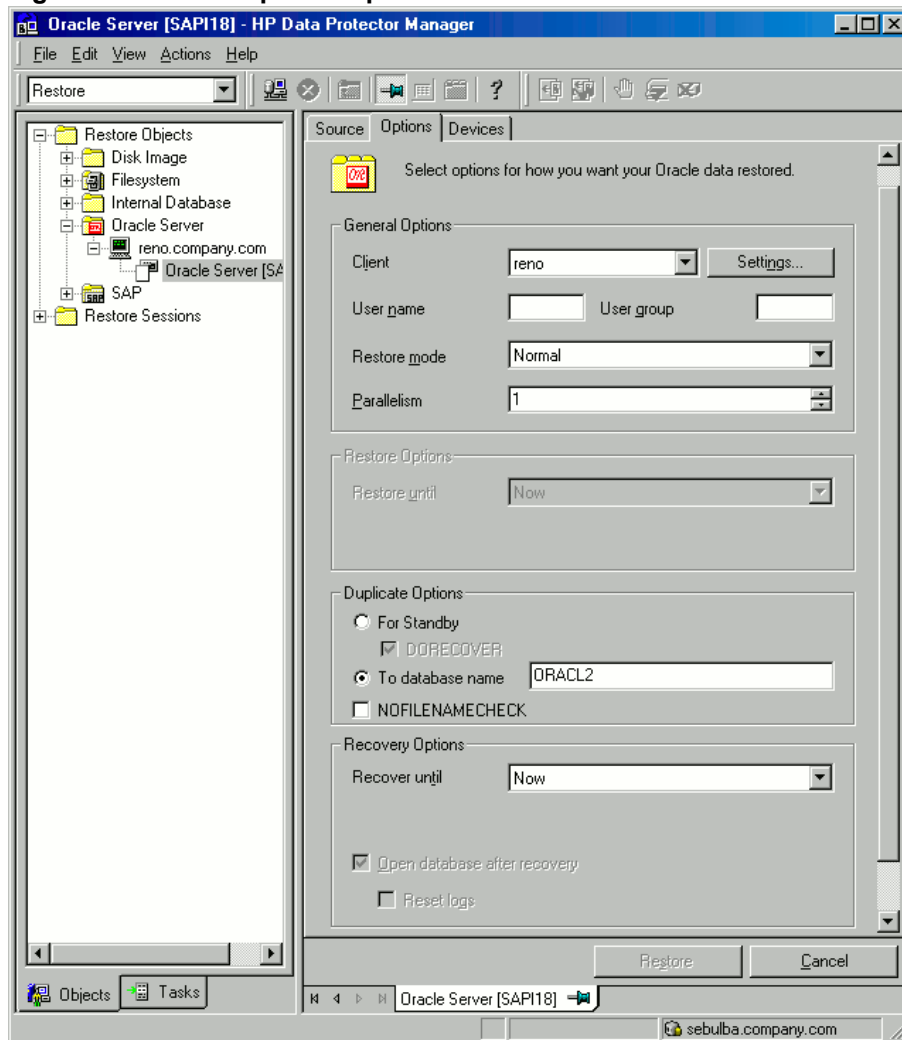
In **User name** and **User group**, specify the user name and group for the OSDBA account, which will be used by the Data Protector Oracle integration agent.

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

Set duplicate options. For information, see [Duplicate options on page 79](#) or press **F1**.

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

Figure 25: Oracle duplicate options



6. Click **Restore**.

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

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

Restore, recovery, and duplicate options

Restore action options

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

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

Perform Restore	Use this option to only restore (but not recover) the database objects using Data Protector. After restore, recover the database manually using RMAN. For information on recovering the database using RMAN, see Restoring Oracle using RMAN on page 82 .
Perform Restore and Recovery	Use this option to perform both the restore and recovery of the database objects using Data Protector.
Perform Recovery Only	Use this option to only recover the database. This action can only be performed on the whole database.
Perform RMAN Repository Restore	Use this option to restore the recovery catalog or the control file when the database objects are not available in the Source page.
Perform Duplication	This option is used to perform duplication of a production database. This action can only be performed on the whole database.

General options

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

Settings	<p>Click Settings to specify the login information (user name, password, and net service name) for the target database (in case of restore and recovery) or auxiliary database (in case of duplication) where you want the selected database objects to be restored or duplicated.</p> <p>If this is not specified in the case of restore or recovery, the login information of the selected database that resides on the selected system will be used.</p> <p>If this is not specified in the case of duplication, the duplication session will fail.</p>
User name, User group (UNIX systems only)	<p>Specify the operating system user account under which you want the restore to start.</p> <p>Ensure that this user has Oracle rights to restore the database (for example, it is in the DBA user group). The user must also be in the Data Protector <code>admin</code> or <code>operator</code> user group (actually, the <code>Start restore</code> and <code>See private objects</code> user rights suffice).</p>
Restore mode	<p>This drop-down list allows you to specify which type of restore you would like perform. The options are:</p> <ul style="list-style-type: none"> • Normal <p>This option should be used when a conventional backup or ZDB using the backup set method was performed.</p> • Proxy copy <p>This option should be used when the original Oracle backup was made using the Oracle RMAN proxy-copy method.</p> <p>This option is disabled when you perform recovery only.</p>
Parallelism	<p>This field is used to specify the number of concurrent data streams that can read from the backup device. The default value is one.</p> <p>In case of Normal restore mode, to optimize restore performance, specify the same number of data streams as were used during the backup. For example, if you set the backup concurrency to 3, set the number of parallel data streams to 3 as well. Note that if a very high number of parallel data streams is specified this may result in a resource problem because too much memory is being used.</p>

Duplicate options

Available if **Perform Duplication** was selected.

For Standby	<p>Select this option to create a standby database.</p> <p>Default: selected.</p>
--------------------	---

DORECOVER	<p>Available if For Standby was selected.</p> <p>Select this option if you want RMAN to recover the database after creating it.</p>
To database name	<p>Select this option to create a new database copy. In the text box, specify its name. The name should match the name in the initialization parameter file that was used to start the auxiliary database instance. By default, the database name is set to the database name of the currently selected target database.</p>
NOFILENAMECHECK	<p>Select this option to disable RMAN to check whether the target datafiles share the same names with the duplicated datafiles.</p> <p>Select this option when the target datafiles and duplicated datafiles have the same names, but reside on different systems.</p> <p>Default: not selected.</p>

Restore and recovery options

Restore until	<p>The options in this drop-down list allow you to limit the selection to those backups that are suitable for an incomplete recovery to the specified time.</p> <ul style="list-style-type: none"> • Now <p>Use this option to restore the most recent full backup. By default, this option is selected.</p> • Selected time <p>Use this option to specify an exact time to which you wish the database to be restored. Data Protector restores the backup that can be used in recovery to the specified time.</p> • Selected logseq/thread number <p>A logseq number is a redo log sequence number. Use this option to specify a particular redo log sequence and a thread number which will act as an upper limit of redo logs to restore. Data Protector restores the backup that can be used in recovery to the specified log sequence number.</p> • Selected SCN number <p>Use this option to specify the SCN number to which you wish the database to be restored. Data Protector restores the backup that can be used in recovery to the specified SCN number.</p>
----------------------	---

Recover until	<p>The options in this drop-down list allow you to specify to which point in time you would like the recovery to be performed.</p> <ul style="list-style-type: none">• Now Data Protector starts RMAN to recover the database to the most recent time possible by applying all archived redo logs. By default, this option is selected.• Selected time Use this option to specify an exact time to which the archive logs are applied.• Selected logseq/thread number A logseq number is a redo log sequence number. Use this option to specify a particular redo log sequence and a thread number which will act as an upper limit of redo logs to recover.• Selected SCN number Use this option to specify the SCN number to which you perform the recovery. <p>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 execute:</p> <pre>rman target <i>Target_Database_Login</i> catalog <i>Recovery_CatLog_Login</i> RMAN> RESET DATABASE; RMAN> exit</pre>
Open database after recovery	<p>Opens the database after a recovery is performed.</p>

Reset logs	<p>Resets the archive logs after the database is opened.</p> <p><i>Always reset the logs:</i></p> <ul style="list-style-type: none">• After an incomplete recovery (not Recover until now).• If a backup of a control file is used in recovery or restore and recovery. <p><i>Do not reset the logs:</i></p> <ul style="list-style-type: none">• After a complete recovery (Recover until now) when the backup of a control file was not used in recovery or restore and recovery.• On the primary database, if the archive logs are used for a standby database. However, if you must reset the archive logs, you will need to recreate the standby database. <p>If you reset the logs when the Recover until option is set to Now, a warning is displayed, stating that you should reset the logs only if you use an older control file for restore.</p> <div data-bbox="386 800 1403 905" style="background-color: #f0f0f0; padding: 5px;"><p>Note: Oracle recommends that you perform a complete backup immediately after a database was opened with the Reset Logs option.</p></div>
-------------------	---

Restoring Oracle using RMAN

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

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

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

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

The following examples of restore are given:

- [Example of full database restore and recovery on page 85](#)
- [Example of point-in-time restore on page 86](#)
- [Example of tablespace restore and recovery on page 87](#)
- [Example of datafile restore and recovery on page 89](#)
- [Example of archive log restore on page 91](#)

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

Oracle database you are using. For detailed steps on how to perform the restore of control files, see the *Recovery Manager User's Guide and References*.

Preparing the Oracle database for restore

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

Prerequisites

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

- If you use the recovery catalog database, make sure that the recovery catalog database is open. If the recovery catalog database cannot be brought online, you will probably need to restore the recovery catalog database. See [Restore on page 63](#) for details of how to restore the recovery catalog database.
- Control files must be available. If the control files are not available, you must restore them. See the *Oracle Recovery Manager User's Guide and References* for more details.

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

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

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

Windows systems example

```
ORACLE_BASE=Oracle_home  
ORACLE_HOME=Oracle_home\product\10.1.0  
ORACLE_TERM=HP  
DB_NAME=PROD  
PATH=$PATH:Oracle_home\product\10.1.0\bin
```

```
NLS_LANG=american  
NLS_DATE_FORMAT='Mon DD YYYY HH24:MI:SS'
```

UNIX systems example

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

HPOpenVMS systems example

```
ORACLE_HOME=DKA400:[ORACLE10G]  
ORACLE_TERM=HP  
DB_NAME=PROD
```

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

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

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

HP OpenVMS systems with Oracle 10g:

```
Oracle_home /oratab  
TEST:/DKA400/ORACLE10G:N CAT:/DKA400/ORACLE10G:N
```

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

Connection strings used in the examples

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

- Target connection string for target database:

```
sys/manager@PROD
```

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

- Recovery catalog connection string for recovery catalog database:

```
rman/rman@CATAL
```

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

SBT_LIBRARY parameter

On Windows and UNIX systems, set the `SBT_LIBRARY` RMAN script parameter to point to the correct platform-specific Data Protector MML. The parameter must be specified for each RMAN channel separately. For details on the Data Protector MML location, see [Specify the parms operand: on page 58](#).

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

Example of full database restore and recovery

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

1. Log in to the Oracle RMAN:

If you use the recovery catalog database, execute:

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

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

```
HP OpenVMS systems: rman target sys/manager@PROD sys/manager@PROD catalog  
rman/rman@CAT
```

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

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

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

```
HP OpenVMS systems: rman target sys/manager@PROD nocatalog
```

2. Start the full database restore and recovery:

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

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

1. Create a `restore_datafile` file in the default Data Protector temporary files directory.
2. Start the full database restore:

If you use the recovery catalog database, execute:

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

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

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

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

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

Example of point-in-time restore

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

1. Log in to the Oracle RMAN:

If you use the recovery catalog database, execute:

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

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

HP OpenVMS systems: `rman target sys/manager@PROD sys/manager@PROD catalog rman/rman@CAT`

If you do not use the recovery catalog, execute:

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

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

HP OpenVMS systems: `rman target sys/manager@PROD nocatalog`

2. Start the point-in-time restore:

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

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

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

1. Create a restore_PIT file in the default Data Protector temporary files directory.
2. Start the point-in-time restore:

If you use the recovery catalog database, execute:

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

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

If you do not use the recovery catalog, execute:

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

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

Example of tablespace restore and recovery

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

1. Log in to the Oracle RMAN:

If you use the recovery catalog database, execute:

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

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

HP OpenVMS systems: `rman target sys/manager@PROD sys/manager@PROD catalog
rman/rman@CAT`

If you do not use the recovery catalog, execute:

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

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

HP OpenVMS systems: `rman target sys/manager@PROD nocatalog`

2. Start the tablespace restore and recovery.

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

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

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

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

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

1. Create a restore_TAB file in the default Data Protector temporary files directory.
2. Start the tablespace restore.

If you use the recovery catalog database, execute:

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

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

If you do not use the recovery catalog, execute:

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

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

Example of datafile restore and recovery

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

To restore and recover a datafile:

1. Log in to the Oracle RMAN.

If you use the recovery catalog database, execute:

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

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

HP OpenVMS systems: `rman target sys/manager@PROD sys/manager@PROD catalog rman/rman@CAT`

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

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

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

HP OpenVMS systems: `rman target sys/manager@PROD nocatalog`

2. Start the datafile restore and recovery:

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

UNIX systems

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

```
    '/opt/oracle/data/oradata/DATA/temp01.dbf';  
recover datafile  
    '/opt/oracle/data/oradata/DATA/temp01.dbf';  
sql "alter database datafile  
    '/opt/oracle/data/oradata/DATA/temp01.dbf' online";  
release channel dev1;  
}
```

Windows systems

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

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

UNIX system

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

Windows system

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

```
recover datafile
  'Oracle_home\data\oradata\DATA\temp01.dbf';
release channel dev1;
}
```

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

1. Create a restore_dbf file in the default Data Protector temporary files directory.
2. Start the datafile restore:

If you use the recovery catalog database, execute:

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

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

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

Windows systems: `ORACLE_HOME/bin\rman target sys/manager@PROD nocatalog cmdfile=Data_Protector_home\tmp\restore_dbf`

UNIX systems: `ORACLE_HOME/bin/rman target sys/manager@PROD nocatalog cmdfile=/var/opt/omni/tmp/restore_dbf`

Example of archive log restore

To restore an archive log:

1. Log in to the Oracle RMAN:

If you use the recovery catalog database, execute:

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

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

HP OpenVMS systems: `rman target sys/manager@PROD sys/manager@PROD catalog rman/rman@CAT`

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

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

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

HP OpenVMS systems: `rman target sys/manager@PROD nocatalog`

2. Start the archive log restore:

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

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

1. Create a restore_arch file in the default Data Protector temporary files directory.
2. Start the archive log restore:

If you use the recovery catalog database, execute:

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

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

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

Windows systems: `ORACLE_HOME\bin\rman target sys/manager@PROD nocatalog cmdfile=Data_Protector_home\tmp\restore_arch`

UNIX systems: `ORACLE_HOME/bin/rman target sys/manager@PROD nocatalog cmdfile=/var/opt/omni/tmp/restore_arch`

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

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

1. Log in to the Oracle RMAN:

Windows systems: `ORACLE_HOME\bin\rman target sys/manager@TIN`

UNIX systems: `ORACLE_HOME/bin/rman target sys/manager@TIN`

HP OpenVMS systems: `rman target sys/manager@TIN`

2. Execute:

```
run {
allocate channel 'dev_0' type 'sbt_tape'
parms 'SBT_LIBRARY=C:/PROGRA~1/OmniBack/bin/orasbt.dll,
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=CAN,OB2BARLIST=test)';
```

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

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

Restoring using another device

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

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

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

where

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

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

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

Example

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

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

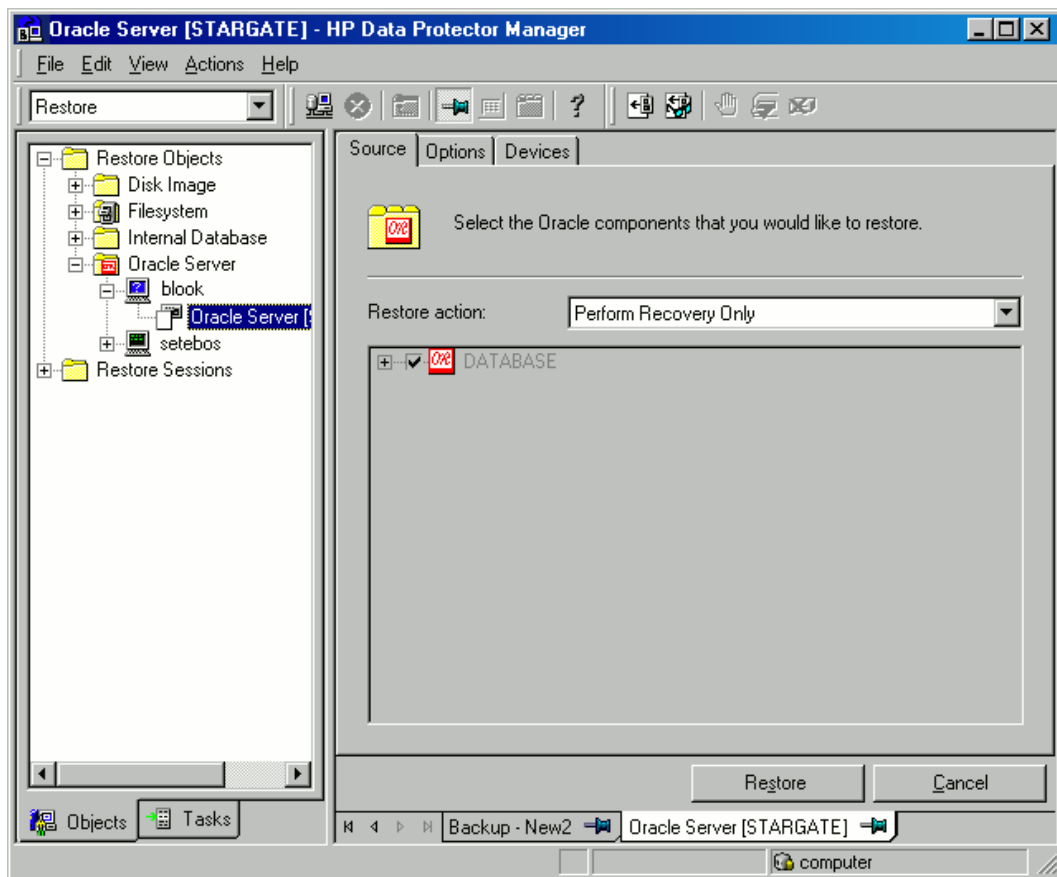
Disaster recovery

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

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

This is a general procedure for recovering an application:

1. Complete the recovery of the operating system.
2. Install, configure, and initialize the database/application so that data on the Data Protector media can be loaded back to the system. Consult the documentation from the database/application vendor for a detailed procedure and the steps needed to prepare the database.
3. Ensure that the database/application server has the required Data Protector client software installed and is configured for the database/application. Follow the procedures in this chapter and in the section. See also the section of this manual about the Data Protector Restore GUI for Oracle for information about using this to restore database items, [Restoring Oracle using the Data Protector GUI on page 65](#).
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.



1. ■

Monitoring sessions

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

Monitoring current sessions

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

1. In the Context List, click **Monitor**.

In the Results Area, all currently running sessions are listed.

2. Double-click the session you want to monitor.

Clearing sessions

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

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

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

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

Monitoring tools

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

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

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

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 `oracle8.log` file, located on the Oracle Server system in the default Data Protector log files directory.

Oracle Server writes the logs in the `Oracle_user_dump_directory\sbtio.log` file.

Resuming sessions

Backup and restore sessions that did not complete successfully can be resumed. The functionality enables you to back up or restore only the files that failed to be backed up or restored in the original session. Consequently, a session started using the resume session functionality (**resumed session**) generally takes less time to complete.

You can resume a session using the Data Protector GUI or CLI.

Resuming backup sessions

When you resume an Oracle Server integration backup session, Data Protector starts a new backup session using the same backup specification. Thus, you can create multiple sessions to finish a backup. Only the last of these sessions can be resumed, while the previous sessions that have already been resumed once, cannot be resumed again. Because the resume session functionality uses the same backup specification as the original session, the changes made to the backup specification between the abort and the resume, affect the resume session.

The main difference compared to a standard backup session is that, during a resumed session, Data Protector modifies the RMAN script before the actual backup is started, adding the clause `NOT BACKED UP SINCE Time` for each backup command, where *Time* is the original backup session start time. See the following example:

```
run{
    allocate channel 'dev_0' type 'sbt_tape'
    parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORCL,OB2BARLIST=New1)';
    allocate channel 'dev_1' type 'sbt_tape'
    parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORCL,OB2BARLIST=New1)';
    allocate channel 'dev_1' type 'sbt_tape'
    parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=ORCL,OB2BARLIST=New1)';
    backup incremental level <incr_level>          format 'New1<ORCL_%s:%t:%p>.dbf'
    NOT BACKED UP SINCE TIME "TO_DATE('5/15/2009 15:30:00',
'MM/DD/YY HH24:MI:SS')"
    database;
    sql 'alter system archive log current';
    backup
    format 'New1<ORCL_%s:%t:%p>.dbf'
    NOT BACKED UP SINCE TIME "TO_DATE('5/15/2009 15:30:00',
'MM/DD/YY HH24:MI:SS')"
    archive log all;
    backup
    format 'New1<ORCL_%s:%t:%p>.dbf'
    NOT BACKED UP SINCE TIME "TO_DATE('5/15/2009 15:30:00',
'MM/DD/YY HH24:MI:SS')"
    recovery area;
    backup
    format 'New1<ORCL_%s:%t:%p>.dbf'
    NOT BACKED UP SINCE TIME "TO_DATE('5/15/2009 15:30:00',
'MM/DD/YY HH24:MI:SS')"
    current controlfile;
```


Consequently, RMAN skips the backup sets that were backed up successfully in the original session.

Suppose that you run the following sessions:

1. 2009/05/13-1 (original backup session)
2. 2009/05/13-2 (resuming 2009/05/13-1)
3. 2009/05/13-3 (resuming 2009/05/13-2)

The *Time* in the RMAN clause `NOT BACKED UP SINCE Time` is always the original backup session start time. Consequently, the RMAN script created in the third session (2009/05/13-3) does not use the start time of the session 2009/05/13-2 but the start time of the original backup session (2009/05/13-1). This ensures that each backup set is backed up only once after the original backup session was started.

Note: Ensure that the Cell Manager and your Oracle Server system are synchronized. Otherwise, if *Time* is not correct, the resume session functionality does not work properly.

Note: The smallest backup unit is a backup set. Therefore, consider the following for the RMAN option `FILESERSET`:

- If the option is set to 1, RMAN creates a separate backup set for each file. In this case, you benefit the most from the resume session functionality. However, note that restore is significantly prolonged if files are backed up with many streams.
- If RMAN creates only one backup set for the files to be backed up and some files fail to be backed up, the whole backup set fails. When you resume such a session, the whole backup set is backed up again, including the files that were backed up successfully.

Resuming restore sessions

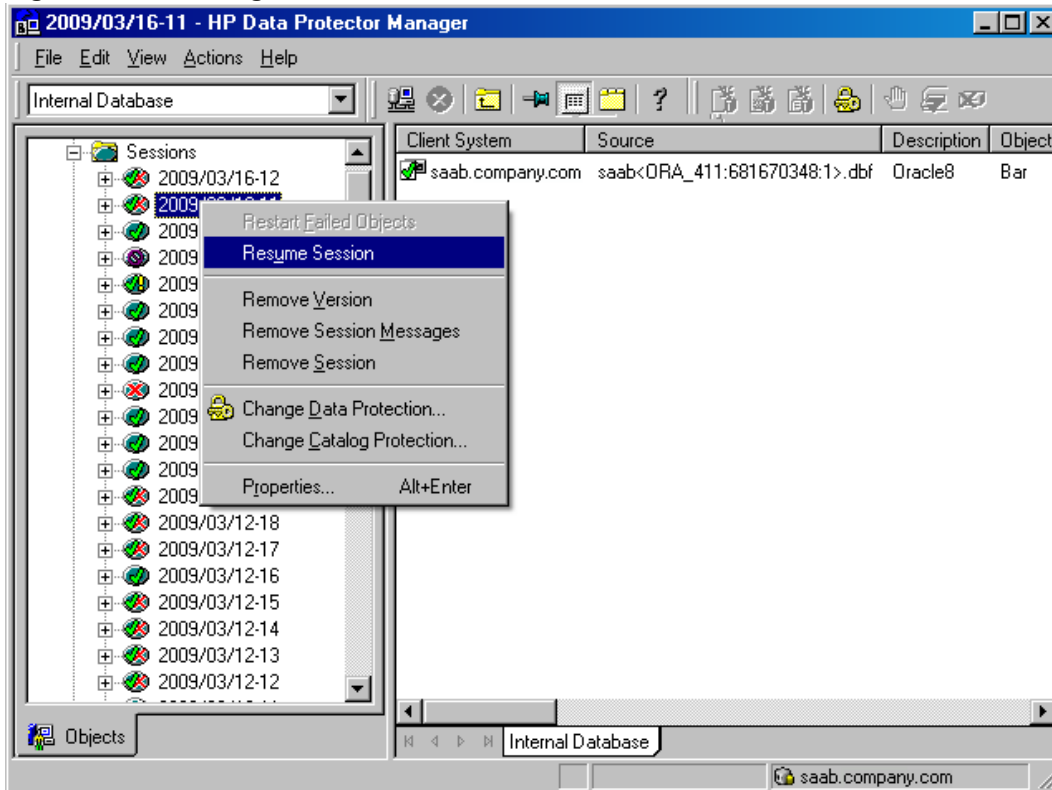
The main benefit of resuming a restore session is that you do not need to specify, all over again, what to restore, which devices to use, and so on. However, in reality, there is no difference between a standard restore session and a resumed restore session. In both cases, Oracle Server first checks if files to be restored already exist at the target location and then restores only the missing ones.

Note: Once you open the Oracle database with the `RESETLOGS` option, it is pointless to use the resume session functionality for sessions that restored old backups (backups created before the logs were reset).

Using the Data Protector GUI

1. In the Internal Database context, expand **Sessions**.
2. Right-click the session that you want to resume and click **Resume Session**. See [Resuming a session on the next page](#).

Figure 26: Resuming a session



Using the Data Protector CLI

1. Log on to the Cell Manager or to any system with the User Interface component installed.
2. Go to the following directory:

Windows systems: `Data_Protector_home\bin`

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

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

3. To resume a backup session, execute:

```
omnib -resume SessionID
```

To resume a restore session, execute:

```
omnir -resume SessionID
```

For details, see the `omnib` and `omnir` man pages or the *HP Data Protector Command Line Interface Reference*.

Example

To resume the backup session `2013/05/13-1`, execute:

```
omnib -resume 2013/05/13-1
```

Aborting sessions

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

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

For details of how to set environment variables, see [Setting environment variables on page 35](#). For details of how to set the corresponding `omnirc` options, see the *HP Data Protector Help* index: "omnirc option". Note that environment variables override `omnirc` options.

Oracle RMAN metadata and Data Protector Media Management Database synchronization

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

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

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

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

```
allocate channel for maintenance type 'sbt_tape' parms 'SBT_LIBRARY=Path_to_Data_Protector_MML, ENV=(OB2MAINTENANCE=1)';  
  
crosscheck backup completed after "TO_DATE('01/13/06 10:30:00', 'MM/DD/YY HH24:MI:SS')"  
";  
  
release channel;
```

The `SBT_LIBRARY` parameter should be specified only on UNIX and Windows systems.

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, execute the following command using RMAN:

```
delete expired backup;
```

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

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

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

Troubleshooting

This section contains a list of general checks and verifications and a list of problems you might encounter when using the Data Protector Oracle integration. You can start at [Problems on page 107](#) and if you cannot find a solution there, perform general checks and verifications.

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

Before you begin

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

Checks and verifications

For more detailed information on performing any of the following procedures, see the Oracle documentation.

If your configuration, backup, or restore failed:

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

- a. Perform the following:

Windows systems: Set the `ORACLE_HOME` and `DB_NAME` variables.

UNIX systems: Export the `ORACLE_HOME` and `DB_NAME` variables as follows:

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

```
ORACLE_HOME="ORACLE_HOME"
```

```
export ORACLE_HOME
```

```
DB_NAME="DB_NAME"
```

```
export DB_NAME
```

- If you are using a `csch`-like shell, enter the following commands:

```
setenv ORACLE_HOME "ORACLE_HOME"
```

```
setenv DB_NAME "DB_NAME"
```

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

```
sqlplus /nolog
```

- c. Start SQL*Plus and type:

```
connect user_name/password@service as sysdba;
```

```
select * from dba_tablespace;
```

```
exit
```

If this fails, open the Oracle target database.

- Verify that you can access the recovery catalog (if used) and that it is opened as follows:
 - a. Export or set the `ORACLE_HOME` and `DB_NAME` variables as described in [Perform the following: on the previous page](#).

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

```
sqlplus /nolog
```

- c. Start SQL*Plus and type:

```
connect Recovery_CatLog_Login
```

```
select * from rcver;
```

```
exit
```

If this fails, open the recovery catalog.

- Verify that the listener is correctly configured for the Oracle target database and the recovery catalog database. This is required to properly establish network connections:
 - a. Export or set the `ORACLE_HOME` variable as described in [Perform the following: on the previous page](#).


- b. Start the listener from the `bin` directory in the `ORACLE_HOME` directory:

```
lsnrctl status service
```

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

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

Figure 27: 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.

- c. Start SQL*Plus from the bin directory in the *ORACLE_HOME* directory:

```
sqlplus /nolog
```

- d. Start SQL*Plus and type:

```
connect Target_Database_Login
```

```
exit
```

and then

```
connect Recovery_Catalog_Login
```

```
exit
```

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

- Verify that the Oracle target database and the recovery catalog database are configured to allow remote connections with the system privileges:
 - a. Export or set the *ORACLE_HOME* and *DB_NAME* variables as described in [Perform the following: on page 100](#).
 - b. Start SQL*Plus from the bin directory in the *ORACLE_HOME* directory:

```
sqlplus /nolog
```

- c. Start SQL*Plus and type:

```
connect Target_Database_Login as SYSDBA
```

```
exit
```

and

```
sqlplus connect Recovery_Catalog_Login as SYSDBA
```

```
exit
```

Repeat the procedure using SYSOPER instead of SYSDBA.

If this fails, see the Oracle documentation for instructions on setting up the password file and any relevant parameters in the `initDB_NAME.ora` file.

- ■ If you do not use the recovery catalog database:

Export or set the `ORACLE_HOME` and `DB_NAME` variables as described in [Perform the following: on page 100](#).

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

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

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

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

- a. Export or set the `ORACLE_HOME` variable as described in [Perform the following: on page 100](#).

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

```
sqlplus /nolog
```

- c. Start SQL*Plus and type:

```
connect Recovery_Catalog_Login;
```

```
select * from rc_database;
```

```
exit
```

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

- Verify backup and restore directly to disk using an RMAN channel type disk:

If you use the recovery catalog:

- a. Export or set the `ORACLE_HOME` variable as described in [Perform the following: on page 100](#).
- b. Start RMAN from the `bin` directory in the `ORACLE_HOME` directory:

```
rman target Target_Database_Login catalog Recovery_Catalog_Login cmd_file=rman_
script
```

If you do not use the recovery catalog:

- a. Export or set the `ORACLE_HOME` variable as described in [Perform the following: on page 100](#).
- b. Start RMAN from the `bin` directory in the `ORACLE_HOME` directory:

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

An example of the RMAN backup script is presented below:

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

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

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

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

Additionally, if your configuration or backup failed:

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

For details, see the *HP Data Protector Installation and Licensing Guide*.

- Check if the `SYSDBA` privilege is granted to the Oracle administrator.
- If you have special Oracle environment settings, ensure that they are entered in the Data Protector Oracle configuration files on the Cell Manager. For information on setting the variables in the Data

Protector Oracle configuration files, see the `util_cmd` man page or the *HP Data Protector Command Line Interface Reference*.

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

For details on performing filesystem backups, see the *HP Data Protector Help* index: “standard backup procedure”.

- **Windows systems:** Check the Data Protector Inet service startup parameters on the Oracle Server system:

Go to **Control Panel > Administrative Tools > Services > Data Protector Inet**.

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

- Examine the system errors reported in the following file on the Oracle Server system into the `debug.log` file.

Additionally, if your backup or restore failed:

- Test the Data Protector internal data transfer using the `testbar2` utility:
 - a. Verify that the Cell Manager name is correctly defined on the Oracle Server system. Check the `cell_server` file located in the default Data Protector client configuration directory, which contains the name of the Cell Manager system.
 - b. From the `bin` directory in the `ORACLE_HOME` directory, execute:

If backup failed:

```
testbar2 -type:Oracle8 -appname:DB_NAME-perform:backup -bar:backup_
specification_name
```

If restore failed:

```
testbar2 -type:Oracle8 -appname:DB_NAME-perform:restore -object:object_name -
version:object_version-bar:backup_specification_name
```

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

- c. You should see only `NORMAL` messages displayed on your screen, otherwise examine the errors reported by the `testbar2` utility by clicking the **Details** button in the Data Protector **Monitor** context.

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

- Check if the user under which the backup or restore session was started has appropriate Oracle permissions (for example, belongs to the `DBA` group). This user must also be in the Data Protector `operator` or `admin` user group.

- Check that the respective Data Protector user group has the `See private objects` user right enabled.
- **If backup failed:** Create an Oracle backup specification to back up to a null device or file. If the backup succeeds, the problem may be related to the backup devices. See the *HP Data Protector Troubleshooting Guide* for instructions on troubleshooting devices.
- **If restore failed:** Execute the `omnidb` command to see objects in the database.

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

Additionally, if your restore failed:

- Verify that an object exists on the backup media.

This can be done by executing the following command on the Oracle server system from the `bin` directory in the `ORACLE_HOME;` directory:

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

The output of the command lists detailed information about the specified Oracle object, as well as the session IDs of the backup sessions containing this object and a list of the media used. For detailed syntax of the `omnidb` command, see its man page.

- Ensure that the database is in the correct state.

If you are trying to restore a database item using the Data Protector GUI and the GUI stops responding, try one of the following:

- If you are restoring the control file, the database should be in the `NoMount` state.

Open a command window and enter the following:

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

- If you are restoring datafiles, the database should be in the `Mount` state.

Open a command window and enter the following:

```
sqlplus/nolog  
  
SQL>connect user/password@service as sysdba
```

```
SQL>shutdown immediate
```

```
SQL>startup mount
```

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

For information, see [Restoring Oracle using RMAN on page 82](#).

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

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

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

Open a SQLplus window and use the following command:

```
sqlplus/nolog
```

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

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

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

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

Problems

Problem

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

The following errors are listed in the Data Protector monitor:

```
EXP-00008: ORACLE error 6550 encountered
```

```
ORA-06550: line 1, column 13:
```

```
PLS-00201: identifier 'SYS.LT_EXPORT_PKG' must be declared
```

```
ORA-06550: line 1, column 7:
```

```
PL/SQL: Statement ignored
```

```
EXP-00083: The previous problem occurred when calling
```

```
SYS.LT_EXPORT_PKG.schema_info_exp
```

```
. exporting statistics
```

```
Export terminated successfully with warnings.
```

```
[Major] From: ob2rman.pl@machine "MAKI" Time: 10/01/01 16:07:53
```

```
Export of the Recovery Catalog Database failed.
```

Action

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

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

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

Restart the failed backup session.

Problem

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

Backup fails and the following error message is displayed:

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

Action

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

Problem

Backup fails after a point in time restore and recovery

The following error is displayed:

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

Action

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

```
rman target Target_Database_Login catalog Recovery_Catalog_Login
```

```
RMAN> RESET DATABASE;
```

```
RMAN> exit
```

Problem

Backup of archive logs on RAC cannot be performed

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

Action

Edit the archive logs backup specification:

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

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

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

Problem

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

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

Action

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

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

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

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

Problem

How to modify the RMAN restore script

When you start a restore of an Oracle database using the Data Protector GUI or CLI, an RMAN restore script is created, which is instantly run, so you cannot edit it first.

Action

To edit the script before it is run, set the Data Protector `omnirc` option `OB2RMANSAVE` to point to an existing directory. When the variable is set and you start a restore, the RMAN restore script, which is created at run time, is saved to the specified location under the name `RMAN_restore_backup_specification_name.rman`, and the actual restore is skipped. Then you can edit the script and run it manually afterwards. On how to set the `omnirc` options, see the *HP Data Protector Help* index: “omnirc options”.

To start a restore using Data Protector again, clear the `OB2RMANSAVE` option by deleting its content or commenting or removing the whole option. If you comment or remove the option on a Windows system, restart the Data Protector `Inet` service for the settings to take effect.

Problem

“IPC Invalid Hostname or IP Address” error message is displayed when browsing Oracle restore sessions

The following error message is displayed when browsing the Oracle database for restore sessions in the Data Protector GUI Restore context:

```
IPC Invalid Hostname or IP Address
```

The problem can appear in the following cases:

- When restoring database items to a different client.
- When importing Data Protector media containing backups of Oracle database from another Data Protector cell.
- When restoring 64-bit Oracle version 10.2.0.4 in the RAC environment, on HP-UX 11.23 PA-RISC systems. If `util_orarest` is present on the system, this error can mean that the `util_orarest` agent ends abnormally while trying to load the 32-bit OCI library from the `ORACLE_HOME/lib32` directory.

Actions

- To successfully restore database items to a different client, make sure that the system on which the Data Protector Oracle integration agent will be started is configured as Data Protector Oracle database instance (`ORACLE_SID`).

To verify, check if it is listed in the **Client** drop-down list in the **Options** page.

Select the system and proceed with [In the Options page, from the Client drop-down list, select the system on which the Data Protector Oracle integration agent will be started. To restore the database objects to a different database than it is selected, click Settings and specify the login information for the target database. on page 71 of the Restoring Oracle database objects on page 69 procedure.](#)

- When restoring 64-bit Oracle database in RAC environment, on HP-UX 11.23, resolve the problem as follows:

In the directory `ORACLE_HOME/lib` remove the soft link `libclntsh.sl`, which points to the 64-bit OCI library `ORACLE_HOME/lib/liblntsh.sl.10.1`.

Problem

When editing the RMAN script section of a backup specification using the Data Protector GUI, an RMAN backup script error is displayed

In the Data Protector GUI, when you edit the RMAN script section of a Data Protector backup specification, the following error message may display:

Cannot proceed, invalid RMAN backup script.

The error is displayed if you have specified an Oracle RMAN parameter, which has not been recognized by the Data Protector parser or a parsing error has occurred.

Action

Disable Oracle RMAN script parsing in the Data Protector GUI by setting the Data Protector NoGUIRMANScriptParsing global option to 1.

For details of how to set the option, see the *HP Data Protector Help* index: "global options".

Chapter 2: Data Protector SAP R/3 integration

Introduction

This chapter explains how to configure and use the Data Protector SAP R/3 integration (**SAP R/3 integration**). It describes concepts and methods you need to understand to back up and restore the following files of the SAP R/3 database environment (**SAP R/3 objects**):

- Data files
- Control files
- Online redo logs
- Offline (archived) redo logs
- SAP R/3 logs and parameter files

Data Protector supports offline and online backups. During an online backup, the SAP R/3 application is actively used.

Data Protector offers interactive and scheduled backups of the following types:

Table 7: Backup types

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

You can start backups using:

- The Data Protector user interface
- The SAP BRTOOLS interface

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

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

You can restore Data Protector backups using:

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

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

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

Integration concepts

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

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

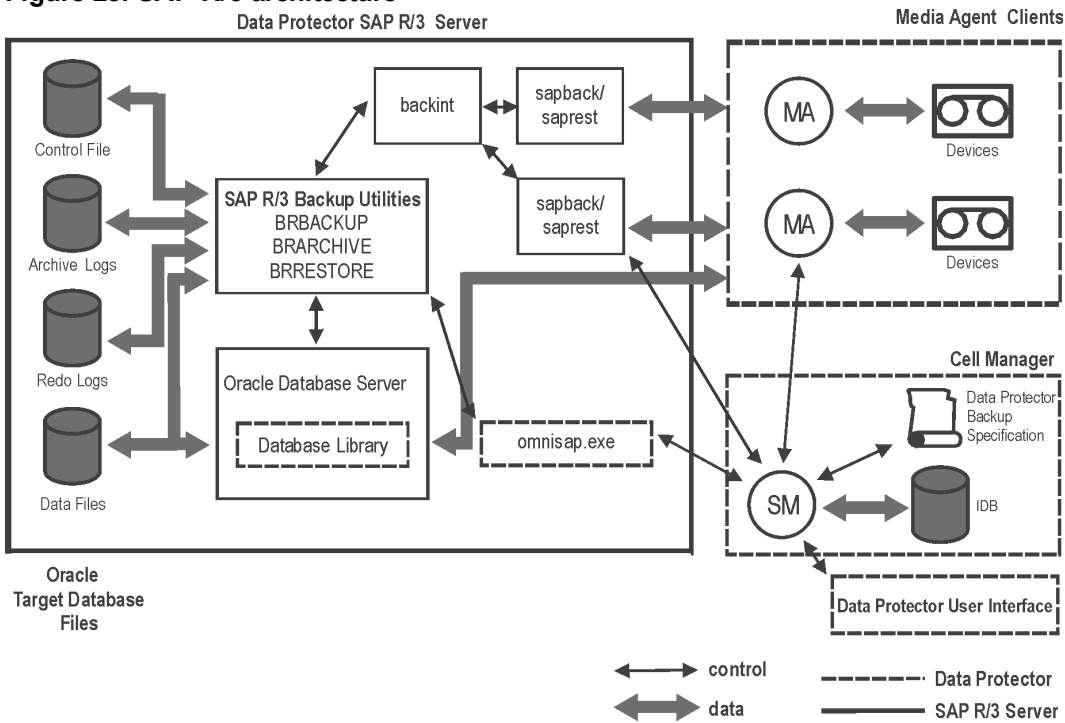
Table 8: SAP backup and restore utilities

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

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

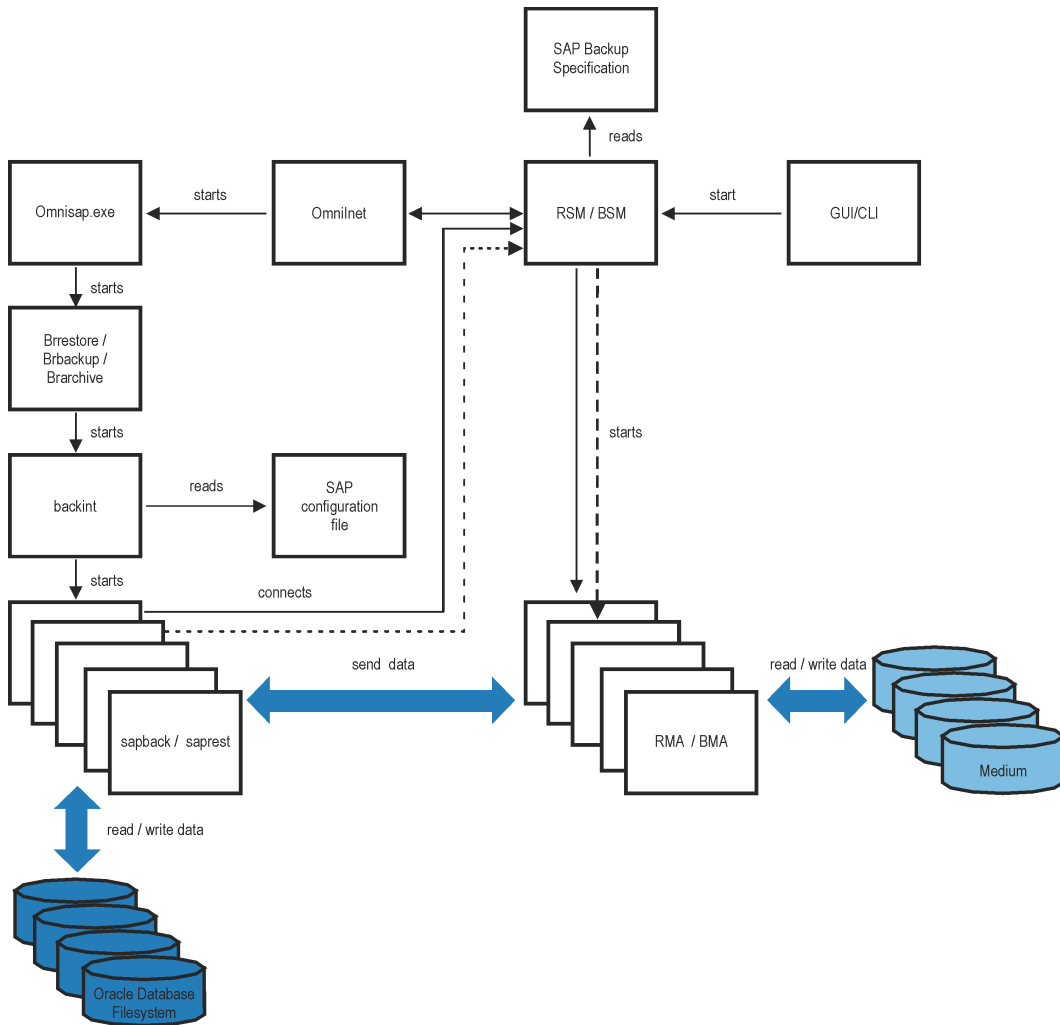
backint	Data is backed up using the Data Protector SAP R/3 integration.
RMAN	Data is backed up using the Oracle Recovery Manager (RMAN). The main benefit of the RMAN mode is that you can back up Oracle database incrementally.

Figure 28: SAP R/3 architecture



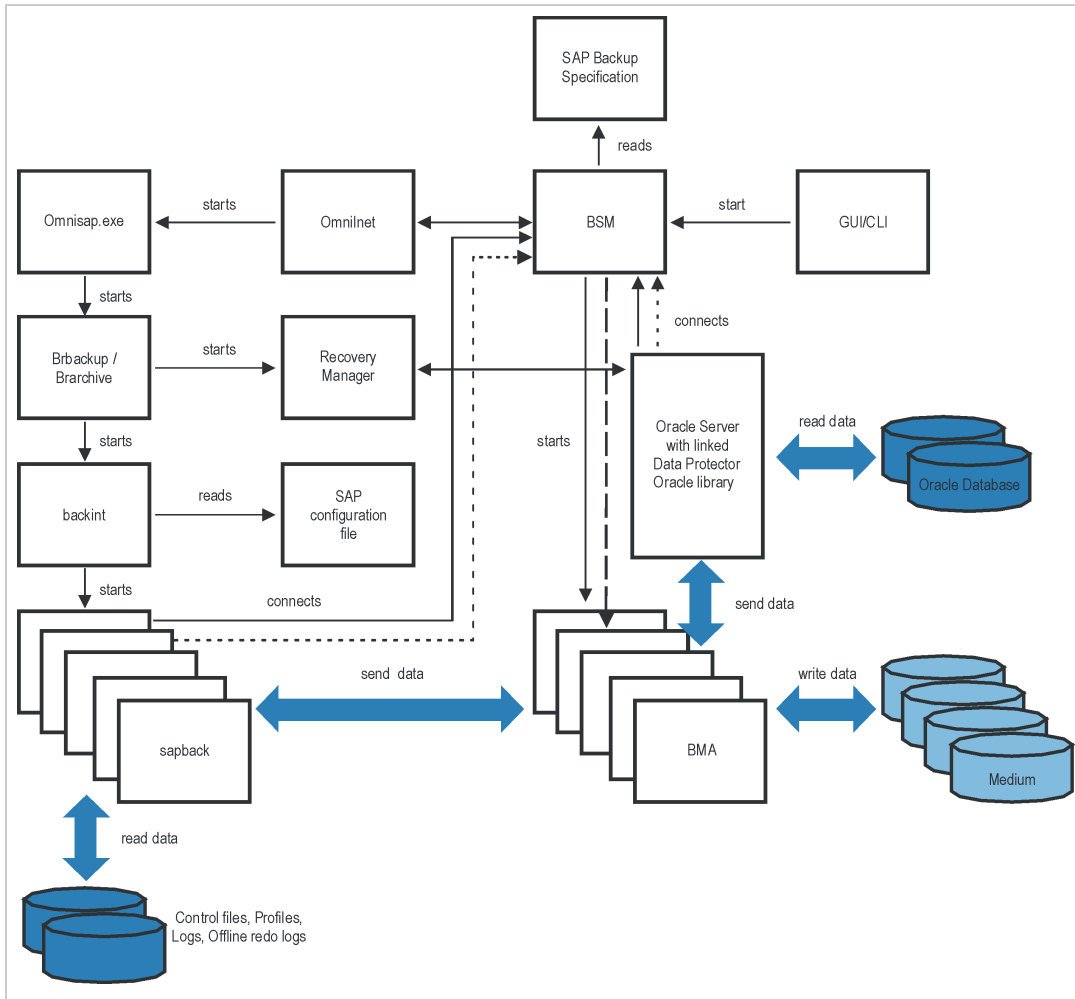
Legend	
SM	Data Protector Session Manager: Backup Session Manager during backup and Restore Session Manager during restore.
Database Library	A set of Data Protector executables that enable data transfer between Oracle Server and Data Protector. Required only if Oracle data files are backed up in the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

Figure 29: SAP R/3 architecture: backint mode



Legend	
BSM/RSM	Data Protector Backup Session Manager/Restore Session Manager
BMA/RMA	Data Protector backup/restore Media Agent
GUI/CLI	Data Protector graphical user interface/command-line interface

Figure 30: SAP R/3 architecture: RMAN mode



Backup flow

1. If the backup session is started:
 - **Using the Data Protector interface (or the scheduler):** BSM is started, which reads the appropriate Data Protector backup specification, checks if the devices are available, and starts `omnisap.exe` on the SAP R/3 client. The `omnisap.exe` agent exports the appropriate environment variables and starts BRBACKUP or BRARCHIVE.
 - **Using SAP BRTOOLS interface:** BRBACKUP or BRARCHIVE are started directly.
2. BRBACKUP does the following:
 - Changes the state of the Oracle Target Database (opened or closed), according to the backup type (online or offline).
 - Switches the Oracle Target Database to the ARCHIVELOG mode.

The archived redo log files are written to the archiving directory by Oracle and are backed up later using BRARCHIVE.

- Creates the BRBACKUP log during the backup session, which contains information about backed up files and the backup ID. This information is needed to determine the location of the database files and archived redo log files during restore.
 - Sets the tablespace mode (BEGIN / END BACKUP) in the case of online backup using backint. In this way, the SAP R/3 application puts a tablespace in the backup mode just before it is backed up and returns it to the normal mode immediately after the backup completes.
3. ■ If BRBACKUP is started:
- i. BRBACKUP starts a backint command (backint mode) or RMAN (RMAN mode), which backs up Oracle data files and control files.
 - ii. BRBACKUP starts a backint command (in the backint and RMAN mode), which backs up the SAP parameter file and the SAP R/3 history files that have been created during the backup of Oracle data files and control files.
- If BRARCHIVE is started (in the backint or RMAN mode), BRARCHIVE starts a backint command, which backs up archived redo log files. In addition, a copy of control files is created, which is also backed up.

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

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

Restore flow

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

1. When you select objects to be restored and start a restore using SAP BRTOOLS, the following happens (depending on which mode you use):
 - **Backint mode:** BRRESTORE checks if enough free disk space is available and starts a backint command to restore Oracle data files.

If the backups of files to be restored reside on different media, backint starts a separate `saprest` process for each medium, so that the files are restored in parallel (provided that the specified concurrency is large enough). The first `saprest` process starts a RSM, while the subsequent

saprest processes connect to the same RSM. RSM checks if the restore devices are available and starts the data flow.

- **RMAN mode:** BRRESTORE starts RMAN, which connects to Data Protector via Data Protector Database Library and Oracle Server processes and enables data transfer of Oracle data files.
2. When all the General Media Agents finish with data transfer, the RSM waits for a timeout (`smWaitForNewRestoreClient` global option) and completes the restore session if no `backint` command is started within this time frame.

Data Protector SAP R/3 configuration file

Data Protector stores the integration parameters for every configured SAP R/3 database in the following file on the Cell Manager:

Windows systems: `Data_Protector_program_data\Config\Server\Integ\Config\Sap\ClientName%ORACLE_SID`

UNIX systems: `/etc/opt/omni/server/integ/config/SAP/ClientName%ORACLE_SID`

The parameters stored are:

- Oracle home directory
- encoded connection string to the target database
- BRTOOLS home directory
- the variables which need to be exported prior to starting a backup
- SAPDATA home directory
- user name and user group
- temporary directory used for the copy of the control file or redo logs
- list of control files and redo logs that will be copied to a safe location
- character set (`ORA_NLS_CHARACTERSET`)
- concurrency number and balancing (for each backup specification), and number of channels for RMAN backup
- speed parameters (time needed for a specific file to back up - in seconds)
- manual balancing parameters

The configuration parameters are written to the Data Protector SAP R/3 configuration file:

- during configuration of the integration
- during creation of a backup specification
- when the configuration parameters are changed

Important: To avoid problems with your backups, take extra care to ensure the syntax and punctuation of your configuration file match the examples.

Note: You can set up the parameters in the Environment section (sublist) of the file by referring to other environment variables in the following way:

```
SAPDATA_HOME=${ORACLE_HOME}/data
```

Syntax

The syntax of the Data Protector SAP R/3 configuration file is as follows:

```
ORACLE_HOME='ORACLE_HOME';
ConnStr='ENCODED_CONNECTION_STRING_TO_THE_TARGET_DATABASE';
BR_directory='BRTTOOLS_HOME';
SAPDATA_HOME='SAPDATA_HOME';
ORA_NLS_CHARACTERSET='CHARACTER_SET';
OSUSER='USER_NAME';
OSGROUP='USER_GROUP';
Environment={
  [ENV_var1='value1'];
  [ENV_var2='value2';
  ...]
}
SAP_Parameters={backup_spec_name=(' -concurrency #_of_concurrency
' | '-time_balance' | '-load_balance' | '-manual_balance' | '-channels
#_of_RMAN_channels');
}
speed={
  AVERAGE=1;
  'filename'=#_of_seconds_needed_to_back_up_this_file;
}
compression={'filename'=size_of_the_file_in_bytes_after_the
_compression;
}
manual_balance={backup_specification_name={
  'filename'=device_number;
  }
}
```

The ORA_NLS_CHARACTERSET parameter is set automatically by Data Protector during SAP R/3 database configuration. For details of how to configure SAP R/3 database for use with Data Protector, see [Configuring SAP R/3 databases on page 128](#).

Example

This is an example of the file:

```
ORACLE_HOME='/app/oracle805/product';
ConnStr='EIBBKIBBEIBBFIBBGHBB0HBB
QDBBOFBBCFBBPFBBFCBBIFBBGFBBDBBBFBBCFBBDFFBCCFBB';
BR_directory='/usr/sap/ABA/SYS/exe/run';
SAPDATA_HOME='/sap';
ORA_NLS_CHARACTERSET='USASCII7';
OSUSER='orasid';
OSGROUP='dba';

Environment={
  SAP_Parameters={
    sap_weekly_offline=('-concurrency 1','-no_balance');
    sap_daily_online=('-concurrency 3','-load_balance');
    sap_daily_manual=('-concurrency 3','-manual_balance');
  }
  speed={
    AVERAGE=203971;
    '/file1'=138186;
    '/file2'=269756;
  }
  compression={
    '/file1'=1234;
    '/file2'=5678;
  }
  manual_balance={
    sap_daily_manual={
    '/file1'=1; /* file 1 is backed up by the first sapback */
    '/file2'=2; /* file 2 is backed up by the second sapback */
    '/file3'=1; /* file 3 is backed up by the first sapback */
    '/file4'=1;
    }
  }
}
```

Setting, retrieving, listing, and deleting Data Protector SAP R/3 configuration file parameters using the CLI

The Data Protector SAP R/3 configuration file parameters are normally written to the Data Protector SAP R/3 configuration file after:

- the Data Protector configuration of the Oracle instance that is run by SAP R/3 is completed.
- a new backup specification is created.
- a backup that uses balancing by time algorithm is completed.

The *util_cmd* command

You can set, retrieve, list, or delete the Data Protector SAP R/3 configuration file parameters using the *util_cmd -putopt* (setting a parameter), *util_cmd -getopt* (retrieving a parameter), or *util_cmd -getconf* (listing all parameters) command on the Data Protector SAP R/3 client.

Cluster-aware clients

In a cluster environment, the environment variable *OB2BARHOSTNAME* must be defined as the virtual hostname before executing the *util_cmd* command from the command line (on the client). The *OB2BARHOSTNAME* variable is set as follows:

Windows systems: `set OB2BARHOSTNAME=virtual_hostname`

UNIX systems: `export OB2BARHOSTNAME=virtual_hostname`

The *util_cmd* synopsis

The syntax of the *util_cmd* command is as follows:

```
util_cmd -getconf[ig] SAP oracle_instance [-local filename]
```

```
util_cmd -getopt[ion] [SAP oracle_instance] option_name [-sub[list] sublist_name] [-local filename]
```

```
util_cmd -putopt[ion] [SAP oracle_instance] option_name [option_value] [-sub[list] sublist_name] [-local filename]
```

where:

option_name is the name of the parameter

option_value is the value for the parameter

[-sub[list] sublist_name] specifies the sublist in the configuration file to which a parameter is written to or taken from.

[-local filename] specifies one of the following:

- When it is used with the *-getconf*[ig] option, it specifies the filename for the output of the command to be written to. If the *-local* option is not specified, the output is written to the standard output.
- When it is used with the *-getopt*[ion], it specifies the filename of the file from which the parameter and its value are to be taken and then written to the standard output. If the *-local* option is not specified, the parameter and its value are taken from the Data Protector SAP R/3 configuration file and then written to the standard output.
- When it is used with the *-putopt*[ion] option, it specifies the filename for the output of the command to be written to. If the *-local* option is not specified, the output is written to the Data Protector SAP R/3 configuration file.

Note: If you are setting the *option_value* parameter as a number, the number must be put in single quotes, surrounded by double quotes.

Return values

The `util_cmd` command displays a short status message after each operation (writes it to the standard error):

- Configuration read/write operation successful.

This message is displayed when all the requested operations have been completed successfully.

- Configuration option/file not found.

This message is displayed when either an option with the specified name does not exist in the configuration, or the file specified as the `-local` parameter does not exist.

- Configuration read/write operation failed.

This message is displayed if any fatal errors occurred, for example: the Cell Manager is unavailable, the Data Protector SAP R/3 configuration file is missing on the Cell Manager, and so on.

Setting parameters

To set the Data Protector `OB2OPTS` and the Oracle `BR_TRACE` parameters for the Oracle instance `ICE` that is run by SAP R/3, use the following commands on the Data Protector SAP R/3 client:

Windows, HP-UX, Solaris, and Linux systems

```
util_cmd -putopt SAP ICE OB2OPTS '-debug 1-200 debug.txt' -sublist Environment  
util_cmd -putopt SAP ICE BR_TRACE "'10'" -sublist Environment
```

Other UNIX systems

```
util_cmd -putopt SAP ICE NLS_LANG 'US7ASCII' -sublist Environment  
util_cmd -putopt SAP TOR BR_TRACE "'10'" -sublist Environment
```

Retrieving parameters

To retrieve the value of the `OB2OPTS` parameter for the Oracle instance `ICE`, use the following command on the Data Protector SAP R/3 client:

```
util_cmd -getopt SAP ICE OB2OPTS -sublist Environment
```

Listing parameters

To list all the Data Protector SAP R/3 configuration file parameters for the Oracle instance `ICE`, use the following command on the Data Protector SAP R/3 client:

```
util_cmd -getconf SAP ICE
```

Deleting parameters

To remove the value of the `OB2OPTS` parameter for the Oracle instance `ICE`, use the following command on the Data Protector SAP R/3 client:

```
util_cmd -putopt SAP ICE OB2OPTS "" -sublist Environment
```

Configuring the integration

To configure the integration:

1. Configure the required user accounts. See [Configuring user accounts on page 125](#).
2. Check the connection to the Oracle database. See [Checking the connection on page 125](#).
3. Enable the use of the authentication password file. See [Authentication password file on page 126](#).
4. Optionally, set the archived logging mode to enable online backups. See [Enabling archived logging on page 126](#).
5. Configure every SAP R/3 database you intend to back up from or restore to. See [Configuring SAP R/3 databases on page 128](#).

Prerequisites

- Ensure that you have correctly installed and configured the SAP R/3 application. The database used by the SAP R/3 application must be an Oracle database. If any other database is used, you can back it up using the corresponding Data Protector integration (for example, Informix). It is assumed that you are familiar with the SAP R/3 application and Oracle database administration.
 - For supported versions, platforms, devices, and other information, see the latest support matrices at <http://support.openview.hp.com/selfsolve/manuals>.
 - For information on installing, configuring, and using the SAP R/3 application and the SAP backup and restore tools (BRBACKUP, BRRESTORE, and BRARCHIVE), see the SAP R/3 application documentation.
- Ensure that you have a license to use the Data Protector SAP R/3 integration. For information, see the *HP Data Protector Installation and Licensing Guide*.
- Ensure that you have correctly installed Data Protector.
 - For information on how to install the Data Protector SAP R/3 integration in various architectures, see the *HP Data Protector Installation and Licensing Guide*.
 - For information on the Data Protector Cell Manager package configuration in the HP SG cluster, see the *HP Data Protector Help* index: "HP Serviceguard integration".

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

Before you begin

- Configure devices and media for use with Data Protector.
- To test whether the SAP R/3 system and the Cell Manager communicate properly, configure and run a Data Protector filesystem backup and restore.
- **Windows systems:**
 - On Windows Server 2003 system, you need to restart the Data Protector Inet service under the Oracle operating system user account described in [Configuring user accounts on the next page](#).

For information on changing the user account under which the Data Protector Inet service is running, see the *HP Data Protector Help* index: “Inet, changing account”.

- On other Windows operating systems, configure the Data Protector Inet service user impersonation for the user that has the appropriate SAP R/3 permissions for running backups and restores.

For details, see the *HP Data Protector Help* index: “Inet user impersonation”.

If there are several SAP R/3 instances running on the same system with different SAP administrator accounts configured for each instance, create an additional, common SAP administrator account. Configure the Data Protector Inet service to use this account as the service startup account.

Cluster-aware clients

- Configure SAP R/3 databases only on one cluster node, since the configuration files reside on the Cell Manager.

Windows systems: During the configuration, Data Protector copies the Data Protector `backint` program from `Data_Protector_home\bin` to the directory that stores the SAP backup tools. This is done only on the currently active node. On the other node, do it manually.

UNIX systems: During the configuration, Data Protector creates a link to the Data Protector `backint` program on the currently active node. On all the other nodes, do it manually. Execute:

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

- If you intend to use the Data Protector CLI, set the Data Protector environment variable `OB2BARHOSTNAME` to the virtual server name as follows:

Windows systems: `set OB2BARHOSTNAME=virtual_server_name`

UNIX systems: `export OB2BARHOSTNAME=virtual_server_name`

Note: SAP recommends installing SAP backup utilities on all cluster nodes.

Configuring user accounts

To enable backup and restore of SAP R/3 database files, you need to configure or create several user accounts.

Oracle operating system user account	Operating system user account that is added to the following user groups: Windows systems: ORA_DBA and ORA_SID_DBA local groups UNIX systems: dba and sapsys For example, user oraSID. UNIX systems: Ensure that this user is the owner of the filesystem or of the raw logical volume on which the database is mounted. The minimum permissions should be 740.
User account root (UNIX systems only)	Default operating system administrator's user account added to the dba user group.
Oracle database user account	Database user account granted at least the following Oracle roles: <ul style="list-style-type: none">• sysdba• sysoper For example, user system. Do not configure the Oracle SYS user for backing up SAP R/3 objects. When backing up using the SYS user account, the SAP backup fails with the error ORA-28009: connection as SYS should be as SYSDBA or SYSOPER.

Add the following user accounts to the Data Protector admin or operator user group:

- Oracle operating system user account
(if you are using backup set method, add this user on the application as well as on backup system)
- **UNIX systems:** User account root

In cluster environments, add these user accounts to the Data Protector admin or operator user group for the following clients:

- virtual server
- every node in the cluster

For information on adding Data Protector users, see the *HP Data Protector Help* index: "adding users".

Checking the connection

To check the connection to the Oracle instance:

1. Log in to the SAP R/3 client as the Oracle OS user.
2. Export/set the `ORACLE_HOME` and `ORACLE_SID` variables.
3. Start `sqlplus`.
4. Connect to the Oracle target database as the Oracle database user, first with the `sysdba` role and then with the `sysoper` role.

Example

For the following configuration:

Oracle instance: `PROORACLE_HOME: /app/oracle816/product`

execute:

```
id
uid=102(oraprod) gid=101(dba)
export ORACLE_SID=PRO
export ORACLE_HOME=/app/oracle816/product
export SHLIB_PATH=/app/oracle816/product/lib:/opt/omni/lbin
sqlplus /nolog
SQLPLUS> connect system/manager@PRO as sysdba;
Connected.
SQLPLUS> connect system/manager@PRO as sysoper;
Connected.
```

Authentication password file

Enable the use of the authentication password file for the database administrator:

1. Shut down the Oracle target database.
2. In the `initORACLE_SID.ora` file, specify:

```
remote_login_passwordfile = exclusive
```

For instructions on how to set up the password file, see the Oracle documentation.

Enabling archived logging

When you set the database to the archived logging mode, you protect the unsaved online redo logs from being overwritten. Online backup of data files is useless without the related redo logs because you cannot recover the database to a consistent state.

Tip: Archive the redo log files generated during the online backup immediately after BRBACKUP completes.

To protect the archive directory from overflowing, clear the directory regularly.

To enable archived logging:

1. In the `initORACLE_SID.ora` file, set

```
log_archive_start = true
```

and specify the `log_archive_dest` option.

Example

This is an example of the `initORACLE_SID.ora` file for the Oracle instance PRO:

```
# @(#)initSID.ora 20.4.6.1 SAP 13/03/30
#####
# (c)Copyright SAP AG, Walldorf
#####
. . . .
. . . . .
. . . . .
. . . . .
### ORACLE Authentication Password File
remote_login_passwordfile = exclusive
### ORACLE archiving
log_archive_dest = /oracle/PRO/saparch/PROarch
log_archive_start = true
. . . .
```

2. Mount the Oracle database and start the archived logging mode using the Oracle Server Manager.
Execute:

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

Example

For the Oracle instance PRO, execute:

Windows systems: set ORACLE_SID=PRO

UNIX systems: export ORACLE_SID=PRO

Any operating system:

```
sqlplus /nolog
SQLPLUS> connect user/passwd@PRO;
Connected.
SQLPLUS> startup mount
ORACLE instance started.
Total System Global Area          6060224 bytes
Fixed Size                          47296 bytes
Variable Size                      4292608 bytes
```

```
Database Buffers          1638400 bytes
Redo Buffers              81920 bytes
Database mounted.
SQLPLUS> alter database archive log;
Statement processed.
SQLPLUS> archive log start;
Statement processed.
SQLPLUS> alter database open;
```

Linking Oracle Server with the Data Protector MML

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

- When you start backups or restores using the Data Protector GUI or CLI, Data Protector automatically links Oracle Server with the correct platform–specific Data Protector MML.

Note: For testing purposes, you can override this automatic selection. You can manually specify which Data Protector MML should be used by setting the Data Protector SBT_LIBRARY parameter. The parameter is saved in the Data Protector SAP R/3 instance configuration file. On how to set the parameter, see the `util_cmd` man page.

- To start backups using the Oracle Recovery Manager or BRBACKUP utility directly, you need to manually link Oracle Server software with the correct platform–specific Data Protector MML as described in [Backing up using Oracle Recovery Manager on page 146](#).

Choosing authentication mode

Data Protector SAP R/3 integration supports two authentication modes for accessing Oracle databases that are used by SAP R/3:

- database authentication mode
- operating system authentication mode

With database authentication mode, you need to re-configure the SAP R/3 integration for an SAP R/3 database with the new Oracle login information each time the corresponding Oracle database user account changes. Such a reconfiguration is not needed if operating system authentication mode is used.

You select the preferred authentication mode when you configure a particular SAP R/3 database.

Configuring SAP R/3 databases

You need to provide Data Protector with the following configuration parameters:

- Oracle Server home directory
- SAP R/3 data home directory
- Optionally, if you choose database authentication mode, Oracle database user account. The user account is used by BRBACKUP and BRARCHIVE during backup.
- Directory in which the SAP backup utilities are stored

Data Protector then creates the configuration file for the SAP R/3 database on the Cell Manager and verifies the connection to the database. On UNIX systems, Data Protector also creates a soft link for the `backint` program from the directory that stores the SAP backup utilities to.:

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

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

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

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

To configure an SAP R/3 database, use the Data Protector GUI or CLI.

Before you begin

- Ensure that the SAP R/3 database is open.

Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **SAP R/3**, and click **Add Backup**.
3. In the **Create New Backup** dialog box, select the template.

Click **OK**.

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

Specify other disk array specific backup options (see). For details on the options, press **F1**.

Note: Select **Track the replica for instant recovery**.

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

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

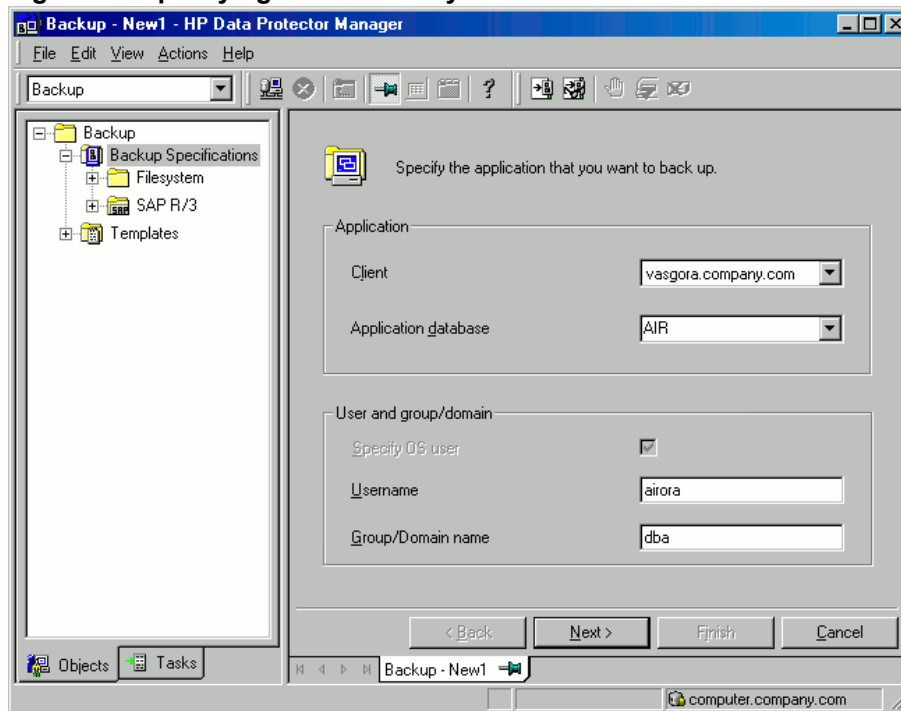
Windows Server 2008: In **Username** and **Group/Domain name**, specify the operating system user account under which you want the backup session to run (for example, the user name Administrator, domain DP).

UNIX systems: In **Username**, type the Oracle OS user described in [Configuring user accounts on page 125](#). In **Group/Domain name**, type dba.

Ensure that this user has been added to the Data Protector admin or operator user group, has the SAP R/3 backup rights, and has been set up for the Data Protector Inet service user impersonation. This user becomes the backup owner.

For details on setting accounts for the Inet service user impersonation, see the *HP Data Protector Help* index: "Inet user impersonation".

Figure 31: Specifying an SAP R/3 system and Oracle instance



Click **Next**.

6. In the **Configure SAP** dialog box, specify the pathname of the Oracle Server home directory and SAP R/3 data home directory. If you leave the fields empty, the default *ORACLE_HOME* directory is used.

Under **Oracle login information to target database**, specify the following:

- For the database authentication mode, specify **Username**, **Password**, and **Service**.
- For the local operating system authentication mode, leave **Username**, **Password**, and **Service** empty.
- For the remote operating system authentication mode, specify only **Service** (leave **Username** and **Password** empty).

The following are the option descriptions:

- **Username** and **Password**: Specify the user name and password of the Oracle database user account described in [Configuring user accounts on page 125](#).
- **Service** : Specify the Oracle service name.

In **Backup and restore executables directory**, specify the pathname of the directory in which the SAP backup utilities reside. By default, the utilities reside in:

Windows systems: \\SAP_system\sapmnt\ORACLE_SID\sys\exe\run

UNIX systems: /usr/sap/ORACLE_SID/SYS/exe/run

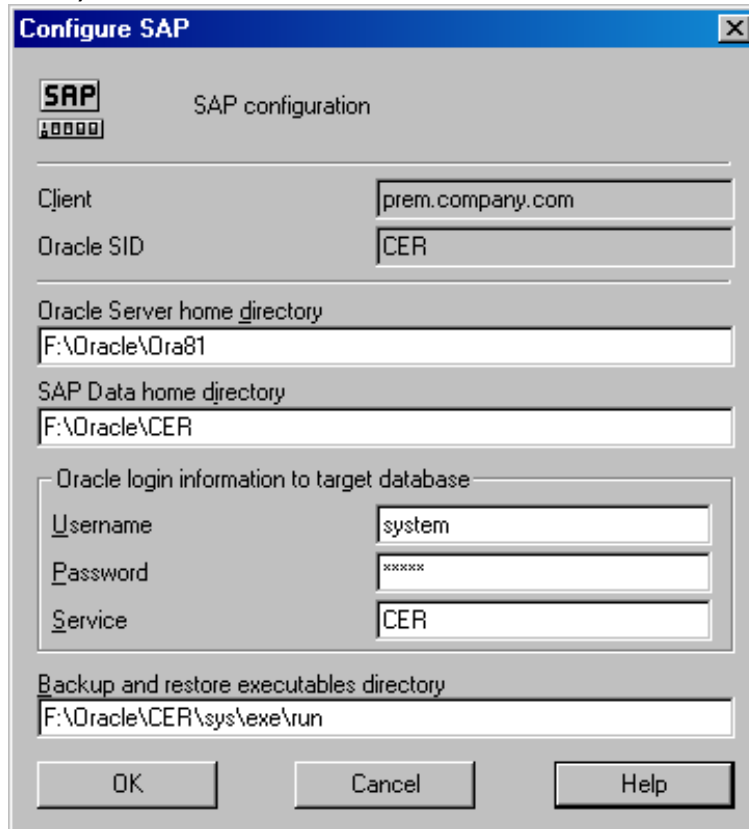
Figure 32: Configuring an SAP R/3 database on a UNIX system (operating system authentication mode)

The screenshot shows a 'Configure SAP' dialog box with the following fields and values:

- Client:** vasgora.company.com
- Oracle SID:** AIR
- Oracle Server home directory:** /app/oracle/product/81
- SAP Data home directory:** /app/oracle/product/AIR
- Oracle login information to target database:**
 - Username:** (empty)
 - Password:** (empty)
 - Service:** AIR
- Backup and restore executables directory:** /usr/sap/CER/SYS/exe/run

Buttons at the bottom: OK, Cancel, Help.

Figure 33: Configuring an SAP R/3 database on a Windows system (database authentication mode)



Click **OK**.

7. The SAP R/3 database is configured. Exit the GUI or proceed with creating the backup specification at [Select SAP R/3 objects to be backed up](#). You can select individual tablespaces, data files, or archived logs. on page 139.

Using the Data Protector CLI

1. Log in to the SAP R/3 system using the Oracle operating system user account.
2. At the command prompt, change current directory to the following directory:

Windows systems: `Data_Protector_home\bin`

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

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

3. Execute:

```
util_sap.exe -CONFIG ORACLE_SIDORACLE_HOMEtargetdb_connection_stringSAPTOOLS_DIR  
[SAPDATA_HOME][SQL_PATH]
```

Parameter description

<i>ORACLE_SID</i>	Oracle instance name.
<i>ORACLE_HOME</i>	Pathname of the Oracle Server home directory.
<i>targetdb_connection_string</i>	This argument value determines the authentication mode used for accessing the Oracle database: <ul style="list-style-type: none">■ To select the database authentication mode, specify the login information to the target database in the format <i>user_name/password@Oracle_service</i>.■ To select the local operating system authentication mode, specify only the character <i>/</i>.■ To select the remote operating system authentication mode, specify the login information to the target database in the format <i>/@Oracle_service</i>.
<i>SAPTOOLS_DIR</i>	Pathname of the directory that stores the SAP backup utilities.
<i>SAPDATA_HOME</i>	Pathname of the directory where the SAP R/3 data files are installed. By default, this parameter is set to <i>ORACLE_HOME</i> .

The message *RETVAL*0 indicates successful configuration.

Handling errors

If you receive the message *RETVAL**error_number* where *error_number* is different than zero, an error occurred.

To get the error description, execute:

Windows systems:

```
Data_Protector_home\bin\omnigetmsg 12 error_number
```

HP-UX and Linux systems:

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

Other UNIX systems:

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

Tip: To get a list of Oracle instances that are used by the SAP R/3 application, execute:

```
util_sap.exe -APP
```

To get a list of tablespaces of an Oracle instance, execute:

```
util_sap.exe -OBJ$ORACLE_SID
```

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

```
util_sap.exe -OBJ$ORACLE_SID TABLESPACE
```

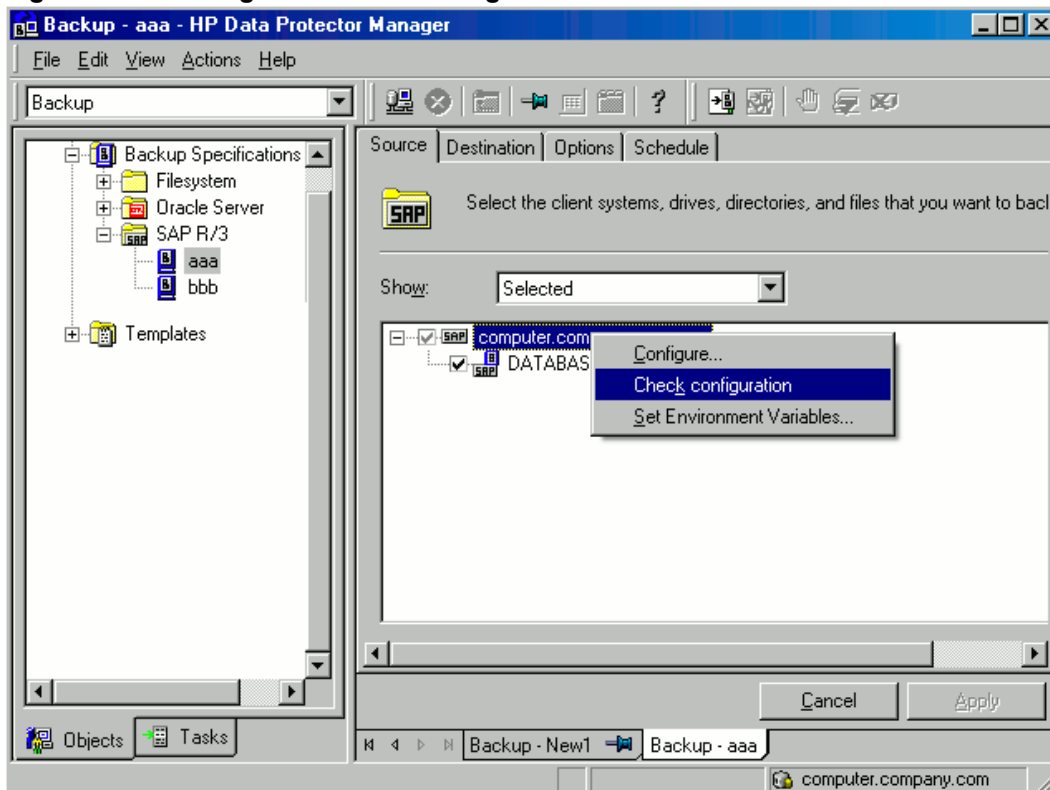
Checking the configuration

You can check the configuration of an SAP R/3 database after you have created at least one backup specification for this database. Use the Data Protector GUI or CLI.

Using the Data Protector GUI

1. In the Context List, select **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **SAP R/3**. Click the backup specification to display the Oracle instance to be checked.
3. Right-click the Oracle instance and click **Check configuration**.

Figure 34: Checking the SAP R/3 configuration



Using the Data Protector CLI

Log in to the SAP R/3 system as the Oracle OS user and execute:

util_sap.exe -CHKCONF ORACLE_SID

where ORACLE_SID is the name of the Oracle instance.

A successful configuration check displays the message *RETVAl*0.

If you receive the message *RETVAl*error_number where error_number is different than zero, an error occurred. On how to get the error description, see [Handling errors on page 133](#).

Parameter	Value/Description
split_cmd	<p>On the application system:</p> <p>UNIX systems: "/opt/omni/lbin/ob2smbsplit \$"</p> <p>Windows systems:"Data_Protector_home\bin\ob2smbsplit \$"</p> <p>On the backup system, you do not need to set the parameter.</p> <p>BRBACKUP uses this parameter to trigger the replica creation. At run time, the optional sign "\$" is replaced with the name of the text file containing the names of files to be backed up.</p> <p>Windows systems: If the pathname contains spaces, use Windows short names instead.</p>
primary_db	<p>On the application system: LOCAL</p> <p>On the backup system: name of the service used for connecting to the Oracle database.</p> <p>This parameter defines the service name of the Oracle database to link the backup system to the application system.</p>

Backup

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

Table 9: Backup types

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

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

To configure a backup, create a backup specification.

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

Table 10: What is backed up

Selected items	Backed up files
ARCHIVELOGS	<ul style="list-style-type: none"> • offline (archived) redo logs • control files
DATABASE or individual tablespaces	<ul style="list-style-type: none"> • data files • control files • SAP R/3 logs and parameter files • online redo logs (only during offline backups)

You can specify SAP R/3 backup options in two different ways:

- Using the BRBACKUP options.
- Using the SAP parameter file.

Note: The BRBACKUP options override the settings in the SAP parameter file.

You can specify BRBACKUP options when you create a backup specification. If no options are specified, the SAP R/3 application refers to the current settings in the SAP parameter file. In such a case, before running a backup, ensure that the SAP parameter file is correctly configured. See examples in [Two alternatives of specifying backup options below](#).

Table 11: Two alternatives of specifying backup options

Backup type	1. BRBACKUP options 2. SAP parameter file settings
offline backup using backint	1. <code>-t offline -d util_file</code> 2. <code>backup_type = offline</code> <code>backup_dev_type = util_file</code>
online backup using backint (tablespaces are in the backup mode during the whole backup session)	1. <code>-t online -d util_file</code> 2. <code>backup_dev_type = util_file</code> <code>backup_type = online</code>

Backup type	1. BRBACKUP options 2. SAP parameter file settings
online backup using backint (tablespaces are in the backup mode only while being backed up)	1. -t online -d util_file_online 2. backup_dev_type = util_file_online backup_type = online
full backup	1. -m full 2. backup_mode = full
backup using RMAN	1. -d rman_util 2. backup_dev_type = rman_util rman_channels = <i>number_of_channels</i> rman_parms = "ENV=(OB2BARTYPE=SAP,OB2APPNAME= <i>DB_Name</i> , OB2BARLIST= <i>Backup_Specification_Name</i>)" For more information, see Backing up using Oracle Recovery Manager on page 146 .

Tip: When you create a backup specification, select a backup template that already contains the desired BRBACKUP options.

Considerations

- Before you start a backup, ensure that the SAP R/3 database is in the open or shutdown mode.
- Backup sessions that back up the same Oracle instance cannot run simultaneously.
- Generally, restore takes longer than backup. The restore is significantly prolonged if files are backed up with many streams. Note that if you start a backup in the RMAN mode with the Oracle RMAN script option FILESPERSET set to 1, RMAN creates a separate backup stream (object) for each database file.

Creating backup specifications

Create a backup specification using the Data Protector Manager.

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **SAP R/3**, and click **Add Backup**.
3. In the **Create New Backup** dialog box, select a template and click **OK**.

Table 12: Backup templates available for standard backup

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

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

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

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

Windows Server 2008: In **Username** and **Group/Domain name**, specify the operating system user account under which you want the backup session to run (for example, the user name Administrator, domain DP).

UNIX systems: In **Username**, type the Oracle OS user described in [Configuring user accounts on page 125](#). In **Group/Domain name**, type dba.

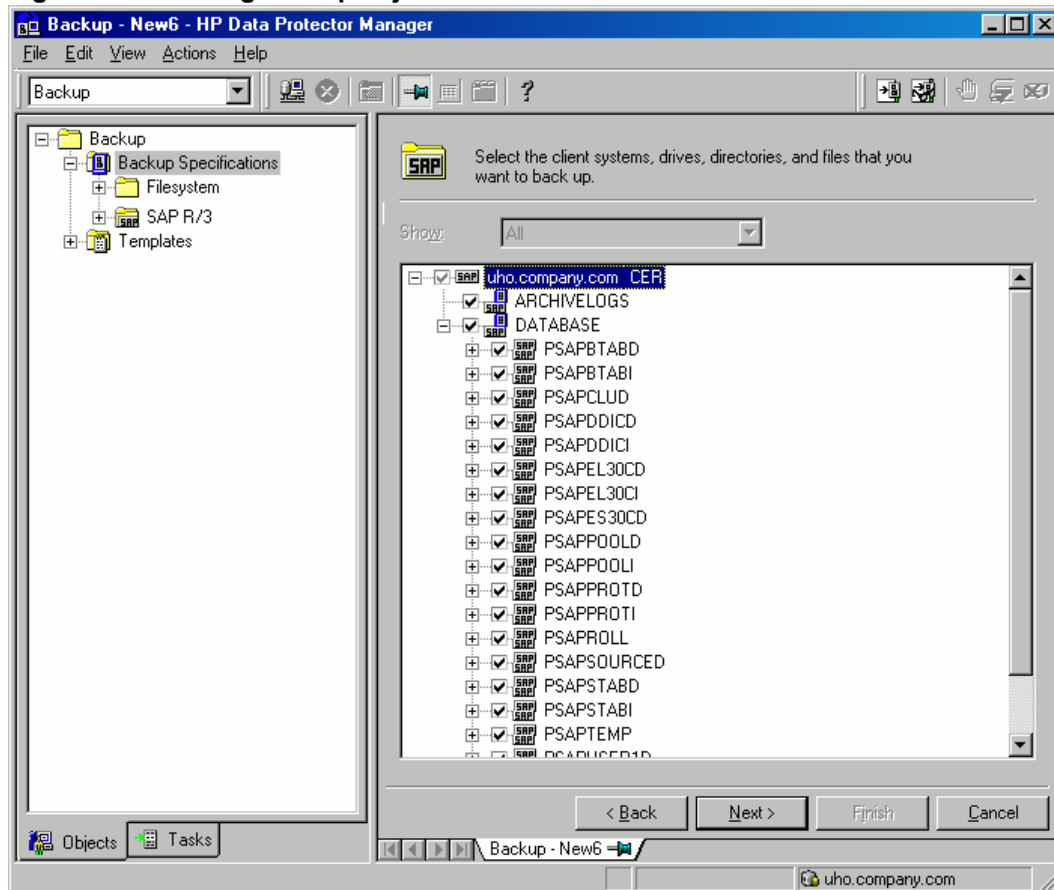
Ensure that this user has been added to the Data Protector admin or operator user group, has the SAP R/3 backup rights, and has been set up for the Data Protector Inet service user impersonation. This user becomes the backup owner.

For details on setting accounts for the Inet service user impersonation, see the *HP Data Protector Help* index: "Inet user impersonation".

Click **Next**.

5. If the SAP R/3 database is not configured yet for use with Data Protector, the **Configure SAP** dialog box is displayed. Configure it as described in [Configuring SAP R/3 databases on page 128](#).
6. Select SAP R/3 objects to be backed up. You can select individual tablespaces, data files, or archived logs.

Figure 35: Selecting backup objects



Click **Next**.

7. Select devices to use for the backup.

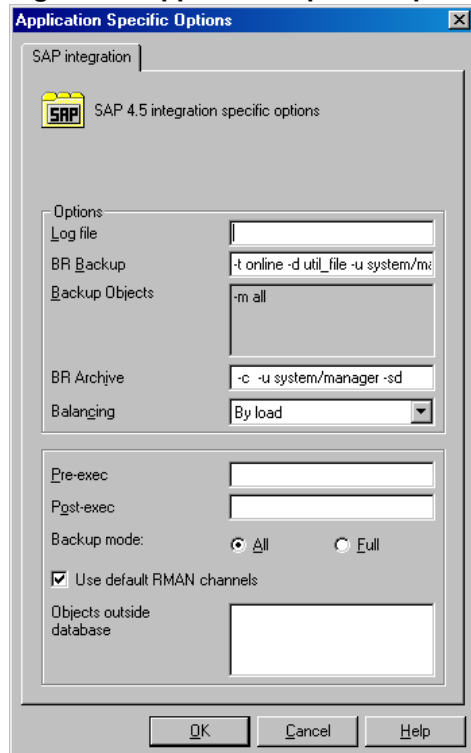
To specify device options, right-click the device and click **Properties**. Specify the number of parallel backup streams in the **Concurrency** tab and the media pool.

Note: Parallelism (the number of streams your SAP R/3 database is backed up with) is set automatically. If load balancing is used, the number of streams equals the sum of concurrencies of the selected devices.

Click **Next**.

8. Set backup options. For information on the application-specific options, see [SAP R/3 backup options on the next page](#).

Figure 36: Application-specific options



Click **Next**.

9. Optionally, schedule the backup. See [Scheduling backup sessions on page 142](#).

Click **Next**.

10. Save the backup specification, specifying a name and a backup specification group.

Tip: Preview your backup specification before using it for real. See [Previewing backup sessions on page 143](#).

Table 13: SAP R/3 backup options

Option	Description
Log file	If you want to create a backint log file during backup, specify a pathname for the file. By default, this file is not created because Data Protector stores all relevant information about backup sessions in the database.
BR Backup	Specifies BRBACKUP options. To run BRBACKUP under a different Oracle database user than the one specified during the configuration, type <code>-u user_name</code> .
Backup Objects	Lists BRBACKUP options passed by <code>omnisap.exe</code> . The list is displayed after you save the backup specification.
BR Archive	Specifies BRARCHIVE options.
Balancing: By Load	Groups files into subsets of approximately equal sizes. The subsets are then backed up concurrently by Data Protector <code>sapback</code> programs. If your backup devices use hardware compression, the sizes of the original and backed up files differ. To inform Data Protector of this, specify the original sizes of the backed up files in the <code>compression</code> section of the Data Protector SAP R/3 configuration file. See Data Protector SAP R/3 configuration file on page 118 .
Balancing: By Time	Groups files into subsets that are backed up in approximately equal periods of time. The duration depends on the file types, speed of the backup devices, and external influences (such as mount prompts). This option is best for environments with large libraries of the same quality. The subsets are backed up concurrently by Data Protector <code>sapback</code> programs. Data Protector automatically stores backup speed information in the <code>speed</code> section of the Data Protector SAP R/3 configuration file. It uses this information to optimize the backup time. This type of balancing may lead to non-optimal grouping of files in the case of an online backup or if the speed of backup devices varies significantly.
Balancing: Manual	Groups files into subsets as specified in the manual balancing section of the Data Protector SAP R/3 configuration file. For more information, see Manual balancing on page 147 . Not applicable for ZDB to disk.

Option	Description
Balancing: None	No balancing is used. The files are backed up in the same order as they are listed in the internal Oracle database structure. To check the order, use the Oracle Server Manager SQL command: <code>select * from dba_data_files</code>
Pre-exec , Post-exec	The command specified here is started by <code>omnisap.exe</code> on the SAP R/3 system before the backup (<code>pre-exec</code>) or after it (<code>post-exec</code>). Do not use double quotes. Provide only the name. The command must reside in the directory: Windows systems: <code>Data_Protector_home\bin</code> HP-UX, Solaris, and Linux systems: <code>/opt/omni/bin</code> Other UNIX systems: <code>/usr/omni/bin</code>
Backup mode	Specifies the RMAN backup type to be used. Available only if the whole database is selected for backup. If <code>All</code> is specified, RMAN backs up the whole database. If <code>Full</code> is specified, RMAN performs a Full backup (level 0), thus enabling RMAN incremental backups.
Use default RMAN channels	Specifies the concurrency value for your backup. Applicable only if RMAN is used for backup. This option overrides the settings in the SAP parameter file.
Objects outside database	Specifies non-database files of the Oracle SAP R/3 environment to be saved. Save these files in a separate backup session.

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

Modifying backup specifications

To modify your backup specification, click its name in the Scoping Pane of the Backup context, then click the appropriate tab, and apply the changes.

Scheduling backup sessions

You can run unattended backups at specific times or periodically. For details on scheduling, see the *HP Data Protector Help* index: "scheduled backups".

Scheduling example

To schedule **Full** backups at 8:00, 13:00, and 18:00 during week days:

1. In the **Schedule** property page, select the starting date in the calendar and click **Add** to open the **Schedule Backup** dialog box.
2. Under **Recurring**, select **Weekly**. Under **Time options**, select **8:00**. Under **Recurring Options**, select **Mon, Tue, Wed, Thu, and Fri**. See [Scheduling backup sessions below](#).

Click **OK**.

3. Repeat Step 1 and Step 2 to schedule backups at 13:00 and 18:00.
4. Click **Apply** to save the changes.

Figure 37: Scheduling backup sessions

Schedule Backup

Specify the desired backup time, frequency, duration, and type.

Recurring

None
 Daily
 Weekly
 Monthly

Time options

Time: 8:00
 Use starting
7/ 3/2007

Recurring options

Every 1 week(s) on

Sun Mon Tue Wed Thu Fri Sat

Session options

Backup type: Full
Network load: Full
Backup protection: Default

OK Cancel Help

Previewing backup sessions

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

Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **SAP R/3**. Right-click the backup specification you want to preview and click **Preview Backup**.
3. Specify **Backup type** and **Network load**. Click **OK**.

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

Using the Data Protector CLI

Execute:

```
omnib -sap_list backup_specification_name -test_bar
```

What happens during the preview?

The `omnisap.exe` command is started, which starts the Data Protector `testbar` command to test the following:

- Communication between the Oracle instance and Data Protector (only if RMAN is used)
- The syntax of the backup specification
- If devices are correctly specified
- If the necessary media are in the devices

Starting backup sessions

Interactive backups are run on demand. They are useful for urgent backups or restarting failed backups.

Backup methods

Start a backup of SAP R/3 objects in any of the following ways:

- Using the Data Protector GUI.
- Using the Data Protector CLI.
- Using the SAP BR*Tools.

Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, and then **SAP R/3**. Right-click the backup specification you want to use and click **Start Backup**.
3. Specify **Backup type** and **Network load**. Click **OK**.

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

Using the Data Protector CLI

Execute:

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

where `SAP_mode` is one of the following:

```
full|incr
```

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

Example

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

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

Using the SAP BRTOOLS

1. Log in to the SAP R/3 system as the Oracle OS user.
2. Export/set the following environment variables:

<code>ORACLE_SID=SAP_ instance_name</code>	
<code>ORACLE_ HOME=Oracle_ software_home_ directory</code>	
<code>[SAPBACKUP_ TYPE=OFFLINE]</code>	Default is ONLINE.

SAPDATA_ HOME= <i>database_ files_directory</i>	
SAPBACKUP= BRTOOLS_ <i>logs_ and_control_ file_copy_ directory</i>	
SAPREORG= BRSPACE_ <i>logs_ directory</i>	
OB2BARLIST= <i>backup_ specification_ name</i>	The backup specification is needed only to specify which Data Protector devices should be used for backup. Other information from the backup specification, like SAP R/3 objects to be backed up or the BRBACKUP options, is ignored and has to be specified manually at run time.
[OB2_3RD_PARTY_ BACKINT=1]	
[OB2BARHOSTNAME= <i>application_ system_name</i>]	Optional if you want to specify a virtual server name in cluster environments.

Alternatively, these variables can be specified in the backint parameter file. If this is required, the location of the file must be specified in the SAP configuration file using the *util_par_file* parameter:

`util_par_file = path\filename`

If you do not supply the path, the system searches for the parameter file in the directory:

Windows systems: `SAPDATA_HOME\database`

UNIX systems: `ORACLE_HOME/dbs`

- If you plan to do backups in the RMAN mode, ensure that the SBT_LIBRARY parameter in the `initSAP_instance.sap` file points to the correct platform-specific Data Protector MML. For details on the Data Protector MML location, see [Specify the parms operand: on page 58](#).
- Execute the BRBACKUP command.

```
brbackup -t {online_split | offline_split | online_mirror | \ offline_mirror} [-q
split] -d \ util_file -m all -c -u user/password
```

Backing up using Oracle Recovery Manager

If RMAN is used directly, consider the following:

- RMAN stores information about backups in the recovery catalog. For security reasons, keep the catalog in a separate database. This requires more administrative work.
- In a disaster situation (such as the loss of a production database and recovery catalog), the restore and recovery of data is complicated. It may be impossible without the help of Oracle Support. If the Recovery Manager does not have administrative data stored in the recovery catalog, it cannot recover the database only by using the backups that have been made.
- For each RMAN channel, set the `SBT_LIBRARY` parameter to point to the correct platform-specific Data Protector MML. For details on the Data Protector MML location, see [Specify the parms operand: on page 58](#).

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

- The recovery catalog is not used. Information about backups is saved in the control file and SAP R/3 log files. After each backup, the control file and SAP R/3 log files are saved. When data is restored, the control file is copied back first, followed by data files. In case of a disaster, restore SAP R/3 log files before you restore any data files.
- Other important files will still be automatically backed up using the backint program.
- All previous SAP R/3 backup strategies can still be used with RMAN. However, RMAN cannot be used for offline redo log backups with BRARCHIVE, or for standby database backups.
- Ensure that the `SBT_LIBRARY` parameter in the `initSAP_instance.sap` file points to the correct platform-specific Data Protector MML. For details on the Data Protector MML location, see [Specify the parms operand: on page 58](#).

Manual balancing

Manual balancing means that you manually group files into subsets, which are then backed up in parallel. To group files into subsets, add the `manual_balance` section to the Data Protector SAP R/3 configuration file as described in the following example.

Example

Suppose that we have a backup specification named `SAP-R3` with the following files to be backed up: `fileA`, `fileB`, `fileC`, `fileD`. To group the files into three subsets (`0={fileA, fileC}`, `1={fileB}`, `2={fileD}`), add the following lines to the Data Protector SAP R/3 configuration file:

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

When you group files into subsets, consider the following:

- Use only one file from the same hard disk at a time.
- The number of files in a subset must be equal to or smaller than the sum of the concurrencies of all devices specified for backup.
- If the backup specification contains files that are not allocated to any subset, Data Protector automatically adds these files to the list of files to be backed up using the load balancing principle. Before the backup, this list is logged in:

Windows systems: `SAPDATA_HOME\sapbackup*.lst`

UNIX systems: `ORACLE_HOME/sapbackup/*.lst`

Restore

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

- Use the Data Protector GUI. See [Restoring using the Data Protector GUI below](#).
- Use the Data Protector CLI. See [Restoring using the Data Protector CLI on page 150](#).
- Use the SAP restore commands. See [Restoring using the SAP commands on page 151](#).

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

Considerations

- Backups created by Oracle RMAN can only be restored using the SAP restore utilities.
- SAP R/3 tablespaces located on raw partitions cannot be restored using the Data Protector GUI. Workaround: Use SAP restore commands (for example, `brrestore`).
- If you are restoring a sparse file, you can improve the performance by setting the sparse option. See [Sparse files on page 153](#).
- If your Oracle database is localized, you may need to set the appropriate Data Protector encoding before you start a restore. For details, see [Localized SAP R/3 objects on page 152](#).
- Restore preview is not supported.

Restoring using the Data Protector GUI

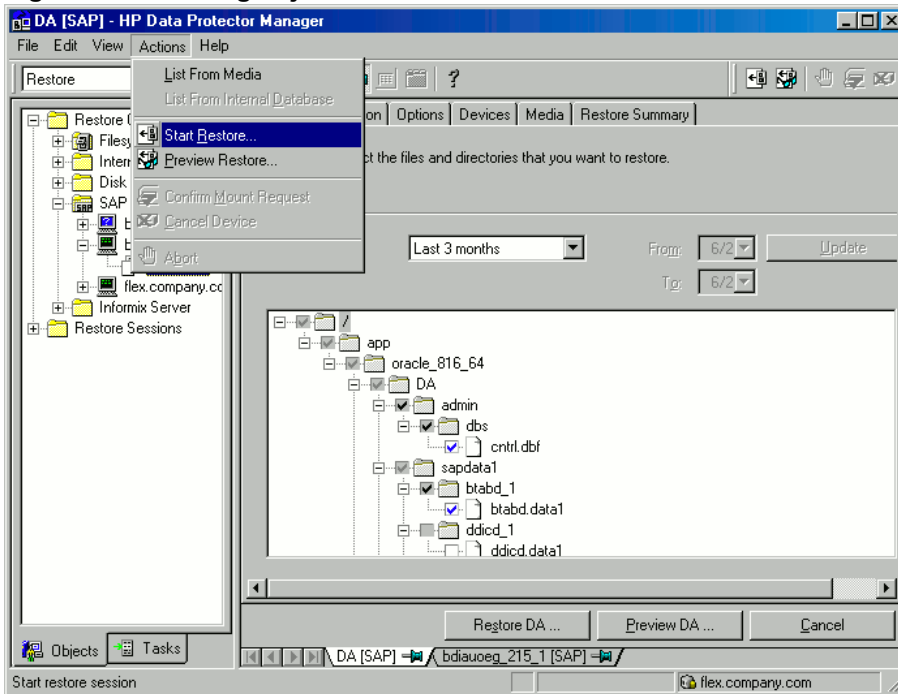
1. In the Context List, click **Restore**.
2. In the Scoping Pane, expand **SAP R/3**, expand the client from which the data was backed up, and then click the Oracle instance you want to restore.

3. In the **Source** page, select SAP R/3 files to be restored.

To restore a file under a different name or to a different directory, right-click the file and click **Restore As/Into**.

To restore a file from a specific backup session, right-click the file and click **Restore Version**.

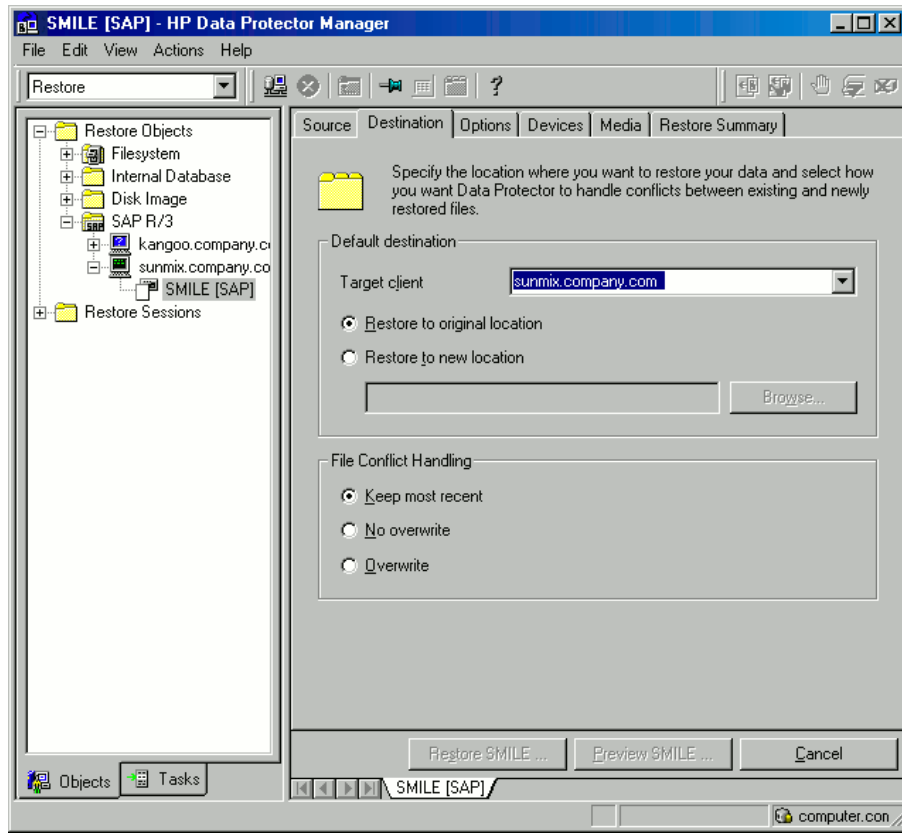
Figure 38: Selecting objects for restore



4. In the **Destination** tab, select the client to restore to (**Target client**). See [Selecting the target client below](#).

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

Figure 39: Selecting the target client



5. In the **Options** page, set the restore options. For information, press **F1**.

6. In the **Devices** page, select the devices to be used for the restore.

For more information on how to select devices for a restore, see the *HP Data Protector Help* index: "restore, selecting devices for".

7. Click **Restore**.

8. In the **Start Restore Session** dialog box, click **Next**.

9. Specify **Report level** and **Network load**.

10. Click **Finish** to start the restore.

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

Restoring using the Data Protector CLI

Execute the following command:

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

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

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

Example (Windows)

To restore the SAP R/3 file `btabd_1.dat` to the original location `C:\oracle\ABA\sapdata1\btabd_1` on the Windows system `computer1.company.com` from the backup session `2011/01/23-1`, execute:

```
omnir -sap computer1.company.com:ABA.0 -session 2011/01/23-1 -tree  
/C:/oracle/ABA/sapdata1/btabd_1/btabd_1.dat
```

Example (UNIX)

To restore the SAP R/3 file `btabd_1.dat` to the original location `/app/oracle/ABA/sapdata1/btabd_1` on the UNIX system `computer2.company.com` from the backup session `2011/01/23-1`, execute:

```
omnir -sap computer2.company.com:ABA.0 -session 2011/01/23-1 -tree  
/app/oracle/ABA/sapdata1/btabd_1/btabd_1.dat
```

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

```
omnidb -sap
```

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

```
omnidb -sap object_name
```

Restoring using the SAP commands

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

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

Windows systems: `set OB2APPNAME=ORACLE_SID`

UNIX systems: `export OB2APPNAME=ORACLE_SID`

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

Windows systems: `set OB2HOSTNAME=client_name`

UNIX systems: `export OB2HOSTNAME=client_name`

4. If you plan do restores in the RMAN mode, ensure that the SBT_LIBRARY parameter in the `initSAP_`

instance.sap file points to the correct platform-specific Data Protector MML. For details on the Data Protector MML location, see [Specify the parms operand: on page 58](#).

5. Execute the SAP restore command.

Restoring using another device

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

Using the Data Protector GUI

For information on how to select another device for a restore using the Data Protector GUI, see the *HP Data Protector Help* index: “restore, selecting devices for”.

Using the Data Protector CLI or SAP commands

If you are restoring using the Data Protector CLI or SAP R/3 commands, specify the new device in the file:

Windows systems: *Data_Protector_program_data\Config\Server\cell\restoredev*

UNIX systems: */etc/opt/omni/server/cell/restoredev*

Use the format:

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

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

Important: Delete this file after use.

On Windows systems, use the Unicode format for the file.

Localized SAP R/3 objects

Oracle Server uses its own encoding, which may differ from the encoding used by the filesystem. In the Backup context, Data Protector displays the logical structure of the Oracle database (with Oracle names) and in the Restore context, the filesystem structure of the Oracle database. Therefore, to display non-ASCII characters correctly, ensure that the Data Protector encoding matches with the Oracle Server encoding during backup and with the filesystem encoding during restore. However, the incorrect display does not impact the restore.

Windows systems: If the current values of DBCS and the default Windows character set for non-Unicode programs do not match, problems arise. See [Restore sessions fail due to invalid characters in filenames on page 163](#).

UNIX systems: To be able to switch between the Data Protector encodings, start the GUI in UTF-8 locale.

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

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

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

Sparse files

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

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

Windows systems: `set OB2SPARSE=sparse`

UNIX systems: `export OB2SPARSE=sparse`

Disaster recovery

For general information, see the *HP Data Protector Disaster Recovery Guide*.

Restoring the control file

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

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

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

2. Execute:

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

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

Monitoring sessions

You can monitor currently running sessions in the Data Protector GUI. When you run an interactive backup or a restore session, a monitor window shows you the progress of the session. Closing the GUI does not affect the session.

You can also monitor sessions from any Data Protector client with the User Interface component installed, using the Monitor context.

On how to monitor a session, see the *HP Data Protector Help* index: “viewing currently running sessions”.

System messages generated during backups are sent to both the SAP R/3 and the Data Protector monitor. However, mount requests are sent only to the Data Protector monitor.

Troubleshooting

This section lists general checks and verifications plus problems you might encounter when using the Data Protector SAP R/3 integration.

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

Before you begin

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

General troubleshooting

Problem

Configuration fails due to a database operation failure

During configuration of an SAP R/3 database, Data Protector reports the following error:

Integration cannot be configured.

The database reported error while performing requested operation.

Action

Review user group membership for the user account which is used in Oracle database access authentication. For details, see [Configuring user accounts on page 125](#).

Problem

Restore session that uses object copies fails

A restore session for a Data Protector SAP R/3 backup object which spans multiple backup media fails and reports the following error:

```
[Major] From: RSM@CMSystemName "" Time: DateTime
```

```
[61:9001] Could not find the object ObjectName named "SAP" in the database. Database error reported is: "Object version not found."
```

This problem occurs only with SAP R/3 backup objects that were copied and when not all original media have been recycled and exported before the session. For such restore sessions, Data Protector selects the original media instead of the media storing the backup object copies. Since some of the media from the original media set can no longer be used, the session fails.

Action

Follow the steps:

1. Recycle and export the remaining media that store original SAP R/3 backup objects.
2. Restart the restore session.

Each time you perform object copy of SAP R/3 backup objects, and start recycling and exporting the original media afterwards, make sure you recycle and export all original media to enable successful restore sessions.

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 SQL Plus from the *ORACLE_HOME* directory:

```
bin\sqlplus
```

At the SQL prompt, type:

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

```
select * from dba_tablespace;
```

```
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:**

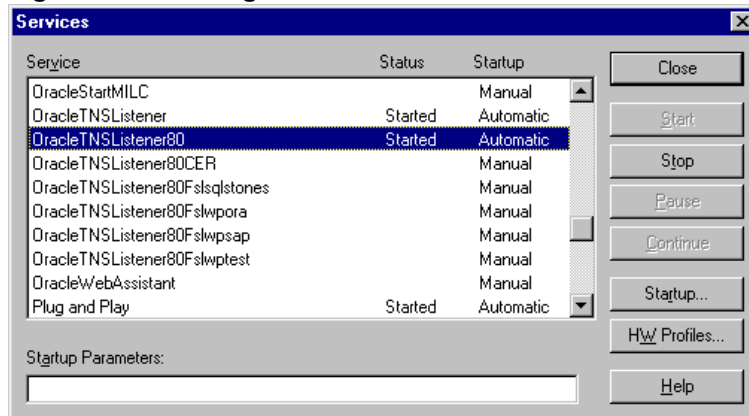
Start the listener from the *ORACLE_HOME* directory:

```
bin\lsnrctl status service  
  
quit
```

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

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

Figure 40: Checking the status of the Oracle listener



- a. The status of the respective listener service in the **Services** window should be **Started**, otherwise you must start it manually.
- b. Start the SQL Plus from the *ORACLE_HOME* directory:

At the SQL prompt, type:

```
connectTarget_Database_Login  
  
exit
```

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

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

Set *ORACLE_HOME* as described in [Verify that you can access the Oracle Target Database and that it is opened, as follows: on page 164](#) and start the Server Manager from the *ORACLE_HOME* directory:

```
bin\svrmgrl
```

At the `wSVRMGR` prompt, type

```
connect Target_Database_Login as SYSDBA;
```

```
exit
```

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

If you are using the recovery catalog:

```
bin\rman target Target_Database_Login rcvcat Recovery_CatLog_Login
```

If you are not using the recovery catalog:

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

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

Prerequisites on the SAP side of the integration

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

1. Verify backup directly to disk as follows:

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

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

2. Verify restore directly to disk as follows:

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

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

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

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

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

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

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

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

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

Protector):

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

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

Then start the following commands:

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

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

See [Configuring SAP R/3 databases on page 128](#).

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.**

For details, see the *HP Data Protector Installation and Licensing Guide*.

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

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

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

See the *HP Data Protector Help* index “standard backup procedure” for details about how to do a filesystem backup.

3. **If the SAP backup utilities are installed in a shared directory, then the inet startup parameter must be specified as described in [If you use the command line to start the Data Protector commands, verify the inet startup parameters. on the next page](#), or the Windows permissions must be set correctly.**

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

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

or

```
dir \\client_name\SAPPEXE\brbackup
```

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

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

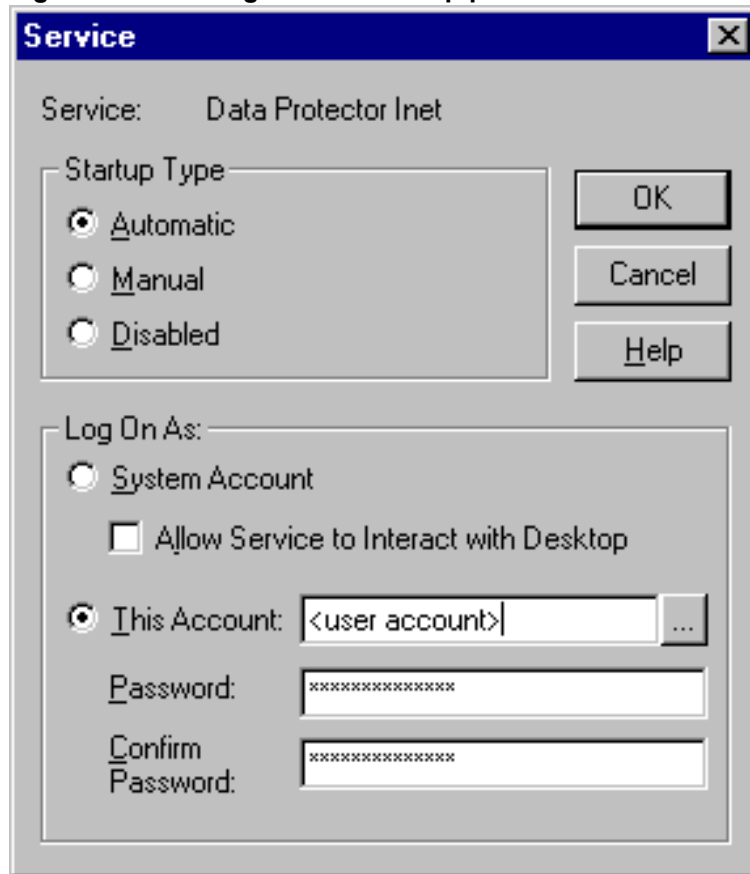
Check the Data Protector Inet service startup parameters on the SAP Database Server system. Proceed as follows:

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

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

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

Figure 41: Checking the Inet start-up parameters



5. **Examine the environment variables.**

If you need to export some variables before starting the Oracle Server Manager, TNS listener, or other Oracle utility, these variables must be defined in the `Environment` section of the Data Protector SAP configuration file on the Cell Manager. See [Data Protector SAP R/3 configuration file on page 118](#).

6. Examine system errors.

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

Problem

Configuration fails due to a script failure

During configuration of an SAP R/3 database, Data Protector reports the following error:

```
Integration cannot be configured.
```

```
Script failed. Cannot get information from remote host.
```

Action

Check the environment settings and ensure Data Protector Inet is running under a user account which has the required privileges. For details, see [Before you begin on page 124](#).

Backup problems

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

1. Check your SAP Server configuration:

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

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

The message `*RETVL*0` indicates successful configuration.

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

Before you run the `testbar2` utility, verify that the Cell Manager name is correctly defined on the SAP Database Server. In the default Data Protector client configuration directory, check the `cell_server` file, which contains the name of the Cell Manager system. Then execute the following command:

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

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

If the messages indicate problems concerning the Data Protector side of the integration, create an SAP backup specification to back up to a null or file device. If the backup succeeds, the problem may

be related to the backup devices. For instructions on troubleshooting devices, see the *HP Data Protector Troubleshooting Guide*. If the test fails again, call support.

3. Verify the backup using backint

```
export OB2BARLIST=barList_name
```

```
export OB2APPNAME=ORACLE_SID
```

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

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

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

Problem

Backup fails with “Connect to database instance failed”

If you start a backup while the database instance is in the unmount or mount mode, the session fails with a message similar to the following:

```
BR0301E SQL error -1033 at location BrDbConnect-2
```

```
ORA-01033: ORACLE initialization or shutdown in progress
```

```
BR0310E Connect to database instance H00H00 failed
```

Action

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

Restore problems

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

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

This can be done by executing the command

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

on the SAP Database Server system.

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

For detailed syntax of the omnidb command, execute:

```
omnidb -help
```

You can also do this using the SAP tools:

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

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

where the specified `input_file` is queried.

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

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

```
backup_ID_1 pathName_1 [targetDirectory_1]
```

```
backup_ID_2 pathName_2 [targetDirectory_2]
```

```
backup_ID_3 pathName_3 [targetDirectory_3]
```

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

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

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

2. Verify the restore using the Data Protector User Interface

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

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

3. Simulate a Restore Session

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

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

In the default Data Protector client configuration directory, check the `cell_server` file, 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
```

```
-apname: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

Execute 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_1pathName_1`
`[targetDirectory_1]backup_ID_2pathName_2` `[targetDirectory_2]backup_ID_3pathName_3`
`[targetDirectory_3]`

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

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

Problem

Restore sessions fail due to invalid characters in filenames

On Windows systems, where the Oracle Database Character Set (DBCS) is not set to the same value as the default Windows character set for non-Unicode programs, and where SAP tools are used to create Oracle datafiles, restore fails if the datafiles contain non-ASCII or non-Latin 1 characters.

Actions

Use any of the following solutions:

- For new Oracle installations, set the DBCS to UTF-8.
- If you do not use other non-Unicode programs, set the language for non-Unicode programs to the same value as DBCS.
- Do not use non-ASCII or non-Latin 1 characters for filenames.

Troubleshooting on UNIX systems

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 SQL Plus from the *ORACLE_HOME* directory:

```
bin\sqlplus
```

At the SQL prompt, type:

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

```
select * from dba_tablespace;
```

```
exit
```

If it fails, open the Oracle Target Database.

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

Export *ORACLE_HOME* as described in [Verify that you can access the Oracle Target Database and that it is opened, as follows: above](#) and start the listener from the *ORACLE_HOME* directory:

```
bin/lsnrctl start service
```

```
exit
```

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

Export *ORACLE_HOME* as described in [Verify that you can access the Oracle Target Database and that it is opened, as follows: on the previous page](#) and start the SQL Plus from the *ORACLE_HOME* directory:

```
bin\sqlplus
```

At the SQL prompt, type:

```
connect Target_Database_Login
```

```
exit
```

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

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

Export *ORACLE_HOME* as described in [Verify that you can access the Oracle Target Database and that it is opened, as follows: on the previous page](#) and start the SQL Plus from the *ORACLE_HOME* directory:

```
bin/svrmgr1
```

At the SQL 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_CatLog_Login
```

If you do not use the Recovery Catalog:

```
bin/rman target Target_Database_Login nocatalog
```

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

4. If you run backups in the RMAN mode, verify backup and restore directly to disk using the Recovery Manager channel type disk.

If you use the Recovery Catalog:

Export *ORACLE_HOME* as described in [Verify that you can access the Oracle Target Database and that it is opened, as follows: on the previous page](#) and start Recovery Manager:

```
bin/rman targetTarget_Database_Loginrcvcat Recovery_Catalog_Login cmd_file=rman_
script
```

If you do not use the Recovery Catalog:

Export `ORACLE_HOME` as described in [Verify that you can access the Oracle Target Database and that it is opened, as follows: on page 164](#) 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 executing 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, see the Oracle documentation to learn how to execute backup and restore directly to disk using the Recovery Manager.

Prerequisites on the SAP side of the integration

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

1. Verify backup directly to disk as follows:

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

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

2. Verify restore directly to disk as follows:

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

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

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

- a. Re-link the Oracle Server with the Database Library provided by SAP (`libobk.sl`).

For each RMAN channel, set the `SBT_LIBRARY` parameter to point to the `libobk.sl` file.

Important: Before you can use Data Protector again in the RMAN mode, you have to re-link the Oracle again with the Data Protector Database Library.

- b. You have to define the parameter `init` in the initialization file `initORACLE_SID.ora`.

Execute the following commands:

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

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

If this fails, see the *SAP Online Help* to learn how to execute backup and restore directly to disk using the SAP backup utility. Check the error message and resolve this issues before you continue.

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

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

```
#!/usr/bin/sh
echo "Test backint called as follows:"
echo "$0 $*"
echo "exiting 3 for a failure"
exit 3
```

Then start the following commands as the Oracle database user described in [Configuring user accounts on page 125](#):

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

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

See [Configuring SAP R/3 databases on page 128](#).

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.

For details, see the *HP Data Protector Installation and Licensing Guide*.

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

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

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

See the *HP Data Protector Help* index “standard backup procedure” for details about how to do a filesystem backup.

3. Examine the environment variables:

If you need to export some variables before starting the Oracle Server Manager, TNS listener, or other Oracle utility, these variables must be defined in the `Environment` section of the Data Protector SAP configuration file on the Cell Manager. See [Data Protector SAP R/3 configuration file on page 118](#).

4. Verify the permissions of the currently used user account:

Your user account has to enable you to perform backup or restore using Data Protector. Use the `testbar2` utility to check the permissions:

```
/opt/omni/bin/utilns/testbar2 -perform:checkuser
```

If the user account holds all required permissions, you will receive only NORMAL messages displayed on the screen.

See also [Configuring user accounts on page 125](#).

5. Examine system errors:

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

Problem

Configuration fails due to a script failure

During configuration of an SAP R/3 database, Data Protector reports the following error:

```
Integration cannot be configured.
```

```
Script failed. Cannot get information from remote host.
```

Action

Resolve the problem by reviewing the user account configuration. For details, see [Configuring user accounts on page 125](#).

Backup problems

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

1. Check your SAP Server configuration:

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

```
/opt/omni/sbin/util_sap.exe -CHKCONF ORACLE_SID (HP-UX, Solaris, and Linux systems) or
```

```
/usr/omni/bin/util_sap.exe -CHKCONF ORACLE_SID (other UNIX systems)
```

In case of an error, the error number is displayed in the form `*RETVAl*Error_number`.

To get the error description, start the command:

```
/opt/omni/sbin/omnigetmsg 12 Error_number(HP-UX, Solaris, and Linux systems) or
```

```
/usr/omni/bin/omnigetmsg 12 Error_number (other UNIX systems)
```

The message `*RETVAl*0` indicates successful configuration.

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

Before you run the `testbar2` utility, verify that the Cell Manager name is correctly defined on the SAP Database Server. Check the `/etc/opt/omni/client/cell_server` (HP-UX, Solaris, and Linux systems) or `/usr/omni/config/cell/cell_server` (other UNIX systems) file, which contains the name of the Cell Manager system. Then execute the following command:

```
/opt/omni/bin/utilns/testbar2 -type:SAP -appname:ORACLE_SID -bar:backup_
specification_name -perform:backup (HP-UX, Solaris, and Linux systems)
```

```
/usr/omni/bin/utilns/testbar2 -type:SAP -appname:ORACLE_SID -bar:backup_
specification_name -perform:backup (other UNIX systems)
```

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

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

- a. Check that the owner of the backup specification is the Oracle OS user described in [Configuring user accounts on page 125](#)
- b. Check that the respective Data Protector user group has the `See private objects` user right enabled.
- c. Create an SAP backup specification to back up to a null or file device. If the backup succeeds, the problem may be related to the backup devices.

For instructions on troubleshooting devices, see the *HP Data Protector Troubleshooting Guide*.

If the test fails again, call support.

3. Verify the backup using backint

```
export OB2BARLIST=barList_name
```

```
export OB2APPNAME=ORACLE_SID
```

```
/opt/omni/sbin/backint -f backup -t file -u ORACLE_SID -i input_file(HP-UX, Solaris,  
and Linux systems)
```

```
/usr/omni/bin/backint -f backup -t file -u ORACLE_SID -i input_file(other UNIX  
systems)
```

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

Backint expects the list of files in the following format: *pathName_1pathName_2pathName_3*

Problem

Util_File_Online SAP backup fails with “semop() error”

When the `util_file_online` option is used with BRBACKUP (for example, if you select the `Brbackup_Util_File_Online` template), the tablespaces are switched into/from backup mode individually. As there can be only one process communicating with BRBACKUP, several `sapback` processes are using a semaphore to synchronize their interaction with BRBACKUP.

The number of `sapback` processes is calculated as the sum of concurrencies of all devices used for backup. With a large number of `sapback` processes, the maximum number of processes that can have undo operations pending on any given IPC semaphore on the system may be exceeded. In such case, several `sapback` agents will fail with the following error:

```
[28] No space left on device.
```

Action

Perform any of the following actions to resolve the problem:

- Reduce the number of backup devices or their concurrency.
- Increase the value of the `semnmu` kernel parameter. After you increase the value, rebuild the kernel and restart the system.

Problem

Backup fails with “Connect to database instance failed”

If you start a backup while the database instance is in the `unmount` or `mount` mode, the session fails with a message similar to the following:

```
BR0301E SQL error -1033 at location BrDbConnect-2
```

```
ORA-01033: ORACLE initialization or shutdown in progress
```

BR0310E Connect to database instance H00H00 failed

Action

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

Restore problems

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

1. Verify a user for the restore:

Verify that user specified for the restore session is the user of backup session and that he/she belongs to the Data Protector operator or admin group.

See [Configuring user accounts on page 125](#).

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

This can be done by executing the command

```
omnidb -sap "object_name" -session "Session_ID" -media (HP-UX, Solaris, and Linux systems) or
```

```
omnidb -sap "object_name" -session "Session_ID" -media (other UNIX systems)
```

on the SAP Database Server system.

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

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

```
omnidb -help (HP-UX, Solaris, and Linux systems)
```

```
omnidb -help (other UNIX systems)
```

You can also do this using the SAP tools:

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

```
/opt/omni/lbin/backint -f inquiry -u ORACLE_SID -i input_file (HP-UX, Solaris, and Linux systems)
```

```
/usr/omni/bin/backint -f inquiry -u ORACLE_SID -i input_file (other UNIX systems)
```

where the specified `input_file` is queried.

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

Backint anticipates a list of files of the following format:

```
backup_ID_1 pathName_1 [targetDirectory_1]
```

```
backup_ID_2 pathName_2 [targetDirectory_2]
```

```
backup_ID_3 pathName_3 [targetDirectory_3]
```

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

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

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

3. Verify the restore using the Data Protector user interface

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

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

4. Simulate a restore session

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

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

Check the */etc/opt/omni/client/cell_server* (HP-UX, Solaris, and Linux systems) or */usr/omni/config/cell/cell_server* (other UNIX systems) file, which contains the name of the Cell Manager system.

Then, test the Data Protector internal data transfer using the *testbar2* utility:

```
/opt/omni/bin/utilns/testbar2 -type:SAP
```

```
-appname:ORACLE_SID
```

```
-perform:restore
```

```
-object:object_name
```

```
-version:object_version
```

```
-bar:backup_specification_name (HP-UX, Solaris, and Linux systems) or
```

```
/usr/omni/bin/utilns/testbar2 -type:SAP
```

```
-appname:ORACLE_SID
```

```
-perform:restore  
  
-object:object_name  
  
-version:object_version  
  
-bar:backup_specification_name (other UNIX systems)
```

You should see only NORMAL messages displayed on your screen, otherwise examine the errors reported by the `testbar2` utility by clicking the **Details** button in the Data Protector **Monitor** context.

5. Verify the restore using backint

Execute the following command:

HP-UX, Solaris, and Linux systems: `/opt/omni/lbin/backint -f restore -u ORACLE_SID -i input_file`

Other UNIX systems: `/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_1pathName_1 [targetDirectory_1]backup_ID_2pathName_2 [targetDirectory_2]backup_ID_3pathName_3 [targetDirectory_3]`

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

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

Problem

Restore of SAP R/3 tablespaces located on raw partitions fails

When restoring SAP tablespaces that are located on raw partitions using the Data Protector GUI, the restore fails with a message similar to the following:

```
[Major] From: VRDA@joca.company.com "SAP" Time: 5/9/06 3:33:51 PM  
/dev/sapdata/rsapdata Cannot restore -> rawdisk section !  
[Warning] From: VRDA@joca.company.com "SAP"  
Time: 5/9/06 3:42:45 PM Nothing restored.
```

Action

Use SAP commands (for example, `brrestore`) to restore these tablespaces.

Chapter 3: Data Protector SAP MaxDB integration

Introduction

This chapter explains how to configure and use the Data Protector SAP MaxDB integration (**SAP MaxDB integration**). It describes the concepts and methods you need to understand to back up and restore SAP MaxDB database objects (**SAP MaxDB objects**).

Data Protector integrates with SAP MaxDB Server to offer online backup of an SAP MaxDB Server instance (**SAP MaxDB instance**). You can back up the following SAP MaxDB objects using the Data Protector SAP MaxDB integration:

- SAP MaxDB data
- SAP MaxDB configuration
- SAP MaxDB archive logs

During backup, the database is online and actively used. It can be in the Admin or Online mode.

Data Protector offers interactive and scheduled backups of the following types:

Table 14: Backup types

Full	SAP MaxDB complete backup. Backs up all selected objects.
Diff	SAP MaxDB incremental backup. Backs up changes made to the database since the last full backup. ¹
Trans	SAP MaxDB log backup. Backs up archived logs. ²

You can restore SAP MaxDB objects:

- To the original location
- To another SAP MaxDB client
- To another SAP MaxDB instance

As part of the restore session, you can also recover the database to a specific point in time or to the last archive log.

You can also back up and restore SAP MaxDB objects using SAP MaxDB utilities.

¹ What is actually backed up depends on which objects you select. For details, see [What is backed up on page 182](#).

² What is actually backed up depends on which objects you select. For details, see [What is backed up on page 182](#).

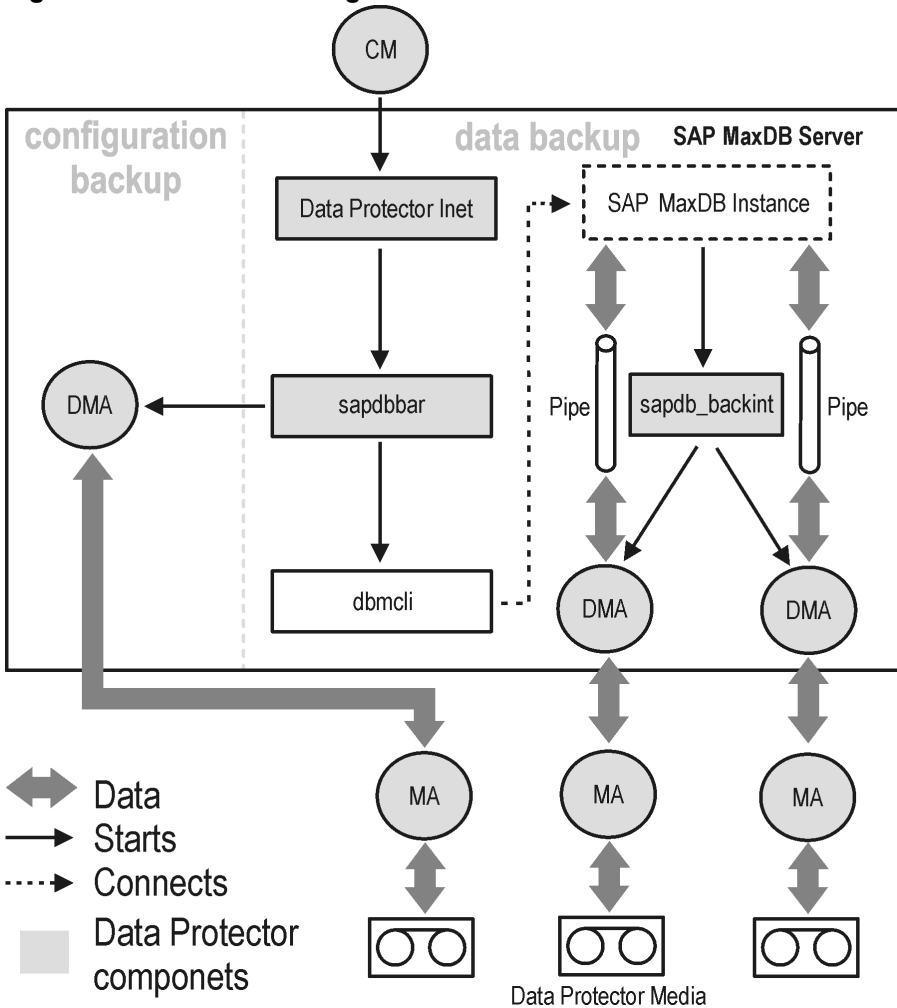
This chapter provides information specific to the Data Protector SAP MaxDB integration. For general Data Protector procedures and options, see the *HP Data Protector Help*.

Integration concepts

Data Protector integrates with the SAP MaxDB Server through the SAP MaxDB integration component using the SAP MaxDB database management server and the `backint` interface.

[SAP MaxDB integration architecture below](#) shows the architecture of the Data Protector SAP MaxDB integration.

Figure 42: SAP MaxDB integration architecture



The Data Protector integration software consists of the following components:

- The `sapdbbar` module, installed on the SAP MaxDB Server system, which controls activities between the SAP MaxDB Server and Data Protector backup and restore processes.
- The `sapdb_backint` component, installed on the SAP MaxDB Server system, is a binary interface between Data Protector and backup and restore functionality of the SAP MaxDB.

- The DMA (Data Mover Agent) component, installed on the SAP MaxDB Server system, is the actual data transferring module, called by the `sapdb_backint`.
- The `util_sapdb` utility, which is used by Data Protector to configure an SAP MaxDB instance to use with Data Protector and check the instance configuration.

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

The integration also takes advantage of the concept of **SAP MaxDB media** and **media groups**, thus providing parallel backup and restore of SAP MaxDB objects. Several SAP MaxDB media are grouped in an SAP MaxDB media group, which is then backed up or restored in streams. This is referred to as SAP MaxDB **parallelism**. See [SAP MaxDB backup options on page 184](#) for more information on the Data Protector Parallelism option.

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

Backup flow

When a backup session is started, the Cell Manager starts the `sapdbbar` with selected backup parameters from the backup specification. The `sapdbbar` module then starts an SAP MaxDB session using the SAP MaxDB `dbmcli`. The `sapdbbar` module issues `dbmcli` commands that configure SAP MaxDB backup media (parallelism), configure `sapdb_backint` and then start the backup using SAP MaxDB `dbmcli`. SAP MaxDB then starts the configured `sapdb_backint` component. For every SAP MaxDB medium (pipe) `sapdb_backint` starts a DMA, which transfers the data from SAP MaxDB media (pipes) to Data Protector media. This procedure is the same for full, differential, and transactional backup. Additionally, if the configuration (including media specification and the backup history) is selected for backup, it is backed up directly by the `sapdbbar` module and DMA. The list of configuration files to be backed up is retrieved through `dbmcli`.

Restore flow

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

Configuring the integration

You need to configure SAP MaxDB users and every SAP MaxDB instance you intend to back up from or restore to.

Prerequisites

- Ensure that you have correctly installed and configured the SAP MaxDB system.
 - See the latest support matrices at <http://support.openview.hp.com/selfsolve/manuals> for supported versions, platforms, devices, and other information.
 - See SAP MaxDB documentation for information on installing, configuring, and using SAP MaxDB Server.

To enable transactional backups (log backups), you need to activate the SAP MaxDB Automatic Log Backup.

- Ensure that you have correctly installed Data Protector. See the *HP Data Protector Installation and Licensing Guide* on how to install Data Protector in various architectures.

Every SAP MaxDB system you intend to back up from or restore to must have the Data Protector SAP MaxDB Integration component installed.

Limitations

The following are not supported:

- Instance names in the Unicode format
- Pre- and post-exec options on the level of the backup specification
- Preview for SAP MaxDB restore sessions
- Integrated offline backup of SAP MaxDB objects

Before you begin

- Configure devices and media for use with Data Protector.
- To test whether the SAP MaxDB system and the Cell Manager communicate properly, configure and run a Data Protector filesystem backup and restore on the SAP MaxDB system.

Cluster-aware clients

Configure SAP MaxDB instances only on one cluster node, since the configuration files reside on the Cell Manager.

If you intend to use the Data Protector CLI, set the Data Protector environment variable OB2BARHOSTNAME to the virtual server name as follows:

Windows systems: `set OB2BARHOSTNAME=virtual_server_name`

UNIX systems: `export OB2BARHOSTNAME=virtual_server_name`

Configuring SAP MaxDB users

Create or identify an **SAP MaxDB database user** with at least the following SAP MaxDB permissions:

- Saving backups (Backup)
- Restoring backups (Recovery)
- Installation management (InstallMgm)
- Parameter access (ParamCheckWrite)

The last two permissions are required for the Data Protector configuration.

UNIX systems: Add the OS user under whose account SAP MaxDB is running (**SAP MaxDB OS user**) or a user belonging to `sapdb admin` group to the Data Protector `admin` or `operator` group. For more information, see the *HP Data Protector Help* index: “adding users”. For example, by default, the SAP MaxDB OS user is the user `sapdb` in the group `sapsys`.

Configuring SAP MaxDB instances

You need to provide Data Protector with the following configuration parameters for the SAP MaxDB instance:

- Username of the SAP MaxDB database user.
- Password of the SAP MaxDB database user.
- Optionally, the SAP MaxDB independent program path parameter

To configure an SAP MaxDB instance, use the Data Protector GUI or CLI.

Data Protector then creates the SAP MaxDB instance configuration file on the Cell Manager and verifies the connection to the instance.

Tip: Once the configuration file is created, you can set, retrieve, and list the configuration file parameters using the Data Protector `util_cmd` command. For details, see the `util_cmd` man page.

To configure an SAP MaxDB instance, use the Data Protector GUI or CLI.

Before you begin

- Ensure that the SAP MaxDB instance is online.

Using the Data Protector GUI

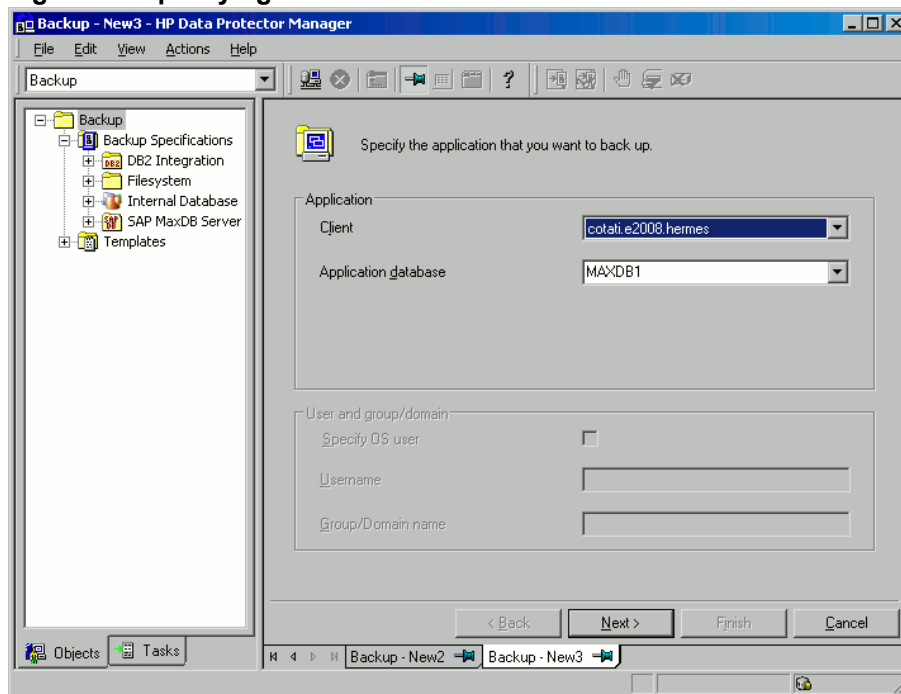
1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **SAP MaxDB Server**, and click **Add Backup**.
3. In the **Create New Backup** dialog box, select the **Blank SAP MaxDB Backup** template. Click **OK**.

In **Client**, select the SAP MaxDB Server system. In a cluster environment, select the virtual server.

In **Application database**, type the SAP MaxDB instance name.

For information on the **User and group/domain** options, press **F1**.

Figure 43: Specifying an SAP MaxDB instance



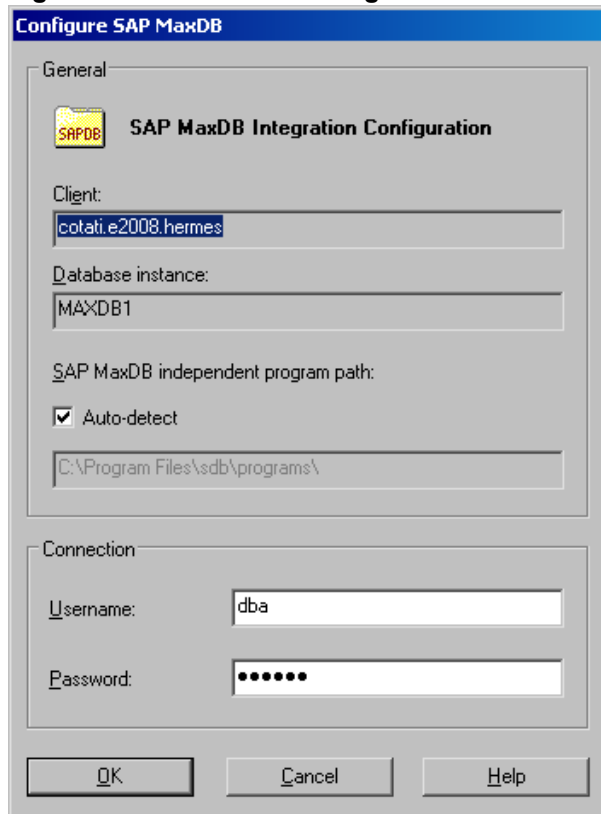
4.

Click **Next**.

5. In the **Configure SAP MaxDB** dialog box, specify the **SAP MaxDB independent program path** parameter. This parameter is the independent program path directory specified during the installation of the SAP MaxDB application. To automatically detect the directory, leave the **Auto-detect** option selected.

Under **Connection**, type the username and password of the SAP MaxDB database user as described in [Configuring SAP MaxDB users on the previous page](#).

Figure 44: SAP MaxDB configuration



Click **OK**.

6. The SAP MaxDB instance is configured. Exit the GUI or proceed with creating the backup specification at [In the Source page, select objects for restore. on page 194.](#)

Using the Data Protector CLI

Log on to the SAP MaxDB Server system as the SAP MaxDB OS user and execute:

```
util_sapdb \[-homedir SAPMaxDB_independent_program_directory] \-config Instance  
Nameusernamepassword
```

Parameter description

<i>SAPMaxDB_independent_program_directory</i>	<p>The SAP MaxDB independent program path parameter. This parameter is the independent program path directory specified during the installation of the SAP MaxDB application on the SAP MaxDB Server.</p> <p>This parameter is optional. If it is not specified, the directory is detected automatically.</p>
<i>Instance_Name</i>	<p>The name of the SAP MaxDB instance to be configured.</p>

<i>username</i>	The username of the SAP MaxDB database user created or identified as described in Configuring SAP MaxDB users on page 178 .
<i>password</i>	The password of the SAP MaxDB database user created or identified as described in Configuring SAP MaxDB users on page 178 .

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

The message *RETVAl*0 indicates successful configuration.

Example

To configure the instance `sapmaxdb_inst` by specifying the database user `sapmaxdb_user` with the password `sapmaxdb_pass`, and the SAP MaxDB independent program path `/opt/sapdb/indep_prog` (UNIX) or `c:\program files\sapdb\indep_prog` (Windows), execute:

Windows systems:

```
util_sapdb -homedir "SAP_MaxDB_independent_program_directory" -config sapdb_inst  
sapdb_user sapdb_pass
```

UNIX systems:

```
util_sapdb -homedir SAP_MaxDB_independent_program_directory/indep_prog -config sapdb_  
inst sapdb_user sapdb_pass
```

Tip: To change the configuration parameters, execute the same command using new values.

Handling errors

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

UNIX systems: To obtain an error description, change the directory to `/opt/omni/lbin` and execute:

```
omnigetmsg 12 Error_number
```

Checking the configuration

Check the configuration of an SAP MaxDB instance after you have created at least one backup specification for the SAP MaxDB instance. Use the Data Protector GUI or CLI.

Using the Data Protector GUI

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

Using the Data Protector CLI

UNIX systems: Log on to the SAP MaxDB Server system as the SAP MaxDB OS user.

Execute the following command:

```
util_sapdb -chkconf Instance_Name
```

where *Instance_Name* is the name of the SAP MaxDB instance.

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

Backup

The integration provides online database backups of different types. What is backed up depends on which objects and backup type you select. See [What is backed up below](#).

Table 15: What is backed up

		SAP MaxDB backup mode		
		Full	Diff	Trans
GUI selection	Data	data	diff on data	archive logs
	Configuration	configuration	configuration	configuration
	Instance	data + configuration	diff on data + configuration	archive logs + configuration

Creating backup specifications

Create a backup specification using the Data Protector Manager.

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **SAP MaxDB Server**, and click **Add Backup**.
3. In the **Create New Backup** dialog box, select the **Blank SAP MaxDB Backup** template. Click **OK**.
4. In **Client**, select the SAP MaxDB Server system. In a cluster environment, select the virtual server.

In **Application database**, type the SAP MaxDB instance name.

For information on the **User and group/domain** options, press **F1**.

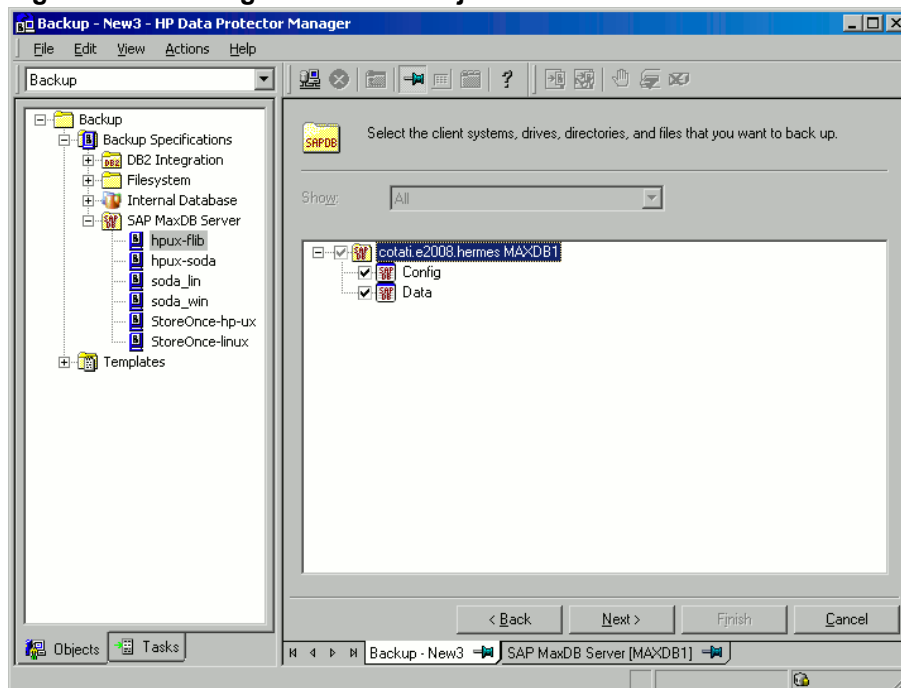
Click **Next**.

5. If the SAP MaxDB instance is not configured yet for use with Data Protector, the **Configure SAP MaxDB** dialog box is displayed. Configure it as described in [Configuring SAP MaxDB instances on page 178](#).

Select the SAP MaxDB objects you want to back up.

Important: To back up SAP MaxDB archive logs, select the **Data** item. The archive log backup is then triggered by selecting the **Trans** backup type when scheduling the backup or running the backup interactively.

Figure 45: Selecting SAP MaxDB objects



- 6.
7. Select devices to use for the backup.

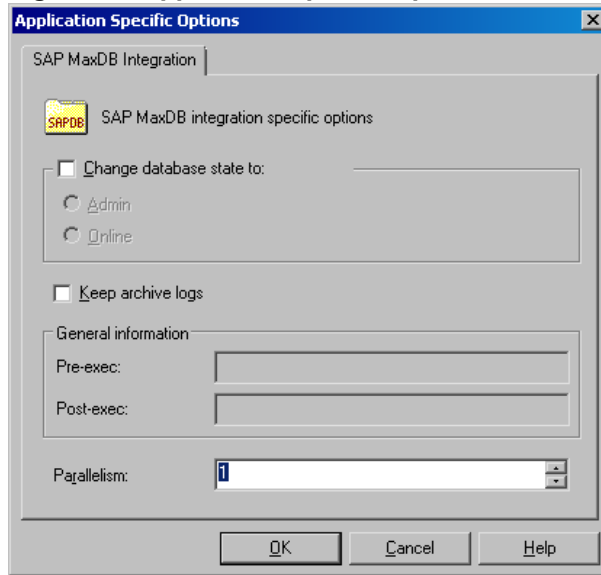
To specify device options, right-click the device and click **Properties**. Specify the device **Concurrency**, media pool, and preallocation policy.

Click **Next**.

8. Set backup options. For information on application-specific options ([Application-specific options on the next page](#)), see [SAP MaxDB backup options on the next page](#).

Click **Next**.

Figure 46: Application-specific options



9. Optionally, schedule the backup. See [Scheduling backup sessions on the next page](#).

Click **Next**.

10. Save the backup specification, specifying a name and a backup specification group.

Tip: Save the backup specifications in the group **SAP MaxDB Integration**.

Tip: Preview your backup specification before using it for real. See [Previewing backup sessions on page 186](#).

Table 16: SAP MaxDB backup options

Option	Description
Change database state	Specifies the SAP MaxDB database mode during backup (Admin or Online). If this option is OFF, the database remains in the current mode.
Keep archive logs	Specifies whether to keep (ON) or delete (OFF) archive logs on the SAP MaxDB Server after the backup has finished.

Option	Description
Parallelism	<p>Specifies the number of SAP MaxDB media created on the SAP MaxDB Server and consequently the number of SAP MaxDB backup data streams.</p> <p>The value must be equal to or lower than:</p> <ul style="list-style-type: none">• The SAP MaxDB MAXBACKUPDEVS parameter.• The sum of concurrency values of all backup devices selected in the backup specification. <p>For more information on the Data Protector Concurrency option, see the <i>HP Data Protector Help</i> index: “concurrency”.</p> <p>Default value: 1.</p> <p>Maximum value: 32.</p> <p>Recommended value: the number of SAP MaxDB data volumes to be backed up.</p>

Modifying backup specifications

To modify your backup specification, click its name in the Scoping Pane of the Backup context, then click the appropriate tab, and apply the changes.

Scheduling backup sessions

You can run unattended backups at specific times or periodically. For details on scheduling, see the *HP Data Protector Help* index: “scheduled backups”.

Scheduling example

To back up SAP MaxDB objects at 8:00, 13:00, and 18:00 during week days:

1. In the **Schedule** property page, select the starting date in the calendar and click **Add** to open the **Schedule Backup** dialog box.
2. Under **Recurring**, select **Weekly**. Under **Time options**, select **8:00**. Under **Recurring Options**, select **Mon, Tue, Wed, Thu, and Fri**. See [Scheduling the backup session on the next page](#).

Click **OK**.

3. Repeat [In the Schedule property page, select the starting date in the calendar and click Add to open the Schedule Backup dialog box.](#) above and Under **Recurring**, select **Weekly**. Under **Time options**, select **8:00**. Under **Recurring Options**, select **Mon, Tue, Wed, Thu, and Fri**. See [Scheduling the](#)

[backup session on the next page.](#) above to schedule backups at 13:00 and 18:00.

4. Click **Apply** to save the changes.

Figure 47: Scheduling the backup session

The screenshot shows the 'Schedule Backup' dialog box. It has a title bar 'Schedule Backup' and a subtitle 'Specify the desired backup time, frequency, duration, and type.' The dialog is divided into several sections: 'Recurring' with radio buttons for 'None', 'Daily', 'Weekly' (selected), and 'Monthly'; 'Time options' with a 'Time' field set to '8:00' and a 'Use starting' checkbox; 'Recurring options' with 'Every' set to '1' week(s) on and checkboxes for 'Sun', 'Mon', 'Tue', 'Wed', 'Thu', 'Fri', and 'Sat'; and 'Session options' with dropdown menus for 'Backup type' (Full), 'Network load' (Full), and 'Backup protection' (Default). At the bottom are 'OK', 'Cancel', and 'Help' buttons.

Previewing backup sessions

Preview the backup session using the Data Protector GUI or CLI to test it.

This interactive test does not back up any data. However, as a result of this test, the *Backup_Specification_Name_TEST_FILE* file is created on the SAP MaxDB Server system in the default Data Protector temporary files directory. Delete the file after the test.

Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **SAP MaxDB Server**. Right-click the backup specification you want to preview and click **Preview Backup**.
3. Specify the **Backup type** and **Network load**. Click **OK**.

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

Using the Data Protector CLI

Execute the following command:

```
omnib -sapdb_list backup_specification_name -test_bar
```

What happens during the preview?

1. The sapdbbar program is started, which then starts the Data Protector testbar2 command.
2. Data Protector tests the Data Protector part of the configuration. The following are tested:
 - Communication between the SAP MaxDB instance and Data Protector
 - The syntax of the backup specification
 - If devices are correctly specified
 - If the necessary media are in the devices

Starting backup sessions

Interactive backups are run on demand. They are useful for urgent backups or restarting failed backups.

Note: If the backup you want to perform is the first backup after the restore, you must select the full backup type.

Prerequisites

- To be able to back up MaxDB history files, make sure that the pathname where these files are located does not contain spaces. If the pathname contains spaces, relocate the history files.

Backup methods

Start a backup of the SAP MaxDB objects selected in the backup specification in any of the following ways:

- Use the Data Protector GUI.
- Use the Data Protector CLI.
- Use the SAP MaxDB utilities.

Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, and then **SAP MaxDB Integration**. Right-click the backup specification you want to use and click Start Backup.
3. Select the **Backup type** and **Network load**. Click **OK**.

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

Using the Data Protector CLI

Log on to the SAP MaxDB Server system as the SAP MaxDB OS user and execute the following command:

```
omnib -sapdb_list ListName [-barmode sapdbmode] [list_options] [-preview]
```

ListName is the name of the backup specification.

sapdbmode specifies the backup type. You can select full, diff, or trans.

For *List_options*, see the omnib man page.

Example

To start a full backup using an existing SAP MaxDB backup specification called TEST, and to set data protection to 10 weeks, execute:

```
omnib -sapdb_list TEST -barmode full -protect weeks 10
```

Using SAP MaxDB utilities

For a description of the variables listed below, see [Parameter description on page 191](#).

1. Create the `bsi_env` file on the SAP MaxDB Server system.

UNIX systems: Grant the SAP MaxDB OS user read permission for this file.

The file must contain the following lines:

Windows systems:

```
BACKINT Data_Protector_home\bin\sapdb_backint  
INPUT Data_Protector_program_data\tmp\inst_name.bsi_in  
OUTPUT Data_Protector_program_data\tmp\inst_name.bsi_out  
ERROROUTPUT Data_Protector_program_data\tmp\inst_name.bsi_err  
PARAMETERFILE name_of_backup_spec  
TIMEOUT_SUCCESS 900  
TIMEOUT_FAILURE 30
```

UNIX systems:

```
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 SAPDB_PARAMETER  
TIMEOUT_SUCCESS 900  
TIMEOUT_FAILURE 30
```

The `SAPDB_PARAMETER` file should contain the following lines:

```
OB2BARLIST=name_of_backup_spec  
OB2APPNAME=inst_name  
OB2BARHOSTNAME=FQDN
```

2. Log in to the SAP MaxDB database manager as the SAP MaxDB database user by executing:

```
dbmcli -d inst_name -u username,password
```

3. In the SAP MaxDB database manager, register the location of the `bsi_env` file created in [Create the bsi_env file on the SAP MaxDB Server system. on the previous page](#) of this procedure by executing:

Windows systems:

```
dbm_configset -raw BSI_ENV Location\inst_name.bsi_env
```

UNIX systems:

```
dbm_configset -raw BSI_ENV Location/inst_name.bsi_env
```

4. Create SAP MaxDB media, grouping them under the same name (*media_group_name*). The number of created media should equal the parallelism you plan to use for backup. To create a medium *medium_name*, execute the following command, depending on the SAP MaxDB version:

- For SAP MaxDB version 7.6:

```
medium_put media_group_name/medium_name pipe_name type backup_type [size [block_size [overwrite [autoLoader [os_command [tool_type]]]]]]
```

- For other SAP MaxDB versions:

```
medium_put media_group_name/medium_name pipe_name medium_type backup_type
```

backup_type can be one of the following:

- DATA for full backup
- PAGES for differential backup
- LOG for log backup

tool_type must be the following:

- "BACK" for backup with Backint for SAP MaxDB

Important: When creating SAP MaxDB media for the purpose of a Data Protector backup and restore, the media group name must begin with the "BACK" string.

Example

The commands below create two media and two pipes (parallelism = 2) in the media group BACKDP-Data[2].

Windows systems, SAP MaxDB version 7.6:

```
medium_put BACKDP-Data[2]/1 \  
\\.\Pipe\inst_name.BACKDP_Data[2].1 PIPE DATA 0 8 \  
NO NO \" \" "BACK"
```

```
medium_put BACKDP-Data[2]/2 \  
\\.\Pipe\inst_name.BACKDP_Data[2].2 PIPE DATA 0 8 \  
NO NO \" \" "BACK"
```

UNIX systems, SAP MaxDB version 7.6:

```
medium_put BACKDP-Data[2]/1 \  
/var/opt/omni/tmp/inst_name.BACKDP_Data[2].1 PIPE \  
DATA 0 8 NO NO \" \" "BACK"
```

```
medium_put BACKDP-Data[2]/2 \  
/var/opt/omni/tmp/inst_name.BACKDP_Data[2].2 PIPE \  
DATA 0 8 NO NO \" \" "BACK"
```

Windows systems, other SAP MaxDB versions:

```
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 systems, other SAP MaxDB versions:

```
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 MaxDB utility session by executing:

```
util_connect
```

6. Start the backup. The following exemplary command starts the full backup for the media created in [Create SAP MaxDB media, grouping them under the same name \(media_group_name\)](#). The number of created media should equal the parallelism you plan to use for backup. To create a medium `medium_name`, execute the following command, depending on the SAP MaxDB version: on the [previous page](#) of this procedure:

```
backup_start BACKDP-Data[2] DATA
```

7. The progress of the session is displayed in the Data Protector **Monitor** context. For more

information, see [Monitoring sessions on page 206](#).

Parameter description

<i>inst_name</i>	Name of the instance to be backed up.
<i>name_of_backup_spec</i>	Name of the Data Protector backup specification to be used for backup.
<i>username, password</i>	Connection string for the SAP MaxDB database user.
<i>location</i>	Location of the <code>bsi_env</code> file.
<i>media_group_name</i>	Name of the SAP MaxDB media group.
<i>medium_name</i>	Name of the SAP MaxDB medium.
<i>pipe_name</i>	Name of the SAP MaxDB pipe.
<i>medium_type</i>	Type of the SAP MaxDB medium.

Restore

Restore SAP MaxDB objects in any of the following ways:

- Use the Data Protector GUI. See [Restoring using the Data Protector GUI on page 194](#).
- Use the Data Protector CLI. See [Restoring using the Data Protector CLI on page 197](#).
- Use the SAP MaxDB utilities. See [Restoring using SAP MaxDB utilities on page 198](#).

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 MaxDB restore options on page 202](#).

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

Depending on the type of restore and on the selected restore and recovery options, the SAP MaxDB database can be switched to the following modes after the restore:

- If the Data Protector Recovery option is selected, the database is switched to the `Online` mode after the restore.
- If the Data Protector Recovery option is *not* selected and archive logs have not been restored (if restore from a full or diff backup session is performed), the database remains in the `Admin` mode after the restore.
- If the Data Protector Recovery option is *not* selected and archive logs have been restored, the

database is, if the restored archive logs allow it, switched to the `Online` mode. If the database, however, cannot be switched to the `Online` mode (because the restored archive logs do not allow it), it remains in the `Admin` mode.

Important: There are several scenarios, depending on the backup option `Keep archive logs` and the recovery option `Use existing archive logs`, in which a gap of transactions between the sequence of redo logs on the SAP MaxDB Server and the restored volumes can occur. When performing recovery (when the database is switched to the `Online` mode), SAP MaxDB always checks whether such a gap exists, regardless of the point in time selected for recovery. If such a gap exists, the recovery is not performed and the database remains in the `Admin` mode, unless the existing redo logs are manually deleted before starting the restore.

If a full or diff backup session is restored, only the data (no archive logs) from the selected backup session is restored. The data on the SAP MaxDB Server is overwritten.

If a trans backup session is restored, only the archive logs (no data) from the selected backup session are restored.

During the restore, the redo logs that existed on the SAP MaxDB Server before the restore are not deleted during the restore.

When restoring, the existing redo logs on the SAP MaxDB Server can be, depending on the Data Protector `Use existing archive logs` option selection (it can be selected only if the `Recovery` option is selected), handled as follows:

- If the `Use existing archive logs` option is selected, the existing archive logs on the SAP MaxDB Server are applied to the redo logs.

When a transactional backup session is selected for restore, or when it is a part of the needed restore chain, and the `Use existing archive logs` option is selected at the same time, the archive logs from Data Protector media are applied to redo logs. Thereafter, the archive logs on the SAP MaxDB Server are applied to redo logs.

- If the `Use existing archive logs` option is not selected, the backed up archive logs on backup media are applied to the redo logs (if trans backup session is restored), or the redo logs are left intact together with the existing archive logs on the SAP MaxDB Server (if full or diff backup session is restored).

Note: The `Use existing archive logs` option is disabled in case of SAP MaxDB migration, thus allowing only for the restore of redo logs from the backed up archive logs on backup media (if trans backup session is restored).

Figure 48: SAP MaxDB restore process

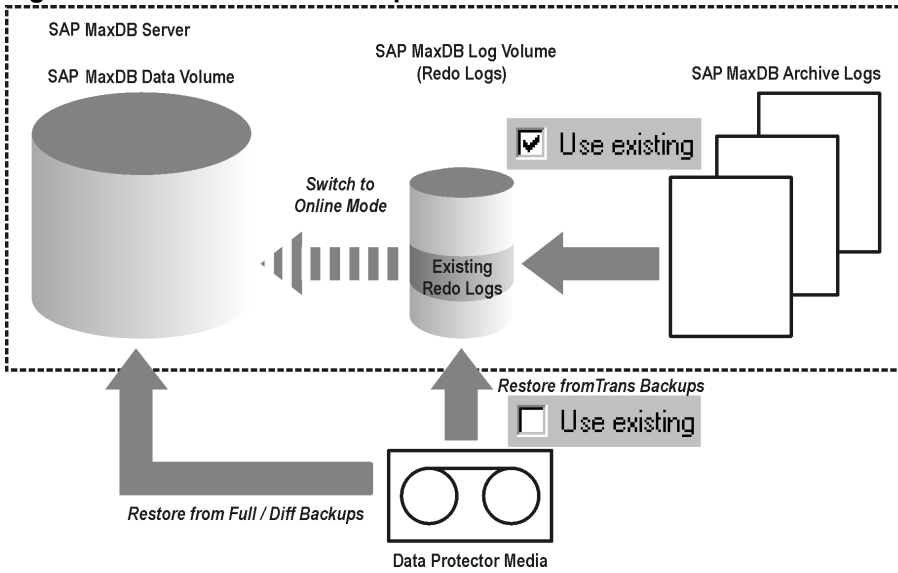
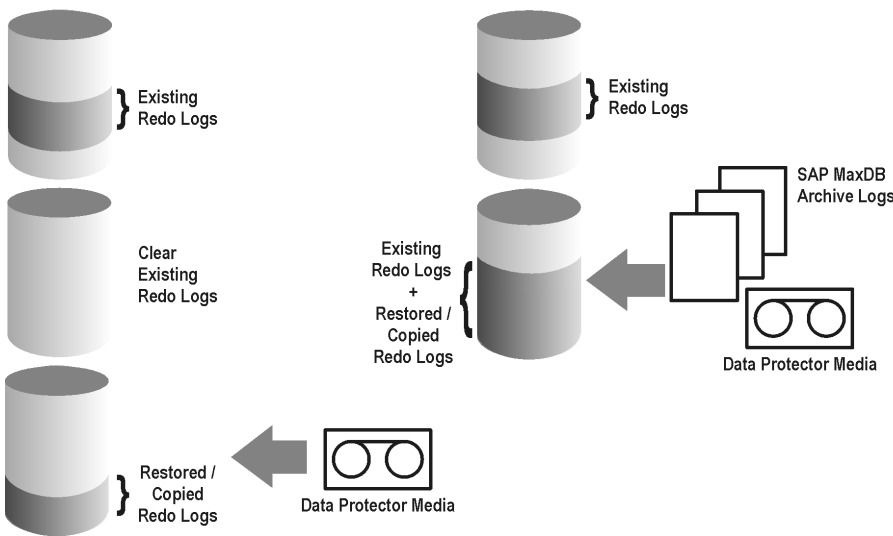


Figure 49: SAP MaxDB archive logs restore process—redo logs details



If you select a differential or a transactional backup session to be restored, you can set the integration to:

- Perform a full database restore. In this case, the integration automatically determines the chain of needed full, differential or transactional backup sessions when performing the restore. After the restore has finished, the database is, if the Recovery option is selected, switched to the Online mode.
- Restore only the selected differential or the selected transactional backup session. If the database is consistent after such a restore and if the Recovery option is selected, it is switched to the Online mode. Otherwise, the database is left in the Admin mode.

Restoring only the selected trans or diff backup session is useful if the database remains offline or in the Admin mode after a restore from full backup session, which is then followed by a restore from diff or trans backup session.

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

Note: After you perform restore or recovery, the first backup you will perform must be a full backup.

Before you begin

If you intend to restore to another SAP MaxDB instance:

- Install the Data Protector SAP MaxDB integration on the SAP Max DB Server system to which you want to restore.
- Add the SAP MaxDB client to the Data Protector cell.
- Configure SAP MaxDB users as described in [Configuring SAP MaxDB users on page 178](#).
- Configure the instance to which you want to perform the restore. See [Configuring SAP MaxDB instances on page 178](#).

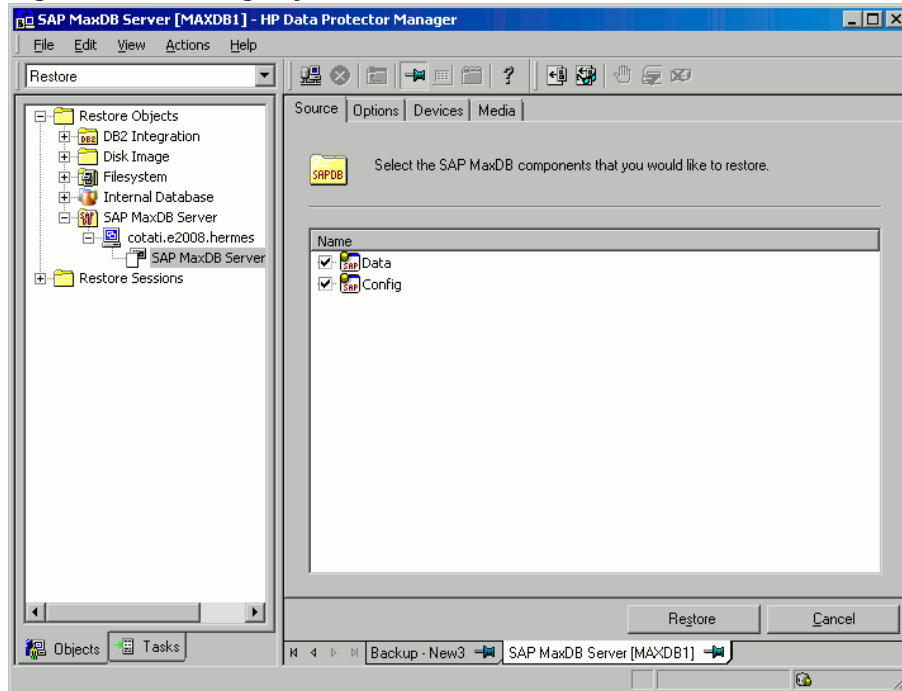
Note: If you are using the Data Protector GUI, you can configure the instance during the restore process.

During the restore to another SAP MaxDB instance, the existing data is overwritten and the existing redo logs are deleted.

Restoring using the Data Protector GUI

1. In the **Context List**, click **Restore**.
2. In the Scoping Pane, expand **SAP MaxDB Server**, expand the client from which the data to be restored was backed up, and then click the SAP Max DB instance you want to restore.
3. In the **Source** page, select objects for restore.

Figure 50: Selecting objects for restore



4. To restore SAP MaxDB objects from a specific backup session, right-click the **Data** item, click **Properties**, and specify **Backup version** in the **Properties for Data** dialog box.

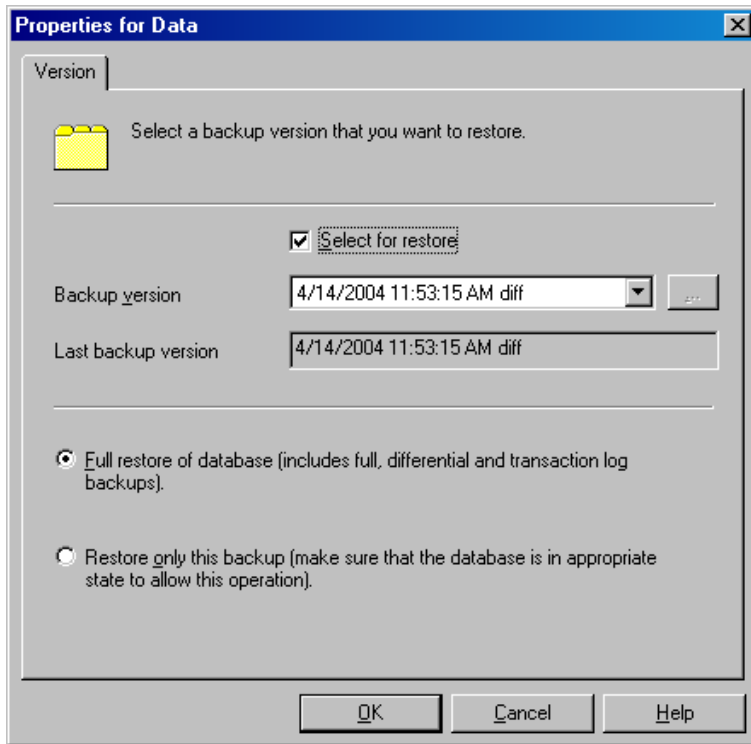
Selecting a Trans or Diff backup session enables you to:

- Perform a full restore of the database (Full restore of database). In this case, the integration automatically determines the chain of needed full, differential, or transactional backup sessions.
- Restore only the selected backup session (Restore only this backup).

Restoring only the selected Trans or Diff backup session is useful if the database remains offline or in the Admin mode after a restore from a full backup session.

5. To restore SAP MaxDB archive logs, select the Data item and a Trans backup session to restore from.

Figure 51: Properties for data



Important: The Configuration item is restored from the same backup session as selected for the Data item, regardless of what you select for the Configuration item.

6. In the **Options** page, set the restore and recovery options. For information, see [SAP MaxDB restore options on page 202](#).
7. In the **Devices** page, select the devices to be used for the restore.

For more information on how to specify devices for a restore, see the *HP Data Protector Help* index: "restore, selecting devices for".

8. In the **Media** page, view the media needed for the restore and verify its availability.
9. Click **Restore**.
10. In the **Start Restore Session** dialog box, click **Next**.
11. Specify the **Report level** and **Network load**.

Note: Select **Display statistical information** to view the restore profile messages in the session output.

12. Click **Finish** to start the restore.

The statistics of the restore session, along with the message `Session completed successfully` is displayed at the end of the session output.

Restoring using the Data Protector CLI

Log on to the SAP MaxDB Server system as the SAP MaxDB OS user and execute the following command:

```
omnir -sapdb -barhost ClientName -instance InstanceName  
[-destination ClientName]  
[-newinstance DestinationInstanceName]  
[-session BackupID]  
[-recover [-endlogs | -time: YYYY-MM-DD.hh.mm.ss] [-from_disk]]  
[-nochain]
```

The `-barhost` option sets the name of the SAP MaxDB Server that was backed up.

The `-instance` option sets the name of the SAP MaxDB instance that was backed up.

The `-session` specifies from which backup data (*BackupID*) to restore. A backup ID is a point in time. All objects (backup data) created in a backup session have the same backup ID, which is the same as the session ID of the backup session.

Mirrored objects and objects created in an object copy session have the same backup ID as objects created in the original backup session. Suppose the media set created in the original backup session no longer exists, but the media set created in an object copy session still exists. To restore the objects, you must specify the session ID of the original backup session (that is, the backup ID) and not the session ID of the object copy session.

The omnir syntax does not let you specify from which object copy to restore if several copies of the same object exist. This is only possible using the Data Protector GUI by setting the media allocation priority list.

If this option is not specified, backup data with the latest backup ID is used, regardless of the `-endlogs` or the `-time` option selection.

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 MaxDB restore options on page 202](#). See also to the `omnir man` page.

Example

To restore an instance named “inst1” (together with configuration), backed up on an SAP MaxDB Server named “srv1.company.com” from the last backup session and then perform a recovery until the end of logs, execute the following command:

```
omnir -sapdb -barhost srv1.company.com -instance inst1 -recover -endlogs
```

On how to find information about backup objects to restore from, see [Finding information for restore on page 202](#).

Restoring using SAP MaxDB utilities

Using this integration, it is also possible to run an integrated Data Protector restore of an SAP MaxDB Server from SAP MaxDB utilities.

To perform a restore to an existing SAP MaxDB Server instance, see [SAP MaxDB restore and recovery below](#).

To migrate an SAP MaxDB instance, see [SAP MaxDB migration on page 202](#).

On how to find information about backup objects to restore from, see [Finding information for restore on page 202](#).

SAP MaxDB restore and recovery

Follow the procedure on the next few pages to restore and recover a database using SAP MaxDB utilities from existing Data Protector SAP MaxDB backup session(s). In the procedure, the following conventions are used:

inst_name is the name of the instance to be restored

name_of_backup_spec is the name of the Data Protector backup specification used at backup.

username , *password* is the connection string for the SAP MaxDB database user created or identified as described in [Configuring SAP MaxDB users on page 178](#).

Location is the location of the *bsi_env* file

media_group_name is the name of the SAP MaxDB media group

medium_name is the name of the SAP MaxDB medium

pipe_name is the name of the SAP MaxDB pipe

medium_type is the type of the SAP MaxDB medium

SessionID is the Data Protector session ID of the session to be restored

Restore

1. Skip this step if the *bsi_env* file is already present and configured on the SAP MaxDB Server.

On the SAP MaxDB Server create the *bsi_env* file in a directory of your choice. It must contain the following lines:

Windows systems:

```
BACKINT Data_Protector_home\bin\sapdb_backint
INPUT Data_Protector_program_data\tmp\inst_name.bsi_in
OUTPUT Data_Protector_program_data\tmp\inst_name.bsi_out
ERROROUTPUT Data_Protector_program_data\tmp\inst_name.bsi_err
PARAMETERFILE SAPDB_PARAMETER
```

```
TIMEOUT_SUCCESS 900  
TIMEOUT_FAILURE 30
```

UNIX systems:

```
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 SAPDB_PARAMETER  
TIMEOUT_SUCCESS 900  
TIMEOUT_FAILURE 30
```

The SAPDB_PARAMETER file should contain the following lines:

```
OB2BARLIST=name_of_backup_spec  
OB2APPNAME=inst_name  
OB2BARHOSTNAME=FQDN
```

2. Login to the SAP MaxDB database manager as the SAP MaxDB database user created or identified as described in [Configuring SAP MaxDB users on page 178](#). On the SAP MaxDB Server, execute the following command to login:

```
dbmcli -d inst_name -u username,password
```

3. In the SAP MaxDB database manager, switch the database to the Admin mode by executing the following command:

```
db_admin
```

4. Skip this step if the location of the bsi_env file is already registered on the SAP MaxDB Server.

Register the location of the bsi_env file as follows:

Windows systems:

```
dbm_configset -raw BSI_ENV Location\inst_name.bsi_env
```

UNIX systems:

```
dbm_configset -raw BSI_ENV Location/inst_name.bsi_env
```

5. Skip this step if the SAP MaxDB media and pipes to be used with Data Protector are already existing on the SAP MaxDB Server.

Note that to restore a Data Protector SAP MaxDB backup session, the number of SAP MaxDB media and pipes required equals the parallelism value used during the backup session.

Create SAP MaxDB media in an SAP MaxDB media group. Execute the following command for every medium to be created, depending on the SAP MaxDB version:

- For SAP MaxDB version 7.6:

```
medium_put media_group_name/medium_namepipe_namemedia_typebackup_type [size  
[block_size [overwrite [autoloader [os_command [tool_type]]]]]]
```

- For other SAP MaxDB versions:

```
medium_put media_group_name/medium_namepipe_namemedia_typebackup_type
```

backup_type can be one of the following:

- DATA for full backup
- PAGES for differential (diff) backup
- LOG for transactional (trans) backup

tool_type must be the following:

- "BACK" for backup with Backint for SAP MaxDB

Important: When creating SAP MaxDB media and pipes for the purpose of a Data Protector backup and restore, the media group name must begin with the "BACK" string. The commands below create two media and two pipes (parallelism = 2) in a media group:

Windows systems, SAP MaxDB version 7.6:

```
medium_put BACKDP-Data[2]/1 \  
\\.\Pipe\inst_name.BACKDP_Data[2].1 PIPE DATA 0 8 \  
NO NO \\\" "BACK"
```

```
medium_put BACKDP-Data[2]/2 \  
\\.\Pipe\inst_name.BACKDP_Data[2].2 PIPE DATA 0 8 \  
NO NO \\\" "BACK"
```

UNIX systems, SAP MaxDB version 7.6:

```
medium_put BACKDP-Data[2]/1 \  
/var/opt/omni/tmp/inst_name.BACKDP_Data[2].1 PIPE \  
DATA 0 8 NO NO \\\" "BACK"
```

```
medium_put BACKDP-Data[2]/2 \  
/var/opt/omni/tmp/inst_name.BACKDP_Data[2].2 PIPE \  
DATA 0 8 NO NO \\\" "BACK"
```

Windows systems, other SAP MaxDB versions:

```
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 systems, other SAP MaxDB versions:

```
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 MaxDB utility session by executing the following command:

```
util_connect
```

7. Start the restore from a Data Protector backup session by executing the following command:

```
recover_start media_group_namebackup_type EBID "inst_nameSessionID:1 pipe_  
name1,inst_nameSessionID:2 pipe_name2[, ...]"
```

Windows systems:

```
recover_start BACKDP-Data[2] DATA EBID "inst_nameSessionID:1 \\.\Pipe\inst_  
name.BACKDP-Data[2].1,TEST SessionID:2 \\.\Pipe\inst_name.BACKDP-Data[2].2"
```

UNIX systems:

```
recover_start BACKDP-Data[2] DATA EBID "inst_nameSessionID:1  
/var/opt/omni/tmp/inst_name.BACKDP-Data[2].1,inst_nameSessionID:2  
/var/opt/omni/tmp/inst_name.BACKDP-Data[2].2"
```

Repeat this step for every session in the required chain of backup sessions.

- a. To recover the database until the specified point in time, execute the `recover_start` command with the UNTIL clause:

```
recover_start BACKDP-Archive LOG EBID "inst_name SessionID:1 pipe_name1,inst_  
name SessionID:2 pipe_name2[, ...]" UNTIL yyyyymmdd hhmmss
```

where the `yyyyymmdd` and `hhmmss` parameters specify which redo log should be applied last.

8. Use one of the commands `recover_start` and `recover_replace`, based on the exit code from the previous execution of `recover_start` or `recover_replace`.

For more information, see the SAP MaxDB documentation.

9. If the command `recover_start` or `recover_replace` from the previous step returned an exit code of `-8020`, and you have already restored all relevant data, execute the following command:

```
recover_ignore
```

For more information, see the SAP MaxDB documentation.

SAP MaxDB migration

When performing an SAP MaxDB migration, some additional tasks must first be done in order to prepare the SAP MaxDB Server or instance. These tasks are described in [Before you begin on page 194](#).

Follow the procedure in the section [SAP MaxDB restore and recovery on page 198](#) to migrate the SAP MaxDB database using SAP MaxDB utilities from existing Data Protector SAP MaxDB backup session (s). When following the mentioned procedure, *before* executing the `recover_start` command, delete the existing redo logs on the SAP MaxDB Server by executing the following command in the SAP MaxDB database manager:

```
util_execute clear log
```

Finding information for restore

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

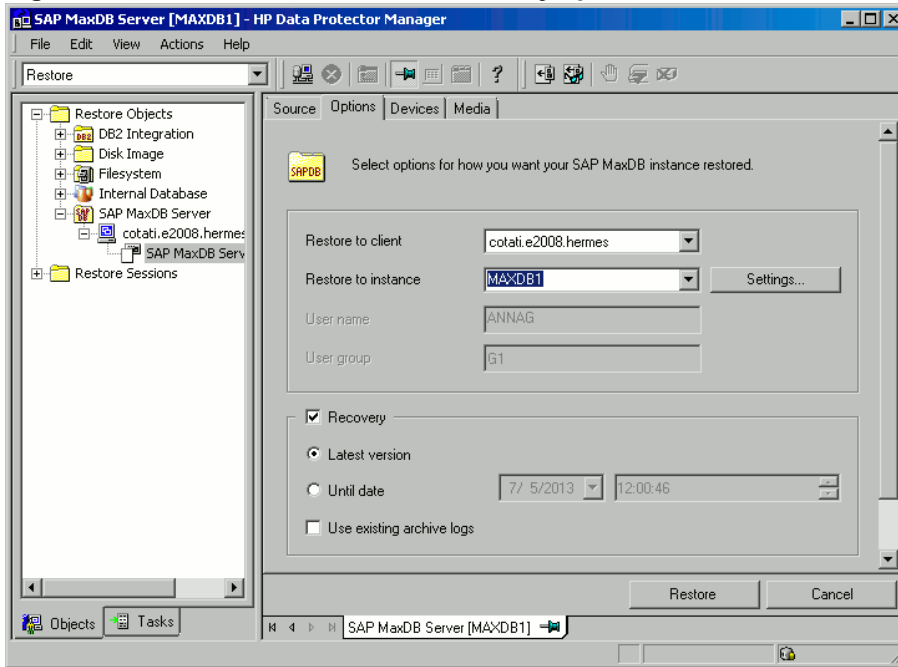
Execute the following Data Protector commands:

- `omnidb -sapdb`
to get a list of SAP MaxDB objects.
- `omnidb -sapdb object_name`
to get details on a specific object, including the SessionID.

SAP MaxDB restore options

[SAP MaxDB restore and recovery options on the next page](#) shows the SAP MaxDB GUI restore and recovery options.

Figure 52: SAP MaxDB restore and recovery options



The following are SAP MaxDB specific backup options:

Migration Options

To restore selected SAP MaxDB object to the same SAP MaxDB Server and instance, leave the migration options as they are. Use the migration options only in case of SAP MaxDB migration (when restoring to some other SAP MaxDB Server or to some other instance than those that were backed up).

The following are descriptions of the migration options. First the GUI option is given, followed by a slash (/), CLI equivalent, and then description.

<p>Restore to client / -destination <i>ClientName</i></p>	<p>When using the GUI, in the drop-down list, select an SAP MaxDB Server to which you want to restore the database.</p> <p>When using the CLI, specify the <code>-destination</code> option and the name of the SAP MaxDB Server as the <i>ClientName</i> argument.</p> <p>The selected SAP MaxDB Server must be a part of the Data Protector cell and must have the Data Protector SAP MaxDB Integration software component installed.</p>
--	---

<p>Restore to instance / - <code>newinstance</code> <i>DestinationInstanceName</i></p>	<p>When using the GUI, you can either:</p> <ul style="list-style-type: none"> • Select an instance in the Restore to instance drop-down list. The drop-down list shows only the instances that are already configured for use with this integration. See Configuring SAP MaxDB instances on page 178 for information on how to configure an SAP MaxDB Server for use with this integration. • Enter the name of an existing instance, not yet configured for use with this integration. In this case, click on the Settings button to configure the specified instance. <p>When using the CLI, the instance specified as the <i>DestinationInstanceName</i> argument to the <code>-newinstance</code> option must already be configured for use with this integration. See Configuring SAP MaxDB instances on page 178 for information on how to configure an SAP MaxDB Server for use with this integration.</p>
<p>User name and User group / N/A</p>	<p>On UNIX, you can change the user name and the group name for the OS system user, under whose account the SAP MaxDB application is running on the SAP MaxDB Server (for example, the <code>sapdb</code> user in the <code>sapsys</code> group). By default, the user that started the Data Protector GUI is set for this option.</p> <p>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.</p>
<p>Settings / N/A</p>	<p>Click this button if the instance you are restoring to is not yet configured for use with this integration. See Configuring SAP MaxDB instances on page 178 for information on parameters that must be entered.</p> <p>When using the CLI, this option is not available. To configure the instance, use the <code>util_sapdb</code> utility as described in Configuring SAP MaxDB instances on page 178.</p>

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 MaxDB Server and the restored volumes can occur. When performing recovery (when the database is switched to the `Online` mode), SAP MaxDB always checks whether such a gap exists, regardless of the point in time selected for recovery. If such a gap exists, the recovery is not performed and the database remains in the `Admin` mode, unless the existing redo logs are manually deleted before starting the restore.

The following are descriptions of the recovery options. First the GUI option is given, followed by a slash (/), CLI equivalent, and then description.

<p>Recovery / -recover</p>	<p>When this option is selected, the database is recovered after the restore (it is switched to <code>Online</code> 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).</p> <div style="background-color: #e0e0e0; padding: 10px; margin: 10px 0;"> <p>Important: When using this option, make sure that the backup ID selected in the Properties for Data dialog box (when using GUI) or specified by the <code>-session</code> option (when using CLI) will restore enough data for the integration to apply the redo logs until the latest version or until the specified date and time. For information on how to access the Properties for Data dialog box, see In the Source page, select objects for restore. on page 194. For information on the <code>-session</code> option, see Restoring using the Data Protector CLI on page 197.</p> </div> <p>When this option is not selected, all other recovery options are disabled and the following happens after the restore:</p> <ul style="list-style-type: none"> • If archive logs are not restored (if restore from a full backup session is performed), the database remains in the <code>Admin</code> mode after the restore. • If archive logs are restored, the database is, if the restored archive logs allow it, switched to the <code>Online</code> mode. If the database, however, cannot be switched to the <code>Online</code> mode (because the restored archive logs do not allow it), it remains in the <code>Admin</code> mode.
<p>Latest version / -endlogs</p>	<p>Select this option to recover the database until the last log.</p> <p>When using the CLI, this is the default option.</p>
<p>Until date / -time: YYYY-MM-DD.hh.mm.ss</p>	<p>When using the GUI, select this option to recover the database until the point you select in the Until date drop-down menu.</p> <p>When using the CLI, specify the <code>-time:</code> option if you want to recover the database until the point specified by the <code>YYYY-MM-DD.hh.mm.ss</code> argument.</p> <div style="background-color: #e0e0e0; padding: 10px; margin: 10px 0;"> <p>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.</p> </div>

<p>Use existing archive logs / -from_ disk</p>	<p>Select this option to copy the existing archive logs on the SAP MaxDB Server to SAP MaxDB Server redo logs.</p> <p>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 MaxDB Server (if full or diff backup session is restored).</p> <p>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 MaxDB Server are applied to redo logs.</p> <div data-bbox="430 604 1404 745" style="background-color: #f0f0f0; padding: 5px;"><p>Note: The Use existing archive logs option is disabled in case of SAP MaxDB migration, thus allowing only for the restore of redo logs from the backed up archive logs on backup media (if trans backup session is restored).</p></div>
---	--

Restoring using another device

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

For information on how to select another device for a restore using the Data Protector GUI, see the *HP Data Protector Help* index: “restore, selecting devices for”.

Monitoring sessions

You can monitor currently running sessions in the Data Protector GUI. When you run an interactive backup or a restore session, a monitor window shows you the progress of the session. Closing the GUI does not affect the session.

You can also monitor sessions from any Data Protector client with the User Interface component installed, using the Monitor context.

On how to monitor a session, see the *HP Data Protector Help* index: “viewing currently running sessions”.

Troubleshooting

This section lists problems you might encounter when using the Data Protector SAP MaxDB integration.

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

Before you begin

- Ensure that the latest official Data Protector patches are installed. See the *HP Data Protector Help* index: “patches” on how to verify this.

- See the *HP Data Protector Product Announcements, Software Notes, and References* for general Data Protector limitations, as well as recognized issues and workarounds.
- See <http://support.openview.hp.com/selfsolve/manuals> for an up-to-date list of supported versions, platforms, and other information.

Problems

Problem

Data Protector reports the following error during backup or restore:

```
[Critical] From: OB2BAR_SAPDBBAR@machine.company.com "INSTANCE"  
Time: 02/06/04 18:17:18 Error: SAPDB responded with:  
-24920,ERR_BACKUPOP: backup operation was unsuccessful  
The database was unable to fulfill a request  
(-2025, Invalid number of backup devices).
```

Action

Increase the value of the SAP MaxDB MAXBACKUPDEVS parameter to a value that is greater than or equal to the value of the Data Protector Parallelism option, or reduce the value of the Data Protector Parallelism option.

Problem

An SAP MaxDB instance cannot be started after restore

Action

Using the SAP MaxDB `db_restartinfo` command, check if the instance can be restarted.

- If the instance cannot be restarted, most probably the existing log volumes do not contain enough data to restart the instance from data volumes. The required differential or transactional backups might not have been restored.
- If the instance can be restarted, check the SAP MaxDB instance kernel error file for errors.

If there was insufficient space for SAP MaxDB logs at some point of time, logs might have been corrupted: delete the logs (using the `dbmcli util_execute clear log` command) or contact SAP MaxDB or Data Protector support.

Problem

A restore session for restoring data from an object copy gets blocked

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:

- a. 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.
- b. 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.
- c. Set the highest media location priority for the newly created copies.

Problem

SAP MaxDB database is in the histlost state

When the log volumes are initialized (for example, after a restore or recovery), the database is in the histlost state. Consequently, restore or recovery in this state cannot be performed successfully. For example, recovery fails with the following error:

```
Error: SAPDB responded with: -24920,ERR_BACKUPOP: backup operation was unsuccessful
The database was unable to fulfill a request (-9407, System error: unexpected error).
```

Action

After a restore or recovery, perform a full backup to start a new backup history. In case this error occurs during recovery or restore, execute the `db_execute clear log` command and repeat the restore or recovery.

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 MaxDB `x_server`. For details, see the SAP MaxDB documentation.

Problem

Data Protector reports the following error:

```
Error: SAPDB responded with:
-24988,ERR_SQL: sql error
1,database not running.
```

Action

Start the SAP MaxDB instance. For details, see the SAP MaxDB documentation.

Problem

Data Protector reports the following error:

```
Error: SAPDB responded with:
-24988,ERR_SQL: sql error1,utility session is already in use.
```


Action

Some other user is connected to the SAP MaxDB instance and is performing administrative tasks (utility session). Such SAP MaxDB tasks are of the "Utility" type and can be displayed using the `dbmcli show task` command. Finish these tasks.

Problem

Data Protector reports the following error:

```
Error: SAPDB responded with:  
-24950,ERR_USRFAIL: user authorization failed.
```

Action

Reconfigure the SAP MaxDB instance as described in the section [Configuring SAP MaxDB instances on page 178](#).

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 executing the following command:

```
util_cmd -putopt SAPDB SAPDB_instance TimeoutSuccess 1000 -sublist Environment
```

For more information, see the `util_cmd` man page.

You can also set the `TimeoutSuccess` environment variable using the Data Protector GUI. Select the backup specification in the Scoping Pane, then right-click the SAP MaxDB instance object in the Results Pane under the **Source** tab and select the **Set Environment Variables** from the pop-up menu.

SAP MaxDB cluster-related troubleshooting

In a cluster environment, the environment variable `OB2BARHOSTNAME` must be defined as the name of the cluster virtual system before performing some procedures executed from the command line (on the client). The `OB2BARHOSTNAME` variable is set as follows:

Windows systems

```
set OB2BARHOSTNAME=VirtualSystemName
```

UNIX systems

```
export OB2BARHOSTNAME=VirtualSystemName
```

Chapter 4: Data Protector SAP HANA Appliance integration

Introduction

This chapter provides information specific to the Data Protector SAP HANA Appliance integration (**SAP HANA integration**). It describes the concepts you need to understand and the procedures you need to follow to successfully back up and restore SAP HANA Appliance objects (**SAP HANA objects**). For general SAP HANA Appliance information, see the SAP HANA documentation. For general Data Protector information and information related to the Data Protector filesystem backup functionality that must be used to back up and restore SAP HANA configuration files, see the *HP Data Protector Help*.

Data Protector integrates with SAP HANA Appliance to offer online backup of SAP HANA Appliance instances (**SAP HANA instances**). You can back up the following entities using the Data Protector SAP HANA integration:

- SAP HANA database
- SAP HANA redo logs

On the other hand, you can back up SAP HANA configuration files only using the Data Protector filesystem backup functionality.

During the backup, the database must be online, meaning it can be actively used.

Data Protector offers backups of the following types:

Table 17: Backup types

Full	This is backup of a complete SAP HANA database. A corresponding session backs up entire database contents. In SAP HANA terminology, this backup type is referred to as <i>data area backup</i> .
Trans	This is backup of SAP HANA redo logs. A corresponding session backs up the redo logs. In SAP HANA terminology, this backup type is referred to as <i>log area backup</i> .

For information on the available target locations to restore SAP HANA objects to (original as opposed to different Data Protector SAP HANA client, original as opposed to different SAP HANA instance on the same client), see the *SAP HANA Administration Guide* and other items of the SAP HANA documentation set.

For more information on the SAP HANA objects deleted from the Data Protector IDB, see the notification saved in the log file at the following location:

```
/var/opt/omni/log/saphana.log
```

As part of the restore process, you can also recover the SAP HANA database either to the state at a specific point in time or to the latest state recreated using the available redo logs.

Important: Current implementation of the Data Protector SAP HANA integration does not include a conventional Data Protector application integration agent. It only enables you to use Data Protector to:

- Define which data needs to be backed up, how it should be backed up, and where the backup images should be stored—by configuring SAP HANA backup specifications in the Data Protector GUI
- Monitor progress of the Data Protector SAP HANA backup and restore sessions
- List SAP HANA objects backed up by Data Protector and their corresponding sessions

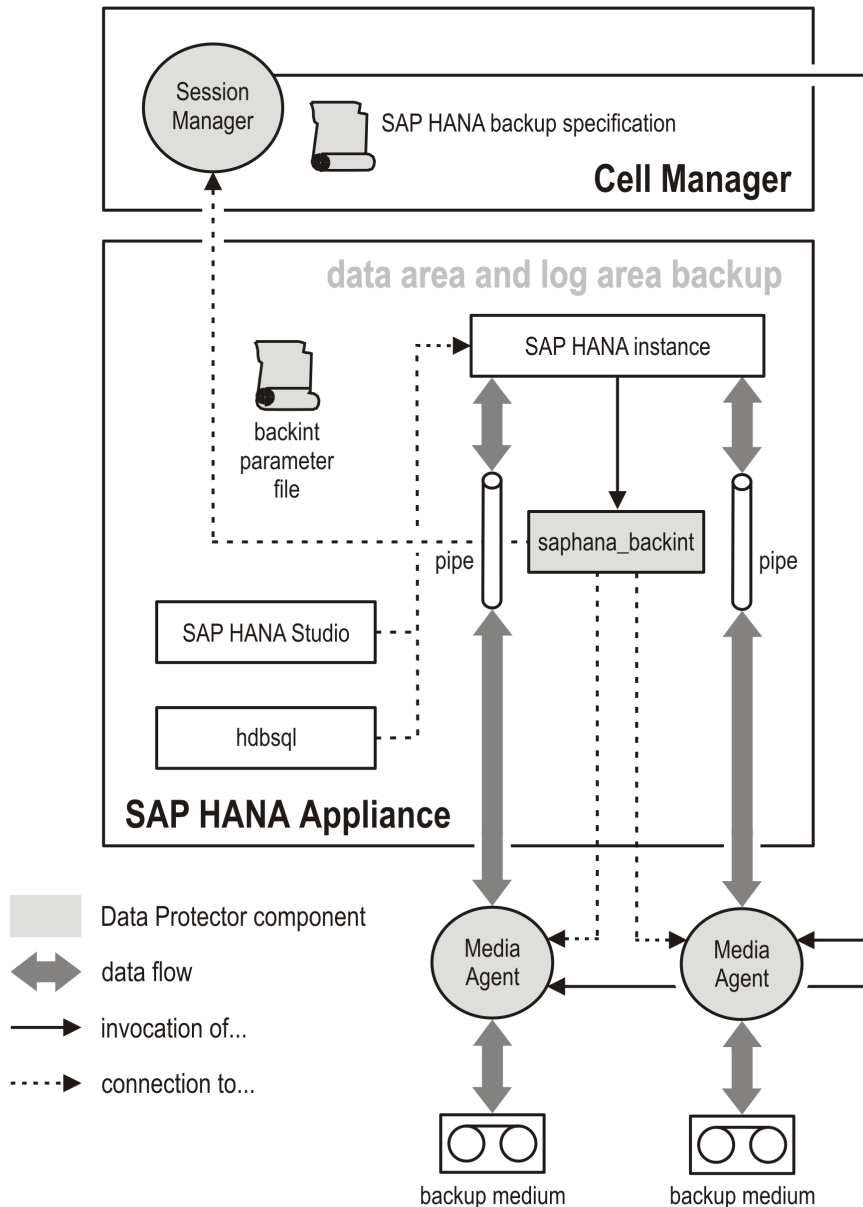
Timing aspects of the backup and restore processes (for example, when the backup sessions are started and when and where the data is restored to) as well as the database recovery options are governed by the SAP HANA Appliance itself. You can only define and control them from the SAP HANA Appliance GUI (SAP HANA Studio) or CLI (the SAP HANA `hdbsql` command-line utility).

Integration concepts

Data Protector integrates with SAP HANA Appliance using the `SAP HANA Integration` component, which provides the Data Protector implementation of the SAP HANA `backint` agent.

[SAP MaxDB integration architecture on page 175](#) shows the architecture of the Data Protector SAP HANA integration.

Figure 53: SAP HANA integration architecture



The Data Protector SAP HANA backint agent consists of the `saphana_backint` binary file that acts as an interface between Data Protector and the SAP HANA Appliance backup and restore functionality. Both SAP HANA database and SAP HANA redo logs are backed up and restored as streams.

Integration limitations

- Data Protector backup specifications for transactional backup cannot be configured.
 A transactional backup is always implicitly performed first whenever a full backup session is invoked from an SAP HANA Appliance user interface.
- Data Protector does not provide means of starting SAP HANA backup sessions (interactive,

scheduled) and SAP HANA restore sessions from its own user interfaces. Previews of SAP HANA backup and restore sessions are not available either.

Configuring the integration

Before configuring the Data Protector SAP HANA integration, check the prerequisites first.

Prerequisites

- Your SAP HANA Appliance must be correctly installed and configured.
 - For a list of versions, platforms, and devices supported by Data Protector, and other related information, see the latest support matrices at <http://support.openview.hp.com/selfsolve/manuals>.
 - For information on installing, configuring, and using SAP HANA Appliance, see the SAP HANA documentation.
- Data Protector must be correctly installed. For instructions on how to install Data Protector to integrate it with your SAP HANA Appliance, see the *HP Data Protector Installation and Licensing Guide*.
- Data Protector backup devices and Data Protector backup media must be configured according to your SAP HANA backup policy needs.
- The SAP HANA operating system user account that you choose for running SAP HANA backup and restore sessions must have a corresponding Data Protector user configured in the Data Protector `admin` or `operator` user group.

Configuration steps

In scope of configuration of the SAP HANA instances for Data Protector backup, you need to:

- Link the Data Protector implementation of the `backint` agent with your SAP HANA Appliance.
- Adjust SAP HANA backup settings for the redo logs as to use the `backint` interface instead of plain files for backup storage.
- Optionally, set the global (CS) `MaxObjectsPerBackupSession` environment variable to limit the number of objects connected to the same backup session. Valid values are 1 through 223, with the default being 223. If the `MaxObjectsPerBackupSession` value reaches 223, backup continues in a new session.

Linking Data Protector backint agent with SAP HANA Appliance

Perform the following:

1. Log on to the SAP HANA system and open a Terminal window.
2. This step depends on your SAP HANA Appliance environment.

Single-server SAP HANA environment:

Change current directory to *InstallationPath/InstanceID/global/hdb* where *InstallationPath* is the path to the installation directory (for example, */usr/sap*) and *InstanceID* is the identifier of your SAP HANA instance (for example, H95).

```
cd InstallationPath/InstanceID/global/hdb
```

Distributed SAP HANA environment:

Change current directory to *SharedInstallationPath/InstanceID/global/hdb* where *SharedInstallationPath* is the path to the directory with shared installation data (for example, */hana/shared*) and *InstanceID* is the identifier of your SAP HANA instance (for example, H95). Execute the command:

```
cd SharedInstallationPath/InstanceID/global/hdb
```

3. Locate the subdirectory */opt*. If it does not exist, create it by executing the command:

```
mkdir opt
```

4. Create a symbolic link to the Data Protector implementation of the backint agent. Execute the commands:

```
cd opt
```

```
ln -s /opt/omni/lbin/saphana_backint hdbbackint
```

You can verify that you properly linked the backint agent in SAP HANA Studio by inspecting the Backint Agent option in the Backup > Configuration > Backint Settings context.

Adjusting SAP HANA backup settings for redo logs

Perform the following:

1. Launch SAP HANA Studio.
2. In the scoping pane, expand any SAP HANA instance and select **Backup**.
3. In the main pane, select the **Configuration** tab, and expand **Backint Settings**.
4. Under Log Backup Settings, select **Backint** for the Destination Type option. Make sure the option **Enable Automatic Log Backup** is selected.
5. Switch to the **Administration Console** view.

6. In the scoping pane, double-click any SAP HANA instance.
7. In the main pane, select the **Configuration** tab, and expand the **global.ini** tree.
8. Expand the **persistence** subtree, and make sure the value of the `log_mode` parameter is `normal`.
9. From the **File** menu, select **Save**.

Backup

The Data Protector SAP HANA integration provides online backup of the SAP HANA database and the SAP HANA redo logs. The backup scope is defined when starting a backup session from an SAP HANA Appliance user interface. See [Backup types on page 210](#).

Note: For information about data that needs to be backed up in scope of SAP HANA configuration files backup and their location, see the SAP HANA documentation.

Creating backup specifications

To create an SAP HANA backup specification, proceed as follows:

1. Launch the Data Protector GUI.
2. In the Context List, click **Backup**.
3. In the Scoping Pane, expand **Backup Specifications**, right-click **SAP HANA Server**, and select **Add Backup**.
4. In the **Create New Backup** dialog box, select the **Blank SAP HANA Backup** template. Click **OK**.
5. From the **Client** drop-down list, select :
 - The SAP HANA system
(if your SAP HANA Appliance is a single-server environment)
 - The SAP HANA master host with the configured name server role `MASTER1`
(if your SAP HANA Appliance is a distributed environment; for instructions on how to determine the SAP HANA system with this role, see the SAP HANA documentation)

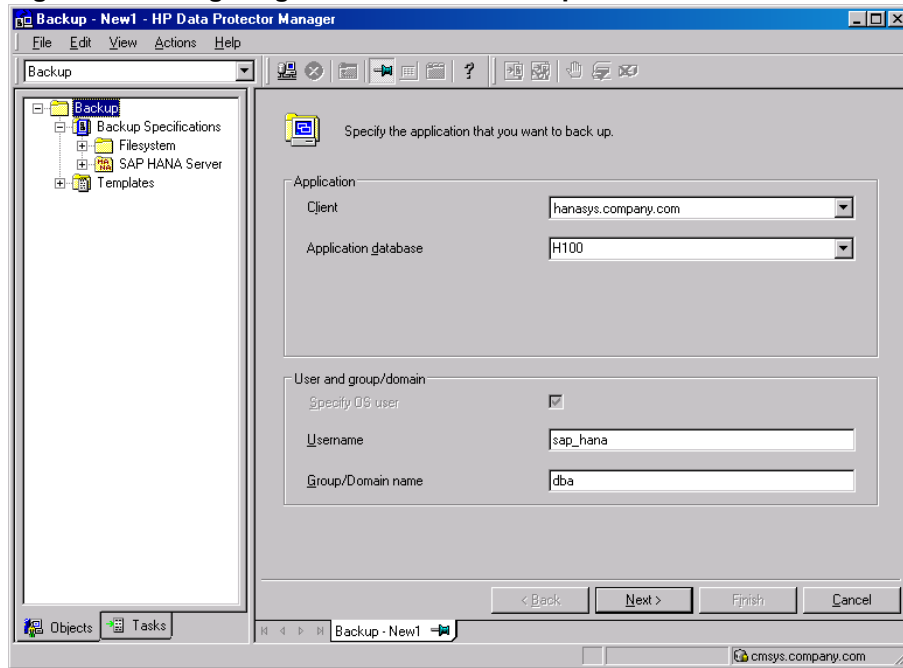
In the **Application database** text box, enter the name of the SAP HANA instance that you want to back up.

In the **Username** and **Group/Domain name** text boxes, enter the appropriate SAP HANA user account. For more information, press **F1**.

Click **Next**.

Leave the only available object selected, and click **Next**.

Figure 54: Configuring an SAP HANA backup



- 6.
7. Select the backup devices to use for the backup session. If needed, change their ordering and adjust load balancing and object mirroring.

To specify device options, right-click the device and select **Properties** to open the Device Properties dialog box, where you can modify device-specific options. When done, click **OK**.

Tip: For the best restore performance of the SAP HANA redo logs, it is recommended to set the device Concurrency option to 1 in a backup specification used for log backup.

Click **Next**.

8. Under Backup Specification Options, enter optional description of you backup specification and click **Advanced** to adjust general backup specification options.

Under Common application options, click **Advanced** to adjust integration backup options. For more information, press **F1**.

Click **Next**.

9. Click **Save as**.
10. In the Save Backup As dialog box, enter a name for you backup specification in the Name text box. From the Group drop-down list, select a backup specification group.
11. Click **OK** to save the backup specification.

Tip: You can keep Data Protector SAP HANA backup specifications better organized by saving them in the same dedicated Data Protector backup specification group that you create yourself.

Creating Data Protector parameter files for SAP HANA

To enable successful execution of SAP HANA backup and restore sessions, you need to further define a few Data Protector parameters, and make them available to your SAP HANA Appliance. Proceed as follows:

1. Start a plain text editor and compose the following lines with parameter configuration:

```
OB2BARLIST='BackupSpecificationName';
```

```
OB2BARHOSTNAME='SystemName';
```

BackupSpecificationName is the name of the SAP HANA backup specification and *SystemName* is the string that you selected from the Client drop-down list when configuring the backup specification.

An example of a properly configured parameters is as follows:

```
OB2BARLIST='sap_hana_backup';
```

```
OB2BARHOSTNAME='hanasys.company.com';
```

2. Save the file to the following directory, depending on your SAP HANA Appliance environment:

Single-server SAP HANA environment:

InstallationPath /InstanceID/global/hdb/opt/hdbconfig where *InstanceID* is the identifier of your SAP HANA instance (for example, H95)

An example of a full pathname of the parameter file in a single-server SAP HANA environment is */usr/sap/H95/global/hdb/opt/hdbconfig/data_protector.par*.

Distributed SAP HANA environment:

SharedInstallationPath /InstanceID/global/hdb/opt/hdbconfig where *SharedInstallationPath* is the path to the directory with shared installation data (for example, */hana/shared*) and *InstanceID* is the identifier of your SAP HANA instance (for example, H95)

3. Launch SAP HANA Studio.
4. In the scoping pane, expand any SAP HANA instance and select **Backup**.
5. In the main pane, select the **Configuration** tab, and expand **Backint Settings**.

The */opt/omni/sbin/saphana_backint* value of the Backint Agent option denotes a properly linked Data Protector backint agent.

6. Under Backint Settings > Data Backup, enter the full pathname of the parameter file in the **Backint Parameter File** text box.
7. From the **File** menu, select **Save**.

Modifying backup specifications

For information on modifying your SAP HANA backup specifications, see the *HP Data Protector Help* index: “backup specifications, modifying”.

Starting backup sessions

With current implementation of the Data Protector SAP HANA integration, you can only use one of the SAP HANA user interfaces to start Data Protector SAP HANA backup sessions: SAP HANA Studio or the `hdbsql` command-line utility.

Prerequisites

- Database of the SAP HANA instance selected for backup must be online.

For instructions on how to start a session for backing up an SAP HANA database and how to configure and adjust settings for the periodic backup of SAP HANA redo logs, see the *SAP HANA Administration Guide* and other items of the SAP HANA documentation set.

Important: Instruct SAP HANA Appliance to use the backint interface instead of plain files for backup storage. For this purpose, select **Backint** for the Destination Type option on the Specify Backup Settings page of the backup wizard.

Restore

The Data Protector SAP HANA integration provides restore of the SAP HANA database and the SAP HANA redo logs. The restore scope is defined when starting a recovery session from an SAP HANA Appliance user interface.

Note: For information about data that needs to be restored in scope of SAP HANA configuration files restore and where, see the SAP HANA documentation.

Finding information for restore

To find the information needed for restoring your SAP HANA data, use the Data Protector `omnidb - saphana [Client:Set]` command. The command retrieves a list of SAP HANA objects backed up by Data Protector and their corresponding sessions. For details, see the `omnidb` reference page in the *HP Data Protector Command Line Interface Reference* or the `omnidb` man page.

Starting restore and recovery sessions

With current implementation of the Data Protector SAP HANA integration, you can only use one of the SAP HANA user interfaces to start Data Protector SAP HANA restore or recovery sessions: SAP HANA Studio or the `hdbsql` command-line utility.

Prerequisites

- Database of the target SAP HANA instance must be offline.

SAP HANA Studio can automatically take the database offline when a restore session is initiated. In contrast, before using the `hdbsql` command-line utility for restoring or recovering an SAP HANA database, you need to take the database offline manually.

For instructions on how to start a Data Protector session for restoring an SAP HANA database, how to perform database recovery, and how to specify other restore options, see the *SAP HANA Administration Guide* and other items of the SAP HANA documentation set.

Important: While the recovery wizard of SAP HANA Studio considers the `backint` interface as the restore source for the SAP HANA database by default, you need to instruct SAP HANA Appliance to use the `backint` interface also as the source for restoring the redo logs. For this purpose, select the **Third-Party Backup Tool (Backint)** option on the Other Settings page of the recovery wizard.

Monitoring sessions

You can monitor currently running Data Protector sessions in the Monitor context of the Data Protector GUI. The Results Pane shows you the progress of the selected session. Closing the GUI does not affect the session.

For monitoring the sessions running in a Data Protector cell or in all cells that are part of a Data Protector Manager-of-Managers environment, you can use any system with the Data Protector User Interface or Manager-of-Managers User Interface component installed, respectively.

For more information, see the *HP Data Protector Help* index: “viewing currently running sessions”.

Troubleshooting

This section lists problems you might encounter when using the Data Protector SAP HANA integration.

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

Before you begin

- Ensure that the latest Data Protector patches are installed. For information on how to verify this, see the *HP Data Protector Help* index: “patches”.

- For general Data Protector limitations as well as recognized issues and workarounds, see the *HP Data Protector Product Announcements, Software Notes, and References*.
- For an up-to-date list of supported versions, platforms, and other information, see <http://support.openview.hp.com/selfsolve/manuals>.

Problems

Problem

Backup session fails in a distributed SAP HANA environment

When a session for backing up a database or the corresponding redo logs is invoked from SAP HANA Studio in a distributed SAP HANA environment, the session fails and reports the following error:

```
Backup of system InstanceID failed.
```

```
Backup could not be completed, Backint cannot execute  
PathnameOfSymbolicLinkToBackintAgent, No such file or directory (2)
```

A possible cause is the missing Data Protector SAP HANA Integration component on some SAP HANA systems.

Action

Install the Data Protector SAP HANA Integration component on all SAP HANA systems where it is missing, and restart the session.

Problem

Recovery session fails with a log error

SAP HANA reports the following error during recovery:

```
RECOVER DATA finished with error:  
recovery could not be completed, volume <volumeID>, reached log position <LSN>,  
A log segment in the log backup for service '<serviceName>' at '<backupID>' could not  
be found, volume id '<volumeID>', state 'end of log reached'
```

This failure may be caused due to the broken restore chain.

Action

You can perform one of the following actions:

- Use the Clear Log feature. It ensures that only log entries in the log backups are recovered. For more information, see the SAP HANA Administration Guide.
- Restore to a specific log position, which is identified by the LSN in the aforementioned error message.

Problem

SAP HANA backup or restore fails in SAP HANA Studio

When SAP HANA backup or restore is started from SAP HANA Studio, the backup or restore fails immediately. As the Data Protector session does not start, you cannot identify the problem.

Action

You can perform the following actions:

- Verify that you have properly linked the backint agent in SAP HANA Studio by inspecting the Backint Agent option in the Backup > Configuration > Backint Settings context.
- Verify that the user performing the SAP HANA backup or restore is added to the Data Protector user list.

For more information, see [Configuring the integration on page 213](#).

Glossary

A

access rights

See user rights.

ACSL (StorageTek specific term)

The Automated Cartridge System Library Server (ACSL) 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.

AES 256-bit encryption

The Data Protector software encryption, based on the AES-CTR (Advanced Encryption Standard in Counter Mode) encryption algorithm that uses random keys of 256-bit length. The same key is used for both encryption and decryption. With AES 256-bit encryption, data is encrypted before it is transferred over a network and before it is written to media.

AML (ADIC/GRAU specific term)

Automated Mixed-Media library.

AMU (ADIC/GRAU specific term)

Archive Management Unit.

application agent

A component needed on a client system 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.

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.

archived log files (Data Protector specific term)

Files that keep track of changes made to the Data Protector Internal Database (IDB). They are used for online or offline IDB restore and recovery where the IDB needs to be recreated either in its latest possible state, beyond the time of the most recent IDB backup session, or in a state between the times of two consecutive IDB backup sessions.

archived redo log (Oracle specific term)

Also called offline redo log. If the Oracle database operates in the ARCHIVELOG mode, as each online redo log is filled, it is copied to an archived log destination. This copy is the archived redo log. The presence or absence of an archived redo log is determined by the mode the database is using: ARCHIVELOG - The filled online redo log files are archived before they are reused. The database can be recovered if an instance or a disk fails. A "hot" backup can be performed only when the database is running in this mode. NOARCHIVELOG - The filled online redo log files are not archived. See also online redo log.

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 the directory <Data_Protector_program_data>\Config\Server\dr\asr (Windows systems) or /etc/opt/omni/server/dr/asr (UNIX systems), as well as on the backup medium. After a disaster occurs, the ASR archive file is extracted to diskettes which you need to perform ASR.

audit logs

Data files to which auditing information is stored.

audit report

User-readable output of auditing information created from data stored in audit log files.

auditing information

Data about every backup session that was performed over an extended, user-defined period for the whole Data Protector cell.

autochanger

See library.

autoloader

See library.

Automatic Storage Management (ASM) (Oracle specific term)

A filesystem and volume manager integrated into Oracle which manages Oracle database files. It eliminates complexity associated with data and disk management and optimizes performance by providing striping and mirroring capabilities.

auxiliary disk

A bootable disk that has a minimal operating system with networking and

Data Protector Disk Agent installed. It can be carried around and used to boot the target system in Phase 1 of Disk Delivery Disaster Recovery of UNIX clients.

B

BACKINT (SAP R/3 specific term)

A Data Protector interface program that lets the SAP R/3 backup programs communicate with the Data Protector software via calls to an open interface. For backup and restore, SAP R/3 programs issue commands through the Data Protector backint interface.

backup API (Oracle specific term)

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 reading and writing of data to media and also to allow creation, searching, and removing of backup files.

backup chain

See restore chain.

backup device

A device configured for use with Data Protector that 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/application 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 (for filesystem objects - the access point in a directory structure on the client where the backup object is located (drive on Windows systems and mount point on UNIX systems), for integration objects - backup stream identification, indicating the backed up database/application items), description (for filesystem objects - uniquely defines objects with identical client name and mount point, for integration objects - displays the integration type), and type (for filesystem objects - filesystem type, for integration objects - "Bar").

backup owner

Each backup object in the IDB has an owner. The default owner of a backup is the user who starts the backup session.

backup session

A process that creates a copy of data on storage media. The activities are specified in a backup specification or an interactive session. The result of a backup session is a set of media, which was written to, also called the backup or media set. See also backup specification, incremental backup, and full backup.

backup set

A complete set of integration objects associated with a backup.

backup set (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; and days and time that you want backups to be performed. The objects are entire disks/volumes or parts of them such as files, directories, or even the Windows Registry (for example). File selection lists, such as include-lists and exclude-lists, can be specified. All clients configured in one backup specification are backed up at the same time in one backup session using the same backup type (full or incremental).

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 of your 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 Continuances are processes that allow customers to access and manage instant copies of EMC Symmetrix standard devices.

BC Process (EMC Symmetrix specific term)

A protected storage environment solution, which uses specially configured EMC Symmetrix devices as mirrors or Business Continuance Volumes to protect data on EMC Symmetrix standard devices. See also BCV.

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 Data Protector 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 containing files required for the initial step of the boot process. Microsoft terminology defines the boot volume/disk/partition as a volume/disk/partition containing the operating system files.

BRARCHIVE (SAP R/3 specific term)

An SAP R/3 backup tool that allows you to archive redo log files. BRARCHIVE also saves all the logs and profiles of the archiving process. See also BRBACKUP and BRRESTORE.

BRBACKUP (SAP R/3 specific term)

An SAP R/3 backup tool that allows an online or offline backup of the control file, of individual data files, or of all tablespaces and, if necessary, of the online redo log files. See also BRARCHIVE and BRRESTORE.

BRRESTORE (SAP R/3 specific term)

An SAP R/3 tool that can be used to restore files of the following type: Database data files, control files, and online redo log files saved with BRBACKUP; Redo log files archived with BRARCHIVE; Non-database files saved with BRBACKUP. You can specify files, tablespaces, complete backups, log sequence numbers of redo log files, or the session ID of the backup. See also BRARCHIVE and BRBACKUP.

BSM

The Data Protector Backup Session Manager, which controls the backup session. This process always runs on the Cell Manager system.

C**CAP (StorageTek specific term)**

The Cartridge Access Port built into the door panel of a library. The purpose is to enter or eject media.

Catalog Database (CDB)

A part of the Data Protector Internal Database (IDB) that contains information about backup, restore, object copy, object consolidation, object verification, and media management sessions. This part of the IDB is always local to the cell. It is stored in the embedded database. See also MMDB.

catalog protection

Defines how long information about backed up data (such as filenames and file attributes) is kept in the IDB. See also data protection.

CDB

See Catalog Database (CDB).

CDF file (UNIX systems 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.

Certificate Server

A Windows certificate server can be installed and configured to provide certificates for clients. It provides customizable services for issuing and managing certificates for the enterprise. These services issue, revoke, and manage certificates employed in public key-based cryptography technologies.

Change Journal (Windows specific term)

A Windows filesystem feature that logs a record of each change as it occurs to the files and directories on a local NTFS volume.

Change Log Provider

A module that can be queried to determine which objects on a filesystem have been created, modified, or deleted.

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' or type 'sbt_tape'. If the specified channel is of type 'sbt_tape' and Oracle is integrated with Data Protector, the server process will attempt to read backups from or write data files to Data Protector.

circular logging (Microsoft Exchange Server and Lotus Domino Server specific term)

Circular logging is a Microsoft Exchange Server database and Lotus Domino Server database mode, in which the transaction log file contents are periodically overwritten after the corresponding data is committed to the database. Circular logging reduces disk storage space requirements.

client backup

A backup of all volumes (filesystems) mounted on a Data Protector client. What is actually backed up depends on how you select objects in a backup specification. If you select the check box next to the client system name, a single backup object of the Client System type is created. As a result, at the time of the backup, Data Protector first detects all volumes that are mounted on the selected client and then backs them up. On Windows clients, CONFIGURATION is also backed up. If you individually select all volumes that are mounted on the client system, a separate backup object of the Filesystem type is created for each volume. As a result, at the time of the backup, only the selected volumes are backed up. Volumes that have been potentially mounted on the client after the

backup specification was created are not backed up.

client or client system

Any system configured with any Data Protector functionality and configured in a cell.

cluster continuous replication (Microsoft Exchange Server specific term)

Cluster continuous replication (CCR) is a high availability solution that uses cluster management and failover options to create and maintain an exact copy (CCR copy) of a storage group. A storage group is replicated to a separate server. CCR removes any single point of failure in your Exchange back-end servers. You can perform backups using VSS on your passive Exchange Server node where a CCR copy is located and thus reducing the load on the active node. A CCR copy is used for disaster recovery since you can switch to the CCR copy in a few seconds. A replicated storage group is represented as a new instance of Exchange writer called Exchange Replication Service and can be backed up (using VSS) like an ordinary storage group. See also Exchange Replication Service and local continuous replication.

cluster-aware application

It is an application that supports the cluster Application Programming Interface. Each cluster-aware application declares its own critical resources (disk volumes (on Microsoft Cluster Server), volume groups (on HP Serviceguard), application services, IP names and addresses ...).

CMD script for Informix Server (Informix Server specific term)

A Windows CMD script that is created in INFORMIXDIR when an Informix Server database is configured. The CMD script

is a set of system commands that export environment variables for Informix Server.

CMMDB

The Data Protector Centralized Media Management Database (CMMDB) is the result of merging MMDBs from several cells in the MoM environment. It allows you to share high-end devices and media across multiple cells in a MoM environment. One cell can control the robotics, serving the devices that are connected to systems in other cells. The CMMDB must reside on the Manager-of-Managers. A reliable network connection between the MoM cell and the other Data Protector cells is highly recommended. See also MoM.

COM+ Class Registration Database (Windows specific term)

The COM+ Class Registration Database and the Windows Registry store COM+ application attributes, class attributes, and computer-level attributes. This guaranties consistency among these attributes and provides common operation on top of these attributes.

command device (HP P9000 XP Disk Array Family specific term)

A dedicated volume in the disk array which acts as the interface between a management application and the disk array's storage system. It cannot be used for data storage and only accepts requests for operations that are then executed by the disk array.

command-line interface (CLI)

A set commands that you can use in shell scripts to perform Data Protector configuration, backup, restore, and management tasks.

concurrency

See Disk Agent concurrency.

container (HP P6000 EVA Disk Array Family specific term)

Space on a disk array, which is pre-allocated for later use as a standard snapshot, vsnap, or snapclone.

control file (Oracle and SAP R/3 specific term)

A data file that contains entries specifying the physical structure of the database. It provides database consistency information used for recovery.

copy set (HP P6000 EVA Disk Array Family specific term)

A pair that consists of the source volumes on a local P6000 EVA and their replica on a remote P6000 EVA. See also source volume, replica, and HP Continuous Access + Business Copy (CA+BC) P6000 EVA.

CRS

The Data Protector Cell Request Server process (service), which runs on the Cell Manager, starts and controls the backup and restore sessions. The service is started as soon as Data Protector is installed on the Cell Manager. The CRS runs under the account root on UNIX systems. On Windows systems it runs under the account of the user, specified at installation time.

CSM

The Data Protector Copy and Consolidation Session Manager process controls the object copy and object consolidation sessions and runs on the Cell Manager system.

D**data file (Oracle and SAP R/3 specific term)**

A physical file created by Oracle that contains data structures such as tables and indexes. A data file can only belong to one Oracle database.

data protection

Defines how long the backed up data on media remains protected, that is, Data Protector will not overwrite it. When the protection expires, Data Protector will be able to reuse the media in one of the next backup sessions. See also catalog protection.

Data 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 replication (DR) group (HP P6000 EVA Disk Array Family specific term)

A logical grouping of HP P6000 EVA Disk Array Family virtual disks. It can contain up to eight copy sets provided they have common characteristics and share a common HP CA P6000 EVA log. See also copy set.

data stream

Sequence of data transferred over the communication channel.

Data_Protector_home

A reference to the directory containing Data Protector program files (on Windows Vista, Windows 7, Windows Server 2008, and Windows Server 2012) or the directory containing Data Protector program files and data files (on other Windows operating systems). Its default path is %ProgramFiles%\OmniBack, but the path can be changed in the Data Protector Setup Wizard at installation time. See also Data_Protector_program_data.

Data_Protector_program_data

A reference to the directory containing Data Protector data files on Windows Vista, Windows 7, Windows Server 2008, and Windows Server 2012. Its default path is %ProgramData%\OmniBack, but the path can be changed in the Data Protector Setup Wizard at installation time. See also Data_Protector_home.

database library

A Data Protector set of routines that enables data transfer between Data Protector and a server of an online database integration, for example, Oracle Server.

database parallelism

More than one database is backed up at a time if the number of available devices allows you to perform backups in parallel.

database server

A computer with a large database stored on it, such as the SAP R/3 or Microsoft SQL database. A server has a database that can be accessed by clients.

Dbobject (Informix Server specific term)

An Informix Server physical database object. It can be a blob space, db space, or logical log file.

DC directory

A directory that contains DC binary files, one for each configured Data Protector backup medium. DC directories constitute the Detail Catalog Binary Files part of the Data Protector Internal Database. See also Detail Catalog Binary Files (DCBF) and Internal Database (IDB).

DCBF

See Detail Catalog Binary Files (DCBF).

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.

Detail Catalog Binary Files (DCBF)

A part of the Data Protector Internal Database that stores names, versions, and metadata of the backed up items. It consists of DC directories with DC binary files. See also DC directory and Internal Database (IDB).

device

See backup device.

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. -Get configuration, status, and performance statistics by device group. -Issue control operations that apply to all devices in the device group.

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 writing to the tape, and so on. If data is written to the tape slower than it is delivered to the device then the device is streaming. Streaming significantly improves the use of space and the performance of the device.

DHCP server

A system running the Dynamic Host Configuration Protocol (DHCP) providing dynamic IP address assignment and network configuration for DHCP clients.

differential backup

An incremental backup that backs up changes made since the last full backup. To perform this type of backup, specify the Incr1 backup type. See also incremental backup.

differential backup (Microsoft SQL Server 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.

directory junction (Windows specific term)

Directory junctions use the reparse point concept of Windows. An NTFS 5 directory junction allows you to redirect a directory/file request to another location.

disaster recovery

A process to restore a client's main system disk to a state close to the time when a (full) backup was performed.

disaster recovery operating system

See DR OS.

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. During an object verification session the Disk Agent receives data from the Media Agent and performs the verification process, but no data is written to disk.

Disk Agent concurrency

The number of Disk Agents that are allowed to send data to one Media Agent concurrently.

disk group (Veritas Volume Manager specific term)

The basic unit of data storage in VxVM system. A disk group can consist of one or more physical volumes. There can be more than one disk group on the system.

disk image (rawdisk) backup

A high-speed backup where Data Protector backs up files as bitmap images. A disk image (rawdisk) backup does not track the files and directory

structure stored on the disk, but stores a disk image structure on byte level. You can perform a disk image backup of either specific disk sections or a complete disk.

disk quota

A concept to manage disk space consumption for all or a subset of users on a computer system. This concept is used by several operating system platforms.

disk staging

The process of backing up data in several phases to improve the performance of backups and restores, reduce costs of storing the backed up data, and increase the data availability and accessibility for restore. The backup stages consist of backing up data to one media type first (for example disk) and later copying it to a different media type (for example tape).

distributed file media format

A media format, available with the file library, which supports a space efficient type of synthetic backup called virtual full backup. Using this format is a prerequisite for virtual full backup. See also virtual full backup.

Distributed File System (DFS)

A service that connects file shares into a single namespace. The file shares can reside on the same or on different computers. DFS provides client access to the resources in a location-transparent manner.

DMZ

The Demilitarized Zone (DMZ) is a network inserted as a "neutral zone" between a company's private network (intranet) and the outside public network (Internet). It prevents outside users from getting direct access to company servers in the intranet.

DNS server

In the DNS client-server model, this is the server containing information about a portion of the DNS database that makes computer names available to client resolvers querying for name resolution across the Internet.

DR image

Data required for temporary disaster recovery operating system (DR OS) installation and configuration.

DR OS

An operating system environment in which disaster recovery runs. It provides Data Protector with a basic runtime environment (disk, network, tape, and filesystem access). It has to be installed on disk or loaded into memory and configured before the Data Protector disaster recovery can be performed. DR OS can be either temporary or active. A temporary DR OS is used exclusively as a host environment for some other operating system restore along with the target operating system configuration data. It is deleted after the target system is restored to the original system configuration. An active DR OS not only hosts the Data Protector disaster recovery process but can also be a part of the restored system because it replaces its own configuration data with the original configuration data.

drive

A physical unit that receives data from a computer system and can write it onto a magnetic medium. 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.

drive-based encryption

The Data Protector drive-based encryption uses the encryption functionality of the drive. While performing the backup, the drive encrypts both the data and the metadata that is written to the medium.

E

EMC Symmetrix Agent

A Data Protector software module that prepares the EMC Symmetrix environment for backup and restore operations.

emergency boot file (Informix Server specific term)

The Informix Server configuration file `ixbar.<server_id>` that resides in the directory `<INFORMIXDIR>/etc` (on Windows systems) or `<INFORMIXDIR>\etc` (on UNIX systems). `<INFORMIXDIR>` is the Informix Server home directory and `<server_id>` is the value of the `SERVERNUM` configuration parameter. Each line of the emergency boot file corresponds to one backup object.

encrypted control communication

Data Protector secure communication between the clients in the Data Protector cell is based on a Secure Socket Layer (SSL) that uses export grade SSLv3 algorithms to encrypt control communication. Control communication in a Data Protector cell is all communication between Data Protector processes, except the data transfer from Disk Agent (and Integrations) to Media Agent, and the other way round.

encryption key

A 256-bit randomly generated number used by the Data Protector encryption algorithm to encode information during backups for which AES 256-bit software encryption or drive-based encryption has been specified. The same key is used for subsequent decryption of the information. Encryption keys for a Data Protector cell are stored in a central keystore on the Cell Manager.

encryption KeyID-StoreID

Combined identifier used by the Data Protector Key Management Server to identify and administer encryption keys used by Data Protector. KeyID identifies the key within the keystore. StoreID identifies the keystore on the Cell Manager. If Data Protector has been upgraded from an earlier version with encryption functionality, there may be several StoreIDs used on the same Cell Manager.

enhanced incremental backup

Conventional incremental backup backs up files that have changed since a previous backup, but has certain limitations in detection of changes. Unlike conventional incremental backup, enhanced incremental backup reliably detects and backs up also renamed and moved files, as well as files with changes in attributes.

enterprise backup environment

Several cells can be grouped together and managed from a central cell. The enterprise backup environment includes all clients located in several Data Protector cells which are managed and administered from a central cell using the Manager-of-Managers concept. See also MoM.

Event Log (Data Protector Event Log)

A central repository of all Data Protector related notifications. By default, all notifications are sent to the Event Log. The Event Log is accessible only to Data Protector users in the 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.

Event Logs (Windows specific term)

Files in which Windows logs all events, such as the starting or stopping of services and the logging on and off of users. Data Protector can back up Windows Event Logs as part of the Windows configuration backup.

Exchange Replication Service (Microsoft Exchange Server specific term)

The Microsoft Exchange Server service that represents storage groups that were replicated using either local continuous replication (LCR) or cluster continuous replication (CCR) technology. See also cluster continuous replication and local continuous replication.

exchanger

See library.

exporting media

A process that removes all data about backup sessions, such as systems, objects, and filenames, which reside on the media from the IDB. Information about the media and their relation to a pool is also removed from the IDB. The data on the media remains unchanged. See also importing media.

**Extensible Storage Engine (ESE)
(Microsoft Exchange Server specific
term)**

A database technology used as a storage system for information exchange in Microsoft Exchange Server.

F**failover**

Transferring of the most important cluster data, called group (on Windows systems) or package (on UNIX systems) from one cluster node to another. A failover can occur mostly because of software or hardware failures or maintenance on the primary node.

**failover (HP P6000 EVA Disk Array
Family specific term)**

An operation that reverses the roles of source and destination in HP Continuous Access + Business Copy (CA+BC) P6000 EVA configurations. See also HP Continuous Access + Business Copy (CA+BC) P6000 EVA.

FC bridge

See Fibre Channel bridge.

Fibre Channel

An ANSI standard for high-speed computer interconnection. Using either optical or copper cables, it allows the high speed bi-directional transmission of large data files and can be deployed between sites kilometers apart. Fibre Channel connects nodes using three physical topologies: point-to-point, loop, and switched.

Fibre Channel bridge

A Fibre Channel bridge or multiplexer provides the ability to migrate existing parallel SCSI devices, like RAID arrays, solid state disks (SSD), and tape libraries

to a Fibre Channel environment. On one side of the bridge or multiplexer there is a Fibre Channel interface while on the other side there are parallel SCSI ports. The bridge or multiplexer enables SCSI packets to be moved between the Fibre Channel and parallel SCSI devices.

file depot

A file containing the data from a backup to a file library device.

file jukebox device

A device residing on disk consisting of multiple slots used to store file media.

file library device

A device which resides on a disk emulating a library with several media, hence containing multiple files, referred to as file depots.

File Replication Service (FRS)

A Windows service that replicates the domain controller store logon scripts and group policies. FRS also enables replication of Distributed File System (DFS) shares between systems and allows any server to perform replication activity.

file tree walk

The process of traversing a filesystem to determine which objects have been created, modified, or deleted.

file version

The same file can be backed up multiple times in case of full backups and incremental backups (if the file changed). If the log level ALL is selected for backup, Data Protector retains one entry in the IDB for the filename itself and one for each version (date/time) of the file.

filesystem

The organization of files on a hard disk. A filesystem is backed up so that the file attributes and the file contents are stored on the backup media.

first-level mirror (HP P9000 XP Disk Array Family specific term)

A mirror of an internal disk (LDEV) of a disk array of the HP P9000 XP Disk Array Family which can be further mirrored itself, producing second-level mirrors. For Data Protector zero downtime backup and instant recovery purposes, only first-level mirrors can be used. See also primary volume and mirror unit (MU) number.

flash recovery area (Oracle specific term)

A directory, filesystem, or Automatic Storage Management (ASM) disk group managed by Oracle that serves as a centralized storage area for files related to backup, restore, and database recovery (recovery files). See also recovery files.

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). Use the Force operation option to reformat Data Protector media with non-protected data. Data Protector media with protected data are not formatted until the protection expires or the media are unprotected/recycled.

free pool

An auxiliary source of media for use by media pools when they run out of media. The media pools must be configured to use free pools.

full backup

A backup in which all selected objects are backed up, whether or not they have been recently modified. See also backup types.

full database backup

A backup of all data in a database, not only the data that has been changed after the last (full or incremental) backup of the database. A full database backup does not depend on any other backup.

full mailbox backup

A full mailbox backup is a backup of the entire mailbox content.

full ZDB

A ZDB-to-tape or ZDB-to-disk+tape session in which all selected objects are streamed to tape, even if there are no changes from the previous backup. See also incremental ZDB.

G

global options

A set of options that define behavior of the entire Data Protector cell. The options are stored in a plain text file on the Cell Manager.

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 graphical user interface provided by Data Protector for easy access to all configuration, administration, and operation tasks. It is available for Microsoft Windows operating systems.

H**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 that resides on the Cell Manager at the following location: <Data_Protector_program_data>\Config\Server\holidays (Windows systems) and /etc/opt/omni/server/Holidays (UNIX systems).

hosting system

A working Data Protector client used for Disk Delivery Disaster Recovery with a Data Protector Disk Agent installed.

HP Business Copy (BC) P6000 EVA (HP P6000 EVA Disk Array Family specific term)

A local replication software solution that enables creation of point-in-time copies (replicas) of the source volumes using the snapshot and clone capabilities of the P6000 EVA firmware. See also replica, source volume, snapshot, and HP Continuous Access + Business Copy (CA+BC) P6000 EVA.

HP Business Copy (BC) P9000 XP (HP P9000 XP Disk Array Family specific term)

An HP P4000 SAN Solutions configuration that enables creation and maintenance of internal copies of LDEVs for various purposes, such as data duplication and backup. The copies (secondary volumes or S-VOLs) can be separated from the primary volumes (P-VOLs) and connected to a different system. For Data Protector zero downtime backup purposes, P-VOLs should be available to the application system, and one of the S-VOL sets should be available to the backup system. See also LDEV, HP Continuous Access (CA) P9000 XP, Main Control Unit (MCU), application system, and backup system.

HP Command View (CV) EVA (HP P6000 EVA Disk Array Family specific term)

The user interface that enables you to configure, manage, and monitor your P6000 EVA storage system. It is used to perform various storage management tasks, for example, creating virtual disk families, managing storage system hardware, and creating snapshots, snapclones, or mirrorclones of virtual disks. The HP Command View EVA software runs on the HP Storage Management Appliance, and is accessed

by a Web browser. See also HP P6000 / HP 3PAR SMI-S Agent and HP SMI-S P6000 EVA Array provider.

HP Continuous Access (CA) P9000 XP (HP P9000 XP Disk Array Family specific term)

An HP P9000 XP Disk Array Family configuration that enables creation and maintenance of remote copies of LDEVs for purposes such as data duplication, backup, and disaster recovery. HP CA P9000 XP operations involve main (primary) disk array units and remote (secondary) disk array units. The main disk array units are connected to the application system and contain primary volumes (P-VOLs), which store original data. The remote disk array units are connected to the backup system and contain secondary volumes (S-VOLs). See also HP BC P9000 XP (HP P9000 XP Disk Array Family specific term), Main Control Unit (MCU), and LDEV.

HP Continuous Access + Business Copy (CA+BC) P6000 EVA (HP P6000 EVA Disk Array Family specific term)

An HP P6000 EVA Disk Array Family configuration that enables creation and maintenance of copies (replicas) of the source volumes on a remote P6000 EVA, and later use of these copies as the source for local replication on this remote array. See also HP BC P6000 EVA, replica, and source volume.

HP P6000 / HP 3PAR SMI-S Agent

A Data Protector software module that executes all tasks required for the HP P6000 EVA Disk Array Family integration. With the P6000 / 3PAR SMI-S Agent, the control over the array is established through an appropriate SMI-S provider, which directs communication between incoming requests and the storage system's native interface. See

also HP Command View (CV) EVA and HP SMI-S P6000 EVA Array provider.

HP P9000 XP Agent

A Data Protector software component that executes all tasks needed by the Data Protector HP P9000 XP Disk Array Family integration. It communicates with the P9000 XP Array storage system via the RAID Manager Library.

HP SMI-S P6000 EVA Array provider

An interface used for controlling HP P6000 EVA Disk Array Family. SMI-S P6000 EVA Array provider runs as a separate service on the HP Storage Management Appliance system and acts as a gateway between incoming requests and HP Command View EVA. With the Data Protector HP P6000 EVA Disk Array Family integration, SMI-S P6000 EVA Array provider accepts standardized requests from the HP P6000 / HP 3PAR SMI-S Agent, communicates with HP Command View EVA for information or method invocation, and returns standardized responses. See also HP P6000 / HP 3PAR SMI-S Agent and HP Command View (CV) EVA.

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

See Internal Database (IDB).

IDB recovery file

A file that maintains information about completed IDB backup sessions and the

backup media and backup devices used in them. If available, it significantly simplifies and speeds up offline recovery of the Internal Database in case of a Cell Manager disaster. Its filename is obdrindex.dat.

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 (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 backup

A backup that selects only files that have changed since a previous backup. Several levels of incremental backup are available, which enables detailed control of restore chain length. See also backup types.

incremental backup (Microsoft Exchange Server specific term)

A backup of the Microsoft Exchange Server data that has changed since the last full or incremental backup. With the incremental backup, only the transaction log files are backed up. See also backup types.

incremental mailbox backup

An incremental mailbox backup backs up all the changes made to the mailbox after the last backup of any type.

incremental restore (EMC Symmetrix specific term)

A BCV or SRDF control operation. In BCV control operations, an incremental restore reassigns a BCV device as the next available mirror of the standard device in the pair. However, the standard devices are updated with only the data that was written to the BCV device during the time of the original pair split, and the data that was written to the standard device during the split is overwritten with data from the BCV mirror. In SRDF control operations, an incremental restore reassigns a target (R2) device as the next available mirror of the source (R1) device in the pair. However, the source (R1) devices are updated with only the data that was written to the target (R2) device during the time of the original pair split, and the data that was written to the source (R1) device during the split is overwritten with data from the target (R2) mirror.

incremental ZDB

A filesystem ZDB-to-tape or ZDB-to-disk+tape session in which only changes from the last protected full or incremental backup are streamed to tape. See also full ZDB.

Incremental1 Mailbox Backup

An incremental1 mailbox backup backs up all the changes made to the mailbox after the last full backup.

Inet

A process that runs on each UNIX system or service that runs on each Windows system in the Data Protector cell. It is responsible for communication

between systems in the cell and for starting other processes needed for backup and restore. The Inet service is started as soon as Data Protector is installed on a system. The Inet process is started by the inetd daemon.

Information Store (Microsoft Exchange Server specific term)

The Microsoft Exchange Server service that is responsible for storage management. Information Store in Microsoft Exchange Server manages two kinds of stores: mailbox stores and public folder stores. A mailbox store consists of mailboxes that belong to individual users. A public folder store contains public folders and messages that are shared among several users. See also Key Management Service and Site Replication Service.

Informix Server (Informix Server specific term)

Refers to Informix Dynamic Server.

initializing

See formatting.

Installation Server

A computer system that holds a repository of the Data Protector installation packages for a specific architecture. The Installation Server is used for remote installation of Data Protector clients. In mixed environments at least two Installation Servers are needed: one for UNIX systems and one for Windows systems.

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.

integration object

A backup object of a Data Protector integration, such as Oracle or SAP MaxDB.

Internal Database (IDB)

An entity in Data Protector that keeps information regarding which data was backed up, to which media it was backed up, how and when backup, restore, and other sessions were run, which devices, libraries, and disk arrays are configured, and so on. It is implemented with an embedded database and a collection of proprietary data files which reside on the Cell Manager. See also DC directory and Detail Catalog Binary Files (DCBF).

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).

ISQL (Sybase specific term)

A Sybase utility used to perform system administration tasks on Sybase SQL Server.

J

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'.

K

Key Management Service (Microsoft Exchange Server specific term)

The Microsoft Exchange Server service that provides encryption functionality for enhanced security. See also Information Store and Site Replication Service.

keychain

A tool that eliminates the supply of a passphrase manually when decrypting the private key. It needs to be installed and configured on the Installation Server if you perform remote installation using secure shell.

keystore

All encryption keys are centrally stored in the keystore on the Cell Manager and administered by the Key Management Server (KMS).

KMS

Key Management Server (KMS) is a centralized service that runs on the Cell Manager and provides key management for the Data Protector encryption functionality. The service is started as soon as Data Protector is installed on the Cell Manager.

L

LBO (Symmetric specific term)

A 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.

LDEV (HP P9000 XP Disk Array Family specific term)

A logical partition of a physical disk of a disk array of the HP P9000 XP Disk Array Family. An LDEV is the entity that can be replicated using the split-mirror or snapshot functionality of such disk array. See also HP Business Copy (BC) P9000 XP, HP Continuous Access (CA) P9000 XP, and replica.

library

Also called autochanger, jukebox, autoloader, or exchanger. A library contains media in repository slots. Each slot holds one medium (for example, DDS/DAT). Media are moved between slots and drives by a robotic mechanism, allowing random access to media. The library can contain multiple drives.

lights-out operation or unattended operation

A backup or restore operation that takes place outside of normal business hours without an operator. This implies that no operator personnel is present to work with the backup application or service mount requests, for example.

LISTENER.ORA (Oracle specific term)

An Oracle configuration file that describes one or more Transparent Network Substrate (TNS) listeners on a server.

load balancing

By default, Data Protector automatically balances the usage of devices selected for backup, so that they are used evenly. Load balancing optimizes the device usage by balancing the number of objects written to each device. Since load balancing is done automatically during

backup time, you do not need to manage how the data is actually backed up. You just specify the devices to be used. If you do not want to use load balancing, you can select which device will be used for each object in the backup specification. Data Protector will access the devices in the specified order.

local and remote recovery

Remote recovery is performed if all Media Agent hosts specified in the SRD file are accessible. If any of them fails, the disaster recovery process fails over to the local mode. This means that the target system is searched for locally attached devices. If only one device is found, it is automatically used.

Otherwise, Data Protector prompts you to select the device, which will be used for restore.

local continuous replication (Microsoft Exchange Server specific term)

Local continuous replication (LCR) is a single-server solution that creates and maintains an exact copy (LCR copy) of a storage group. An LCR copy is located on the same server as the original storage group. When an LCR copy is created, it is kept up to date through change propagation (log replay) technology. The replication feature in LCR guarantees that logs that have not been replicated are not deleted. The implication of this behavior is that running backups in a mode that deletes logs may not actually free space if replication is sufficiently far behind in its log copying. An LCR copy is used for disaster recovery because you can switch to the LCR copy in a few seconds. If an LCR copy is used for backup and if it is located on a different disk than the original data, then the I/O load on a production database is minimal. A replicated storage group is represented as a new instance of Exchange writer called Exchange Replication Service and

can be backed up (using VSS) as a normal storage group. See also cluster continuous replication and Exchange Replication Service.

lock name

You can configure the same physical device several times with different characteristics, by using different device names. The lock name is a user specified string that is used for locking all such device configurations to prevent collision if several such devices (device names) are used concurrently. Use an identical lock name for all device definitions which use the same physical device.

log_full shell script (Informix Server UNIX systems specific term)

A script provided by ON-Bar that you can use to start backing up logical log files when Informix Server issues a logfull event alarm. The Informix Server ALARMPROGRAM configuration parameter defaults to the `<INFORMIXDIR>/etc/log_full.sh`, where `<INFORMIXDIR>` is the Informix Server home directory. If you do not want logical logs to be backed up continuously, set the ALARMPROGRAM configuration parameter to `<INFORMIXDIR>/etc/no_log.sh`.

logging level

An option that determines the amount of details on files and directories written to the IDB during backup, object copying, or object consolidation. You can always restore your data, regardless of the logging level used during backup. Data Protector provides four logging levels: Log All, Log Directories, Log Files, and No Log. The different logging level settings mainly influence the IDB growth and the convenience of browsing data for restore.

logical-log files

This applies to online database backup. Logical-log files are files in which modified data is first stored before being flushed to disk. In the event of a failure, these logical-log files are used to roll forward all transactions that have been committed as well as roll back any transactions that have not been committed.

login ID (Microsoft SQL Server specific term)

The name a user needs to log on to Microsoft SQL Server. A login ID is valid if Microsoft SQL Server has an entry for that user in the system table syslogin.

login information to the Oracle Target Database (Oracle and SAP R/3 specific term)

The format of the login information is <user_name>/<password>@<service>, where: <user_name> is the name by which a user is known to Oracle Server and to other users. Every user name is associated with a password and both have to be entered to connect to an Oracle Target Database. This user must have Oracle SYSDBA or SYSOPER rights. <password> must be the same as the password specified in the Oracle password file (orapwd), which is used for authentication of users performing database administration. <service> is the name used to identify an SQL*Net server process for the target database.

login information to the Recovery Catalog Database (Oracle specific term)

The format of the login information to the Recovery (Oracle) Catalog Database is <user_name>/<password>@<service>, where the description of the user name, password, and service name is the same as in the Oracle SQL*Net V2 login information to the Oracle target database.

In this case, <service> is the name of the service to the Recovery Catalog Database, not the Oracle target database. Note that the Oracle user specified here must be the owner of the Oracle Recovery Catalog.

Lotus C API (Lotus Domino Server specific term)

An interface for the exchange of backup and recovery information between Lotus Domino Server and a backup solution, like Data Protector.

LVM

A Logical Volume Manager is a subsystem for structuring and mapping physical disk space to logical volumes on UNIX systems. An LVM system consists of several volume groups, where each volume group has several volumes.

M**Magic Packet**

See Wake ONLAN.

mailbox (Microsoft Exchange Server specific term)

The location to which e-mail is delivered, which is set up by the administrator for each user. If a set of personal folders is designated as the e-mail delivery location, e-mail is routed from the mailbox to this location.

mailbox store (Microsoft Exchange Server specific term)

A part of the Information Store that maintains information in 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 P9000 XP Disk Array Family specific term)

An HP P9000 XP Disk Array Family unit that contains primary volumes (P-VOLs) for the HP CA P9000 XP Array or HP CA+BC P9000 XP Array configuration and acts as a master device. See also HP Business Copy (BC) P9000 XP, HP Continuous Access (CA) P9000 XP, and LDEV.

maintenance mode

An operating mode you can initiate on the Cell Manager to prevent changes in the Internal Database. It enables you to perform various maintenance tasks, such as upgrading and patching the installation.

make_net_recovery

make_net_recovery is an Ignite-UX command, which allows you to create a recovery archive over the network onto the Ignite-UX server or any other specified system. The target system can be recovered across subnets after booting either from a bootable tape created by the Ignite-UX make_boot_tape command or the system boots directly from the Ignite-UX server. Booting directly from the Ignite-UX server can be automated with the Ignite-UX bootsys command or interactively specified on the boot console.

make_tape_recovery

make_tape_recovery is a command on Ignite-UX which creates a bootable recovery (installation) tape, customized for your system and enables you unattended disaster recovery by connecting the backup device directly to the target system and booting the target system from the bootable recovery tape. The backup device has to be locally connected to the client during the creation of the archive and recovery of the client.

Manager-of-Managers

See MoM.

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.

MCU

See Main Control Unit (MCU).

Media Agent

A process that controls reading from and writing to a device, which reads from or writes to a medium (typically a tape). During a backup session, a Media Agent receives data from the Disk Agent and sends it to the device for writing it to the medium. During a restore or object verification session, a Media Agent locates data on the backup medium and sends it to the Disk Agent for processing. For a restore session, the Disk Agent then writes the data to the disk. A Media Agent also manages the robotics control of a library.

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 assert the condition of a medium.

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.

medium ID

A unique identifier assigned to a medium by Data Protector.

merging

This defines one mode to resolve file conflicts during restore. If the file to be restored already exists at the destination, the one with the more recent modification date is kept. Files not present on the disk are always restored. See also overwrite.

Microsoft Exchange Server

A "client-server" messaging and a workgroup system that offers a transparent connection to many different communication systems. It provides users with an electronic mail system, individual and group scheduling, online forms, and workflow automation tools. It provides a developer with a platform on which to build custom information-sharing and messaging-service applications.

Microsoft Management Console (MMC) (Windows specific term)

An administration model for Windows-based environments. It provides a simple, consistent, and integrated administration user interface allowing management of many applications through the same GUI, provided that the applications adhere to the MMC model.

Microsoft SQL Server

A database management system designed to meet the requirements of distributed "client-server" computing.

Microsoft Volume Shadow Copy Service (VSS)

A software service that provides a unified communication interface to coordinate backup and restore of a VSS-aware application regardless of its specific features. This service collaborates with the backup application, writers, shadow

copy providers, and the operating system kernel to implement the management of volume shadow copies and shadow copy sets. See also shadow copy, shadow copy provider, replica, and writer.

mirror (EMC Symmetrix and HP P9000 XP Disk Array Family specific term)

See target volume.

mirror rotation (HP P9000 XP Disk Array Family specific term)

See replica set rotation.

mirror unit (MU) number (HP P9000 XP Disk Array Family specific term)

A non-negative integer number that determines a secondary volume (S-VOL) of an internal disk (LDEV) located on a disk array of the HP P9000 XP Disk Array Family. See also first-level mirror.

mirrorclone (HP P6000 EVA Disk Array Family specific term)

A dynamic replica of a storage volume, which is kept updated with changes made to the original storage volume via a local replication link. Replication between the original storage volume and its mirrorclone can be suspended. For each storage volume, a single mirrorclone can be created on the disk array.

MMD

The Media Management Daemon process (service) (MMD) 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 and Catalog Database (CDB).

MoM

Several cells can be grouped together and managed from a central cell. The management system of the central cell is the Manager-of-Managers (MoM). The cells are called MoM clients. The MoM enables you to configure and manage multiple cells from a central point.

mount point

The access point in a directory structure for a disk or logical volume, for example /opt or d:. On UNIX systems, the mount points are displayed using the bdf or df command.

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.

MSM

The Data Protector Media Session Manager, which runs on the Cell Manager and controls media sessions, such as copying media.

multisnapping (HP P6000 EVA Disk Array Family specific term)

Simultaneous creation of target volumes so that the backup data is consistent not only on each individual target volume, but also across all the volumes that constitute a snapshot. See also snapshot.



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.

obdrindex.dat

See IDB recovery file.

object

See backup object.

object consolidation

The process of merging a restore chain of a backup object, consisting of a full backup and at least one incremental backup, into a new, consolidated version of this object. The process is a part of the synthetic backup procedure. The result is a synthetic full backup of the specified backup object.

object consolidation session

A process that merges a restore chain of a backup object, consisting of a full backup and at least one incremental backup, into a new, consolidated version of this object.

object copy

A copy of a specific object version that is created during an object copy session or a backup session with object mirroring.

object copy session

A process that creates an additional copy of the backed up data on a different media set. During an object copy session, the selected backed up objects are copied from the source to the target media.

object copying

The process of copying selected object versions to a specific media set. You can

select object versions from one or several backup sessions to be copied.

object ID (Windows specific term)

The object IDs (OIDs) enable access to NTFS 5 files no matter where in the system the files reside. Data Protector treats the OIDs as alternate streams of the files.

object mirror

A copy of a backup object created using object mirroring. Object mirrors are often referred to as object copies.

object mirroring

The process of writing the same data to several media sets during a backup session. Data Protector enables you to mirror all or some backup objects to one or more media sets.

object verification

The process of verifying the data integrity of backup objects, from the Data Protector point of view, and the ability of Data Protector to deliver them to the required destination. The process can be used to provide a level of confidence in the ability to restore object versions created by backup, object copy, or object consolidation sessions.

object verification session

A process that verifies the data integrity of specified backup objects or object versions and the ability of selected Data Protector network components to deliver them to a specified host. Object verification sessions can be run interactively, or as specified in automated post-backup, or scheduled specifications.

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. Cell Manager can only be recovered offline.

offline redo log

See archived redo log.

ON-Bar (Informix Server specific term)

A backup and restore system for Informix Server. ON-Bar enables you to create a copy of your Informix Server data and later restore the data. The ON-Bar backup and restore system involves the following components: the onbar command, Data Protector as the backup solution, the XBSA interface, and ON-Bar catalog tables, which are used to back up dobjects and track instances of dobjects through multiple backups.

ONCONFIG (Informix Server specific term)

An environment variable that specifies the name of the active ONCONFIG configuration file. If the ONCONFIG environment variable is not present, Informix Server uses the configuration values from the onconfig file located in

the directory <INFORMIXDIR>\etc (on Windows systems) or <INFORMIXDIR>/etc/ (on UNIX systems).

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 recovery

A type of the Internal Database recovery that you can use when the Cell Manager is accessible. In this case, the Cell Manager runs sessions, the sessions are logged into the IDB, and you can monitor the session progress using the GUI.

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.

OpenSSH

A set of network connectivity tools used to access remote machines securely, by using a variety of authentication and encryption methods. It needs to be installed and configured on the Installation Server and the client if you perform remote installation using secure shell.

Oracle Data Guard (Oracle specific term)

Oracle Data Guard is Oracle's primary disaster recovery solution. Oracle Data Guard is able to maintain up to nine standby databases, each of which is a real-time copy of the production (primary) database, to protect against corruptions, data failures, human errors, and disasters. If a failure occurs on the production (primary) database, then a failover to one of the standby databases which becomes the new primary database is possible. In addition, planned downtime for maintenance can be reduced because the production processing can be moved from the current primary database to a standby database and back quickly.

Oracle instance (Oracle specific term)

Each installation of an Oracle database on one or more systems. One computer system can have several instances of a database running.

ORACLE_SID (Oracle specific term)

A unique name for an Oracle Server instance. To switch among Oracle Servers, specify the desired <ORACLE_SID>. The <ORACLE_SID> is included in the CONNECT DATA parts of the connect descriptor in a TNSNAMES.ORA file and in the definition of the TNS listener in the LISTENER.ORA file.

original system

The system configuration backed up by Data Protector before a computer disaster hits the system.

overwrite

An option that defines one mode to resolve file conflicts during restore. All files are restored from a backup even if they are older than existing files. See also merging.

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. When copying or consolidating backup objects, the owner of the resulting objects is the user who started the original backup session.

P

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 the

directory <Data_Protector_program_data>\Config\Server\dr\p1s (Windows systems) or /etc/opt/omni/dr/p1s (UNIX systems) with the filename recovery.p1s.

package (HP 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 P9000 XP Disk Array Family specific term)

The status of a disk pair (secondary volume and its corresponding primary volume) of a disk array of the HP P9000 XP Disk Array Family. Depending on the circumstances, the paired disks can be in various states. The following states are particularly important for the operation of the Data Protector HP P9000 XP Agent:

PAIR - The secondary volume is prepared for zero downtime backup. If it is a mirror, it is completely synchronized, and if it is a volume to be used for snapshot storage, it is empty.

SUSPENDED - The link between the disks is suspended. However, the pair relationship is still maintained, and the secondary disk can be prepared for zero downtime backup again at a later time.

COPY - The disk pair is currently busy and making a transition into the PAIR state. If the secondary volume is a mirror, it is re-synchronizing with the primary volume, and if it is a volume to be used for snapshot storage, its contents are getting cleared.

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.

phase 0 of disaster recovery

Preparation for disaster recovery - the prerequisite condition for a successful disaster recovery.

phase 1 of disaster recovery

Installation and configuration of DR OS, establishing previous storage structure.

phase 2 of disaster recovery

Restoration of operating system (with all the configuration information that defines the environment) and Data Protector.

phase 3 of disaster recovery

Restoration of user and application data.

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 systems and as shell scripts on UNIX systems. 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 systems and as shell scripts on UNIX systems.

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 systems and as shell scripts on UNIX systems. See also post-exec.

primary volume (P-VOL) (HP P9000 XP Disk Array Family specific term)

An internal disk (LDEV) of a disk array of the HP P9000 XP Disk Array Family for which a secondary volume (S-VOL), either its mirror or a volume to be used for its snapshot storage, exists. In the HP CA P9000 XP and HP CA+BC P9000 XP configurations, primary volumes are located in the Main Control Unit (MCU). See also secondary volume (S-VOL) and Main Control Unit (MCU).

protection

See data protection and catalog protection.

public folder store (Microsoft Exchange Server specific term)

The part of the Information Store that maintains information in public folders. A

public folder store consists of a binary rich-text .edb file and a streaming native internet content .stm file.

public/private backed up data

When configuring a backup, you can select whether the backed up data will be: public, that is visible (and accessible for restore) to all Data Protector users private, that is, visible (and accessible for restore) only to the owner of the backup and administrators

R

RAID

Redundant Array of Independent Disks.

RAID Manager Library (HP P9000 XP Disk Array Family specific term)

A software library that is used for accessing the configuration, status, and performance measurement data of a P9000 XP Array storage system, and for invoking operations on the disk array. It translates function calls into sequences of low-level SCSI commands. See also HP P9000 XP Agent.

RAID Manager P9000 XP (HP P9000 XP Disk Array Family specific term)

A software application that provides a command-line interface to disk arrays of the HP P9000 XP Disk Array Family. It offers an extensive set of commands for reporting and controlling the status of a P9000 XP Array storage system, and for performing various operations on the disk array.

rawdisk backup

See disk image backup.

RCU

See Remote Control Unit (RCU).

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.

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, and stored scripts.

Recovery Catalog Database (Oracle specific term)

An Oracle database that contains a recovery catalog schema. You should not store the recovery catalog in your target database.

recovery files (Oracle specific term)

Recovery files are Oracle specific files that reside in the flash recovery area: the current control file, online redo logs, archived redo logs, flashback logs, control file autobackups, datafile copies, and backup pieces. See also flash recovery area.

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.

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.

recycle or unprotect

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 P9000 XP Disk Array Family specific term)

An HP P9000 XP Disk Array Family unit that acts as a slave device to the Main Control Unit (MCU) in the HP CA P9000 XP or HP CA+BC P9000 XP configuration. In bidirectional configurations, the RCU can also act as an MCU.

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 or software with which it is created, the image may be an independent exact duplicate (clone) of the storage blocks at the physical disk level (for example, a split mirror), or a virtual copy (for example, a snapshot). From perspective of a basic operating system, the complete physical disk containing a backup object is replicated. However, if a volume manager is used on a UNIX system, the whole volume/disk group containing a backup object is replicated. If partitions are used on a Windows system, the whole physical volume containing the selected partition is replicated. See also snapshot, snapshot creation, split mirror, and split mirror creation.

replica set (ZDB specific term)

A group of replicas, all created using the same backup specification. See also replica and replica set rotation.

replica set rotation (ZDB specific term)

The use of a replica set for regular backup production: Each time the same backup

specification requiring the use of a replica set is run, a new replica is created and added to the set, until the maximum number for the set is reached. After that, the oldest replica in the set is replaced, maintaining the maximum number of replicas in the set. See also replica and replica set.

restore chain

Backup images that are needed to restore a backed up object to the state it was in at the selected point in time. In general, a restore chain of an object consists of its full backup image and one or more related incremental backup images.

restore session

A process that copies data from backup media to a client.

resync mode (HP P9000 XP Disk Array Family VSS provider specific term)

One of two P9000 XP Array VSS hardware provider operation modes. When the P9000 XP Array provider is in the resync mode, the source volume (P-VOL) and its replica (S-VOL) are in the suspended mirror relationship after a backup. The maximum number of replicas (S-VOLs per a P-VOL) rotated is three provided that MU range is 0-2 or 0, 1, 2. Restore from a backup in such a configuration is possible only by re-synchronization of an S-VOL with its P-VOL. See also VSS compliant mode, source volume, primary volume (P-VOL), replica, secondary volume (S-VOL), MU number, and replica set rotation.

RMAN (Oracle specific term)

See Recovery Manager.

RSM

The Data Protector Restore Session Manager controls restore and object

verification sessions. 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.

S

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 P9000 XP Disk Array Family specific term)

An internal disk (LDEV) of a disk array of the HP P9000 XP Disk Array Family which is paired with another LDEV: a primary volume (P-VOL). It can act as a mirror of the P-VOL or as a volume to be used for the P-VOL's snapshot storage. An S-VOL is assigned a SCSI address different from the one used for the P-VOL. In an HP CA P9000 XP configuration, the S-VOLs acting as mirrors can be used as failover devices in a MetroCluster configuration. See also primary volume (P-VOL) and Main Control Unit (MCU).

session

See backup session, media management session, and restore session.

session ID

An identifier of a backup, restore, object copy, object consolidation, object verification, 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 (Microsoft 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 and replica.

shadow copy provider (Microsoft 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 (Microsoft VSS specific term)

A collection of shadow copies created at the same point in time. See also shadow copy and replica set.

shared disks

A Windows disk on another system that has been made available to other users on the network. Systems with shared disks can be backed up without a Data Protector Disk Agent installed.

Site Replication Service (Microsoft Exchange Server specific term)

The Microsoft Exchange Server service that permits compatibility with Microsoft Exchange Server 5.5 by emulating the Exchange Server 5.5 directory service. See also Information Store and Key Management Service.

slot

A mechanical position in a library. Each slot can hold a medium, such as a DLT tape. Data Protector references each slot with a number. To read a medium, a robotic mechanism moves the medium from a slot into the drive.

SMB

See split mirror backup.

SMBF

The Session Messages Binary Files (SMBF) part of the IDB stores session messages generated during backup, restore, object copy, object consolidation, object verification, and media management sessions. One binary file is created per session. The files are grouped by year and month.

SMI-S Agent (SMISA)

See HP P6000 / HP 3PAR SMI-S Agent.

snapshot (HP P4000 SAN Solutions, HP P6000 EVA Disk Array Family, HP P9000 XP Disk Array Family, and HP 3PAR StoreServ Storage specific term)

A type of target volumes created using a specific replication technology.

Depending on the disk array model and the chosen replication technique, a range of snapshot types with different characteristics is available. Basically, each snapshot may be either a virtual copy, still reliant upon the contents of the source volume, or an independent duplicate (clone) of the source volume. See also replica and snapshot creation.

snapshot backup

See ZDB to tape, ZDB to disk, and ZDB to disk+tape.

snapshot creation (HP P4000 SAN Solutions, HP P6000 EVA Disk Array Family, HP P9000 XP Disk Array Family, and HP 3PAR StoreServ Storage specific term)

A replica creation process in which copies of the selected source volumes are created using storage virtualization technology. Such a replica is considered to be created at a particular point in time, and is immediately available for use. However, with certain snapshot types, a background data copying process continues to run on the disk array after the moment of the replica creation. 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 P9000 XP Disk Array Family specific term)

A type of target volumes created using a specific replication technology. A split-mirror replica provides independent duplicates (clones) 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 P9000 XP Disk Array Family specific term)

See ZDB to tape, ZDB to disk, and ZDB to disk+tape.

split mirror creation (EMC Symmetrix and HP P9000 XP Disk Array Family 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 P9000 XP Disk Array Family specific term)

A process in which data backed up in a ZDB-to-tape or a ZDB-to-disk+tape session is first copied from the backup media to a replica, and from the replica to the source volumes afterwards.

Individual backup objects or complete sessions can be restored using this method. See also ZDB to tape, ZDB to disk+tape, and replica.

sqlhosts file or registry (Informix Server specific term)

An Informix Server connectivity information file (on UNIX systems) or registry (on Windows systems) 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

A text file in the Unicode (UTF-16) format, generated during CONFIGURATION backup of a Windows system and stored on the Cell Manager. It contains system information required for installing and configuring the operating system on the target system in the case of a disaster. See also target system.

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 (SSEA)

See HP P9000 XP Agent.

sst.conf file

The file `/usr/kernel/drv/sst.conf` is required on each Data Protector Solaris client to which a multi-drive library device is connected. It must contain an entry for the SCSI address of the robotic mechanism of each library device connected to the client.

st.conf file

The file `/kernel/drv/st.conf` is required on each Data Protector Solaris client with a backup device connected. It must contain device information and a SCSI address for each backup drive connected to that client. A single SCSI entry is required for a single-drive device and multiple SCSI entries are required for a multi-drive library device.

stackers

Devices with multiple slots for media storage usually with only one drive. A stacker selects media from the stack sequentially. In contrast, a library can randomly select media from its repository.

standalone file device

A file device is a file in a specified directory to which you back up data.

Storage Group (Microsoft Exchange Server specific term)

A collection of mailbox stores and public folder stores that share a set of transaction log files. Exchange Server manages each storage group with a separate server process.

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.

Typically, these can be created or exist within a storage system such as a disk array.

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.

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. The 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.

SYMA

See EMC Symmetrix Agent.

synthetic backup

A backup solution that produces a synthetic full backup, an equivalent to a conventional full backup in terms of data, without putting stress on the production servers or the network. A synthetic full backup is created from a previous full backup and any number of incremental backups.

synthetic full backup

The result of an object consolidation operation, where a restore chain of a backup objects is merged into a new, synthetic full version of this object. A synthetic full backup is equivalent to a conventional full backup in terms of restore speed.

System Backup to Tape (SBT) (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 (sybssystemprocs) -model database (model).

System Recovery Data file

See SRD file.

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 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.

T

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 (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 database (Oracle specific term)

In RMAN, the target database is the database that you are backing up or restoring.

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 faulty 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 (Microsoft SQL Server specific term)

An executable entity that belongs to only one process. It comprises a program counter, a user-mode stack, a kernel-mode stack, and a set of register values. Several threads can run at a time within one process.

TimeFinder (EMC Symmetrix specific term)

A business continuation process that creates an instant copy of single or multiple EMC 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

See 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 log table (Sybase specific term)

A system table in which all changes to the database are automatically recorded.

transportable snapshot (Microsoft 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).

U**unattended operation**

See lights-out operation.

user account (Data Protector user account)

You can use Data Protector only if you have a Data Protector user account, which restricts unauthorized access to Data Protector and to backed up data. Data Protector administrators create this account specifying a user logon name, the systems from which the user can log on, and a Data Protector user group membership. This is checked whenever the user starts the Data Protector user interface or performs specific tasks.

User Account Control (UAC)

A security component in Windows Vista and Windows Server 2008 that limits application software to standard user privileges until an administrator authorizes an increase in privilege level.

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.

user_restrictions file

A file that restricts specific user actions, which are available to Data Protector user groups according to the user rights assigned to them, to be performed only on specific systems of the Data Protector cell. Such restrictions apply only to Data Protector user groups other than admin and operator.

V**vaulting media**

The process of storing media to a safe and remote place. The media are brought back to the data center when they are needed for restore or are ready for reuse in subsequent backups. The vaulting procedure depends on your company's backup strategy and policies for data protection/reliability.

verify

A function that lets you check whether the Data Protector data on a specified medium is readable. Additionally, consistency within each block can be checked if the backup was performed with the cyclic redundancy check (CRC) option ON.

Virtual Controller Software (VCS) (HP P6000 EVA Disk Array Family specific term)

The firmware that manages all aspects of storage system operation, including communication with HP Command View EVA through the HSV controllers. See also HP Command View (CV) EVA.

Virtual Device Interface (Microsoft SQL Server specific term)

This is a Microsoft SQL Server programming interface that allows fast backup and restore of large databases.

virtual disk (HP P6000 EVA Disk Array Family specific term)

A unit of storage allocated from a storage pool of a disk array of the HP P6000 EVA Disk Array Family. A virtual disk is the entity that can be replicated using the snapshot functionality of such disk array. See also source volume and target volume.

virtual full backup

An efficient type of synthetic backup where data is consolidated using pointers instead of being copied. It is performed if all the backups (the full backup, incremental backups, and the resulting virtual full backup) are written to a single file library that uses distributed file medium format.

Virtual Library System (VLS)

A disk-based data storage device hosting one or more virtual tape libraries (VTLs).

virtual server

A virtual machine in a cluster environment defined in a domain by a network IP name and address. Its address is cached by the cluster software and mapped to the cluster node that is currently running the virtual server resources. This way all requests for a particular virtual server are cached by a specific cluster node.

virtual tape

An archival storage technology that backs up data to disk drives in the same way as if it were being stored on tape. Benefits of virtual tape systems include improved backup and recovery speed and lower operating costs. See also Virtual Library System (VLS) and virtual tape library (VTL).

virtual tape library (VTL)

Data storage virtualization software that is able to emulate tape devices and libraries thus providing the functionality of traditional tape-based storage. See also Virtual Library System (VLS).

volser

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 (VSS).

VSS

See Microsoft Volume Shadow Copy Service (VSS).

VSS compliant mode (HP P9000 XP Disk Array Family VSS provider specific term)

One of two P9000 XP Array VSS hardware provider operation modes. When the P9000 XP Array provider is in the VSS compliant mode, the source volume (P-VOL) and its replica (S-VOL) are in simplex, unpaired state after a backup. Therefore the number of replicas (S-VOLs per a P-VOL) rotated is not limited. Restore from a backup in such a configuration is possible only by switching of the disks. See also resync mode, source volume, primary volume (P-VOL), replica, secondary volume (S-VOL), and replica set rotation.

VxFS

Veritas Journal Filesystem.

VxVM (Veritas Volume Manager)

A Veritas Volume Manager is a system for managing disk space on Solaris platforms. A VxVM system consists of an arbitrary group of one or more physical volumes organized into logical disk groups.

W

Wake ONLAN

Remote power-up support for systems running in power-save mode from some other system on the same LAN.

Web reporting

The Data Protector functionality that allows you to view reports on backup, object copy, and object consolidation status and Data Protector configuration using the Web interface.

wildcard character

A keyboard character that can be used to represent one or many characters. The asterisk (*), for example, typically represents one or more characters, and the question mark (?) typically represents a single character. Wildcard characters are often used in operating systems as a means of specifying more than one file by name.

Windows configuration backup

Data Protector allows you to back up Windows CONFIGURATION, including Windows Registry, user profiles, Event Logs, and WINS and DHCP server data (if configured on a system) in one step.

Windows Registry

A centralized database used by Windows to store configuration information for the operating system and the installed applications.

WINS server

A system running Windows Internet Name Service software that resolves Windows networking computer names to IP addresses. Data Protector can back up WINS server data as part of the Windows configuration.

writer (Microsoft VSS specific term)

A process that initiates change of data on the original volume. Writers are typically applications or system services that write persistent information on a volume. Writers also participate in the shadow copy synchronization process by assuring data consistency.

X**XBSA Interface (Informix Server specific term)**

ON-Bar and Data Protector communicate with each other through the X/Open Backup Services Application Programmer's Interface (XBSA).

Z**ZDB**

See zero downtime backup.

ZDB 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 with specific disk array families, split mirror restore. See also zero downtime backup (ZDB), ZDB to disk, ZDB to tape, instant recovery, replica, and replica set rotation.

ZDB to tape (ZDB specific term)

A form of zero downtime backup where data in the replica produced is streamed to a backup medium, typically tape. Instant recovery is not possible from such a backup, so the replica need not be retained on the disk array after backup completion. The backed up data can be restored using standard Data Protector restore from tape. With specific disk array families, split mirror restore can also be used. See also zero downtime backup (ZDB), ZDB to disk, ZDB to disk+tape, instant recovery, and replica.

zero downtime backup (ZDB)

A backup approach in which data replication techniques, provided by a disk array, are used to minimize the impact of backup operations on an application system. A replica of the data to be backed up is created first. All subsequent backup operations are performed on the replicated data rather than the original data, while the application system can be returned to normal operation. See also ZDB to disk, ZDB to tape, ZDB to disk+tape, and instant recovery.

Index

A

architecture

SAP HANA Appliance integration 211

SAP MaxDB integration 175

SAP R/3 integration 114

B

backing up Oracle 51

backup options 44

backup specifications, creating 38

backup templates 37

backup types 11

examples, using RMAN 60

offline 51

online 51

recovery catalog 53

resuming backup sessions 96

scheduling backups 54

starting backups 55

starting backups, using CLI 56

starting backups, using GUI 55

starting backups, using RMAN 57

backing up SAP HANA Appliance 215

architecture 211

backup specification, modifying 218

backup specifications, creating 215

backup types 210

concepts, scheme 211

full backups 210

online backups 210

starting backup sessions 218

transaction log backup 210

backing up SAP MaxDB 182

architecture 175

backup flow 176

backup modes 182

backup options 184

backup specification, modifying 185

backup specifications, creating 182

backup types 174

concepts, scheme 175

differential backups 174

full backups 174

online backups 174

parallelism 185

parallelism, concepts 176

previewing backups 186

scheduling backups 185

scheduling backups, example 185

starting backups 187

transaction log backups 174

backing up SAP R/3 135

architecture 114

backup flow 116

backup modes 113

backup options 141

backup specification, modifying 142

backup specifications, creating 137

backup templates 138

backup types 112, 135

full backups 112, 135

- incremental backups 112, 135
- manual balancing 141, 147
- previewing backups 143
- SAP backup utilities 113
- SAP R/3 parameter file 136
- scheduling backups 142
- scheduling backups, example 143
- starting backups 144
- using RMAN 137, 146
- backint mode
 - SAP R/3 integration 113
- backup flow
 - SAP MaxDB integration 176
 - SAP R/3 integration 116
- backup flow, Oracle integration 14
- backup modes
 - SAP MaxDB integration 182
- backup modes, SAP R/3 integration 113
- backup options
 - Oracle integration 44
 - SAP MaxDB integration 184
 - SAP R/3 integration 141
- backup sessions, scheduling
 - Oracle integration 54
 - SAP MaxDB integration 185
 - SAP R/3 integration 142
- backup specifications, creating
 - Oracle integration 38
 - SAP HANA Appliance integration 215
 - SAP MaxDB integration 182
 - SAP R/3 integration 137
- backup specifications, modifying
 - SAP HANA Appliance integration 218
 - SAP MaxDB integration 185
 - SAP R/3 integration 142
- backup templates
 - Oracle integration 37
 - SAP R/3 integration 138
- backup types
 - Oracle integration 11
 - SAP HANA Appliance integration 210
 - SAP MaxDB integration 174
 - SAP R/3 integration 112, 135
- BRARCHIVE
 - SAP R/3 integration 113
- BRBACKUP
 - SAP R/3 integration 113
- BRRESTORE 151
 - SAP R/3 integration 113
- C**
- checking configuration
 - Oracle integration 33
 - SAP MaxDB integration 181
 - SAP R/3 integration 134
- concepts
 - Oracle integration 12
 - SAP HANA Appliance integration 211
 - SAP MaxDB integration 175
 - SAP R/3 integration 113
- configuration file
 - SAP R/3 integration 118

- configuring Oracle 16
 - checking configuration 33
 - example, CLI 28
 - prerequisites 18
- configuring SAP HANA Appliance 213
- configuring SAP MaxDB 176
 - checking configuration 181
- configuring SAP R/3 118
 - authentication modes 128
 - checking configuration 134
 - configuration file 118
- control files, Oracle integration
 - restore 68
- creating backup specifications
 - Oracle integration 38
 - SAP HANA Appliance integration 215
 - SAP MaxDB integration 182
 - SAP R/3 integration 137
- D**
- Data Guard, Oracle integration
 - configuration, example 28
 - limitations 18
 - primary databases, restore 75
 - standby databases, restore 75
- database recovery
 - Oracle integration, options 78
- differential backups
 - SAP MaxDB integration 174
- disaster recovery
 - Oracle integration 65, 93
 - SAP R/3 integration 153

- E**
- examples
 - SAP R/3 integration, starting interactive backups 145
- examples, Oracle integration
 - backing up using RMAN 60
 - restoring using RMAN 82
- examples, SAP MaxDB integration
 - scheduling backups 185
 - starting interactive backups 188
- examples, scheduling backups
 - SAP R/3 integration 143
- F**
- full backup
 - SAP HANA Appliance integration 210
- full backups
 - SAP MaxDB integration 174
 - SAP R/3 integration 112, 135
- I**
- incremental backups
 - Oracle integration 54
 - SAP R/3 integration 112, 135
- Informix backup
 - backup specifications, creating 137
- interactive backups
 - Oracle integration 55
 - SAP MaxDB integration 187
 - SAP R/3 integration 144
- introduction
 - Oracle integration 11
 - SAP HANA Appliance integration 210

- SAP MaxDB integration 174
- SAP R/3 integration 112
- L**
- limitations
 - SAP HANA Appliance integration 212
 - SAP MaxDB integration 177
- M**
- manual balancing
 - SAP R/3 integration 141, 147
- Media Management Library 19
- migration, restore
 - SAP MaxDB integration 194
- MML (Data Protector Media Management Library)
 - linking with Oracle, HP OpenVMS 19
 - linking with Oracle, UNIX 19
- modifying backup specifications
 - SAP HANA Appliance integration 218
 - SAP MaxDB integration 185
 - SAP R/3 integration 142
- monitoring sessions
 - Oracle integration 94
 - SAP HANA Appliance integration 219
 - SAP MaxDB integration 206
 - SAP R/3 integration 154
- O**
- OB2RMANSAVE, Oracle integration 110
- omniintconfig.pl
 - Oracle integration 29
- online backups
 - SAP HANA Appliance integration 210
 - SAP MaxDB integration 174
- Oracle backup 51
 - backup concepts, scheme 15
 - backup specifications, creating 38
 - backup templates 37
 - backup types 11
 - resuming backup sessions 96
 - scheduling backups 54
 - starting backups 55
 - starting backups, using CLI 56
 - starting backups, using GUI 55
 - starting backups, using RMAN 57
- Oracle configuration
 - checking configuration 33
 - example, CLI 28
 - prerequisites 18
- Oracle integration
 - backup 51
 - concepts 12
 - configuration 16
 - disaster recovery 93
 - introduction 11
 - monitoring sessions 94
 - restore 63
 - resuming sessions 96
 - troubleshooting 100
 - viewing sessions 95
- Oracle restore 63
 - control files 68
 - database items 63
 - database objects 69

- disaster recovery 93
- editing RMAN scripts 110
- examples, using RMAN 82
- preparing databases for restore 83
- primary databases, Data Guard 75
- recovery catalog 66
- restorable items 63
- restore flow 14
- restore methods 63
- restore options 78
- restore types 11
- resuming restore sessions 97
- standard restore procedure 63
- standby databases, Data Guard 75
- tablespaces and datafiles 74
- using another device 93
- using GUI 65
- using RMAN 82
- Oracle RMAN metadata 99
- Oracle RMAN script 45
- Oracle troubleshooting 100
- overview, restore
 - SAP MaxDB integration 191
- P**
- parallelism
 - SAP MaxDB integration 185
- parallelism, concepts
 - SAP MaxDB integration 176
- previewing backups
 - SAP MaxDB integration 186
 - SAP R/3 integration 143
- primary databases, Oracle integration
 - restore 75
- R**
- RAC, configuring Oracle Servers
 - on HP-UX 19
 - on other UNIX systems 19
- recovery
 - Oracle integration, options 78
- recovery catalog, Oracle integration
 - backup 53
 - restore 66
- restore flow
 - SAP MaxDB integration 176
 - SAP R/3 integration 117
- restore methods
 - SAP R/3 integration 148
- restore options
 - SAP MaxDB integration 202
- restore types
 - Oracle integration 11
- restoring Oracle 63
 - control files 68
 - database objects 69
 - disaster recovery 93
 - editing RMAN scripts 110
 - methods 63
 - primary databases, Data Guard 75
 - recovery catalog 66
 - restore flow 14
 - resuming restore sessions 97
 - standard restore procedure 63

- standby databases, Data Guard 75
- tablespaces and datafiles 74
- using another device 93
- using GUI 65
- using RMAN 82
- restoring SAP HANA Appliance 218
 - starting restore sessions 219
- restoring SAP MaxDB 191
 - migration 194
 - overview 191
 - parallelism, concepts 176
 - restore flow 176
 - restore options 202
 - using another device 206
 - using CLI 197
 - using GUI 194
 - using SAP MaxDB utilities 198
- restoring SAP R/3 148
 - architecture 114
 - disaster recovery 153
 - restore flow 117
 - restore methods 148
 - SAP restore utilities 113
 - standard restore 148
 - using another device 152
 - using BRRESTORE 151
 - using CLI 150
 - using GUI 148
 - using SAP BRTOOLS 151
- RMAN mode
 - SAP R/3 integration 113
 - RMAN, backup
 - SAP R/3 integration 137, 146
 - RMAN, Oracle integration 57
 - backup 60
 - restore 82
 - scripts, examples 60
 - running backups 144
- S**
 - SAP HANA Appliance backup 215
 - architecture 211
 - backup specification, modifying 218
 - backup specifications, creating 215
 - backup types 210
 - concepts, scheme 211
 - full backups 210
 - online backups 210
 - starting backup sessions 218
 - transaction log backup 210
 - SAP HANA Appliance configuration 213
 - SAP HANA Appliance integration
 - backup 215
 - concepts 211
 - configuration 213
 - introduction 210
 - limitations 212
 - monitoring sessions 219
 - restore 218
 - troubleshooting 219
 - SAP HANA Appliance restore 218
 - starting restore sessions 219
 - SAP HANA Appliance troubleshooting 219

- SAP MaxDB backup 182
 - architecture 175
 - backup flow 176
 - backup modes 182
 - backup options 184
 - backup specification, modifying 185
 - backup specifications, creating 182
 - backup types 174
 - concepts, scheme 175
 - differential backups 174
 - full backups 174
 - online backups 174
 - parallelism 185
 - parallelism, concepts 176
 - previewing backups 186
 - scheduling backups 185
 - scheduling backups, example 185
 - starting backups 187
 - transaction log backups 174
- SAP MaxDB configuration 176
 - checking configuration 181
- SAP MaxDB integration
 - backup 182
 - concepts 175
 - configuration 176
 - introduction 174
 - limitations 177
 - monitoring sessions 206
 - restore 191
 - troubleshooting 206
- SAP MaxDB restore 191
 - migration 194
 - overview 191
 - parallelism, concepts 176
 - restore flow 176
 - restore options 202
 - using another device 206
 - using CLI 197
 - using GUI 194
 - using SAP MaxDB utilities 198
- SAP MaxDB troubleshooting 206
- SAP MaxDB utilities
 - restore 198
- SAP R/3 backup 135
 - architecture 114
 - backup flow 116
 - backup modes 113
 - backup options 141
 - backup specification, modifying 142
 - backup templates 138
 - backup types 112, 135
 - full backups 112, 135
 - incremental backups 112, 135
 - manual balancing 141, 147
 - previewing backups 143
 - SAP backup utilities 113
 - SAP R/3 parameter file 136
 - scheduling backups 142
 - scheduling backups, example 143
 - starting backups 144
 - using RMAN 137, 146

- SAP R/3 configuration 118
 - authentication modes 128
 - checking configuration 134
 - configuration file 118
- SAP R/3 integration
 - backup 135
 - concepts 113
 - configuration 118
 - disaster recovery 153
 - introduction 112
 - monitoring sessions 154
 - restore 148
 - troubleshooting 154
- SAP R/3 restore 148
 - architecture 114
 - disaster recovery 153
 - restore flow 117
 - restore methods 148
 - SAP restore utilities 113
 - standard restore 148
 - using another device 152
 - using BRRESTORE 151
 - using CLI 150
 - using GUI 148
- SAP R/3 troubleshooting 154
 - on UNIX 164
 - on Windows 155
- SBT_LIBRARY, Oracle integration 19, 58, 85
- scheduling backups
 - Oracle integration 54
 - SAP MaxDB integration 185
 - SAP R/3 integration 142
- standby databases, Oracle integration
 - restore 75
- starting backup sessions
 - SAP HANA Appliance integration 218
- starting backups
 - SAP MaxDB integration 187
 - SAP R/3 integration 144
- starting backups, Oracle integration 55
 - using CLI 56
 - using GUI 55
 - using RMAN 57
- starting restore sessions
 - SAP HANA Appliance integration 219
- T**
- transaction log backup
 - SAP HANA Appliance integration 210
- transaction log backups
 - SAP MaxDB integration 174
- troubleshooting Oracle 100
- troubleshooting SAP HANA Appliance 219
- troubleshooting SAP MaxDB 206
- troubleshooting SAP R/3 154
 - on UNIX 164
 - on Windows 155
- U**
- users, configuring
 - Oracle integration 20

V

viewing sessions

Oracle integration 95

We appreciate your feedback!

If you have comments about this document, you can [contact the documentation team](#) by email. If an email client is configured on this system, click the link above and an email window opens with the following information in the subject line:

Feedback on Integration guide for Oracle and SAP (Data Protector 8.10)

Just add your feedback to the email and click send.

If no email client is available, copy the information above to a new message in a web mail client, and send your feedback to AutonomyTPFeedback@hpe.com.

